

Individual Fairness in Advertising Auctions Through Inverse Proportionality

Shuchi Chawla ✉

The University of Texas at Austin, TX, USA

Meena Jagadeesan ✉

University of California, Berkeley, CA, USA

Abstract

Recent empirical work demonstrates that online advertisement can exhibit bias in the delivery of ads across users even when all advertisers bid in a non-discriminatory manner. We study the design ad auctions that, given fair bids, are guaranteed to produce fair outcomes. Following the works of Dwork and Ilvento [11] and Chawla et al. [7], 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 maximization and fairness.

This work makes two conceptual contributions. First, 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 ℓ_∞ distance between the corresponding allocations. Standard auctions do not satisfy this kind of value stability.

Second, we introduce a new class of allocation algorithms called Inverse Proportional Allocation that achieve a near optimal tradeoff between fairness and social welfare for a broad and expressive class of value stability conditions. These allocation algorithms are truthful and prior-free, and achieve a constant factor approximation to the optimal (unconstrained) social welfare. In particular, the approximation ratio is independent of the number of advertisers in the system. In this respect, these allocation algorithms greatly surpass the guarantees achieved in previous work. We also extend our results to broader notions of fairness that we call subset fairness.

2012 ACM Subject Classification Theory of computation → Algorithmic mechanism design

Keywords and phrases Algorithmic fairness, advertising auctions

Digital Object Identifier 10.4230/LIPIcs.ITCS.2022.42

Related Version *Full Version*: <https://arxiv.org/abs/2003.13966>

Funding *Shuchi Chawla*: Supported in part by NSF awards CCF-2008006 and SHF-1704117.

Meena Jagadeesan: Supported by a Paul and Daisy Soros Fellowship.

1 Introduction

Algorithms play an increasingly important role in today’s society and can arguably have far-ranging social, economic, and political ramifications in many different contexts. One such context is access to sponsored information such as advertisements on social media, search pages, and other websites. It has been well documented that online advertisement can exhibit skewed delivery: users that differ on sensitive attributes such as race, gender, age, religion, and national origin can receive very different allocations of ads (see, e.g., [20, 9, 2, 3, 17, 16, 1]). For example, certain employment jobs on Facebook have been found to exclusively target men or exclude older people; and housing ads have been found to exclude users based on race [3, 2, 16].



© Shuchi Chawla and Meena Jagadeesan;

licensed under Creative Commons License CC-BY 4.0

13th Innovations in Theoretical Computer Science Conference (ITCS 2022).

Editor: Mark Braverman; Article No. 42; pp. 42:1–42:21

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

There are two main sources of unfairness in digital ads. The first is explicit or implicit targeting of users based on sensitive attributes by the advertisers. This source of unfairness can be mitigated by examining and auditing each advertiser’s targeting approach separately.¹ The second, more subtle and insidious, 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. Recent empirical studies [17, 1] have observed that even when each advertiser’s targeting parameters are inclusive, skewed delivery can arise as a result of the ad auction mechanism. Our paper focuses on eliminating this second source of unfairness.

We study the design of ad auctions from the perspective of individual fairness. Individual fairness [10] posits that similar users should be treated similarly. In the context of sponsored content, this translates into requiring that similar users should be served a similar mix of ads. We follow the approach of Chawla et al. [7] and isolate the effect of *auction design* on fairness in outcomes by assuming that each advertiser’s bids are non-discriminatory. Then, if two users receive similar bids from all advertisers, we require that they receive similar ad allocations. In other words, the auction does not introduce any further unfairness than what may be present already in advertisers’ bids.

We consider a stylized model of (truthful) ad auctions where users arrive over time and each is shown a single ad. When a user arrives, k advertisers report their values for the single ad slot; denote these values with the vector $\mathbf{v} = (v_1, v_2, \dots, v_k)$.² The auction selects a (potentially random) advertiser based on the values and displays its ad; we denote the auction’s selection by a distribution over $[k]$, $\mathbf{a}(\mathbf{v}) \in \Delta([k])$, where $a_i(\mathbf{v})$ is the probability that advertiser i ’s bid is displayed. We measure the quality of this allocation in terms of its social welfare – the expected sum of values corresponding to ads displayed, with the contribution of each user to this sum given by $\mathbf{v} \cdot \mathbf{a}(\mathbf{v}) = \sum_{i \in [k]} v_i a_i(\mathbf{v})$.

Social welfare maximization for our simple model in the absence of a fairness constraint is easy: just serve each user the ad with the highest value. Real-world ad auctions are essentially variants of this “highest-bid-wins” auction, modified to accommodate click-through-rates, multiple slots, etc. Our main challenge is that these standard auction formats are fundamentally incompatible with fairness. Consider, for example, two job ads A and B competing over two users Alice and Bob. Suppose that Alice and Bob have similar educational qualifications and work experience but belong to different ethnic groups. Further, due to slight differences in their profiles, suppose that A values Alice slightly higher than B and B values Bob slightly higher than A. Then a highest-bid-wins auction would allocate ad A to Alice and ad B to Bob. This violates our fairness constraint: although the two users are similar to each other in most respects, their ad allocations are completely different. Indeed, if job B is much more desirable than job A, then Alice is discriminated against. Note that the unfairness persists even if each advertiser individually places similar values on the two users. As we show in experiments in the full version of the paper, this behavior is not just hypothetical but is frequently observed in real data as well. The highest-bid-wins auction greatly exaggerates minor, subtle differences in input into huge swings in output.

¹ Indeed, in response to a lawsuit brought on by the U.S. Department of Housing and Urban Development and ProPublica against Facebook on discrimination in housing ads, Facebook’s first action was to disallow advertisers to target users based on sensitive attributes [16]. However, despite this action, unfairness in outcomes continued to persist [1].

² We focus in this work on truthful auctions, and so do not make a distinction between bids and values. All of the allocation rules we propose and study are monotone in values and can be implemented in a truthful manner.

1.1 Our contributions

In this paper we propose a new framework for expressing fairness in ad auctions as a value-stability condition. We then design allocation algorithms that achieve near-optimal tradeoffs between value-stability and social welfare.

Fairness as a stability condition. Consider, again, two users Alice and Bob. Let \mathbf{v} denote the values received by Alice and \mathbf{v}' the values received by Bob. We will use a parameter $\lambda \geq 1$ to quantify the similarity between Alice and Bob, and will say that they receive non-discriminatory values if it holds that $v_i' \in [\frac{1}{\lambda}v_i, \lambda v_i]$ for all $i \in [k]$. Given such non-discriminatory values, we will require that Alice and Bob receive allocations that are component-wise close to each other, i.e. $a_i(\mathbf{v}') \in a_i(\mathbf{v}) \pm f(\lambda)$ for all $i \in [k]$.³ Here f is an externally specified function that we call the fairness constraint and that governs the strength of the fairness guarantee. We say that the auction is *value stable* with respect to the constraint f .

We essentially require the allocation function to satisfy a kind of stability condition: when it receives two value vectors that are component-wise close, it should return allocations that are component-wise close. Observe that we measure closeness in values as a ratio, and consequently our stability notion is scale free. On the other hand, we measure closeness in allocations as a difference, specifically the ℓ_∞ distance between the respective distributions. This corresponds to the notion of *multiple-task fairness* defined by Dwork and Ilvento [11] in the context of multi-dimensional allocation algorithms.

Our approach of expressing fairness as a stability constraint offers several benefits. First, our approach assumes that every pair of similar users receives multiplicatively similar values from *all* advertisers. This enables a clear separation of responsibility between the auctioneer and the advertisers, which can simplify the design of both of these components of the ad auction pipeline. We envision that an auditor would check whether advertisers are following fairness guidelines. As a result, the auctioneer only needs to ensure that the allocation is fair when the advertisers' reported values are fair, but is required to provide no fairness guarantee when the values themselves are unfair. Second, and more importantly, the auctioneer can provide this guarantee without access to the specifics of fairness requirements across users. One criticism of the notion of individual fairness is its reliance on an appropriate similarity metric on users that captures exactly what features can be used to differentiate between users. This metric can not only be challenging to learn, but may itself be sensitive information that is unavailable to the auctioneer or illegal to use directly. Our approach allows the auctioneer to provide a meaningful fairness guarantee without knowing or using the similarity metric. In effect, the metric is encoded in the advertisers' reported values when those values are fair.

Finally, we note that in [7], we considered a similar model of fairness in ad auctions but formalized the fairness constraint in a slightly different manner. The algorithmic results in this paper are directly comparable to those in [7], and we make the connection explicit in Section 2.1.

Designing value-stable allocation algorithms. Now that we have defined value-stability, our goal can be restated as designing value stable auctions that achieve high social welfare.

In [7], we proposed a natural candidate for value stability, namely Proportional Allocation (PA) algorithm: given the value vector \mathbf{v} , we assign each ad i to the user with probability $v_i / \sum_j v_j$. It is straightforward to observe that changing each value by a factor of λ changes

³ We emphasize that we do not require that two ads with similar values for the same user be served with similar probabilities: we aim to satisfy fairness across users, not across advertisers.

each allocation multiplicatively by at most λ^2 . In fact, PA actually achieves the additive form of fairness defined above with $f(\lambda) = (\lambda - 1)/(\lambda + 1)$.⁴ Unfortunately, the social welfare of PA degrades as the number of advertisers increases. Suppose, for example, that there are $k + 1$ advertisers; the first one places a value of 1 on the user whereas all the others value the user at $x < 1$. Then, PA assigns the first ad to the user with probability $1/(1 + xk)$ and obtains a social welfare of $(1 + x^2k)/(1 + xk)$. When $x = 1/\sqrt{k}$, this is a factor of \sqrt{k} off from the optimal social welfare of 1. Other proportional allocation algorithms that assign allocations in proportion to other functions of values⁵ suffer from the same problem: their worst case approximation ratio goes to zero as the number k of advertisers goes to infinity. Is it possible to obtain better social welfare while maintaining fairness?

In this paper we introduce a new family of allocation rules that we call **Inverse Proportional Allocation** (IPA). 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 strength of the desired fairness constraint. This new family of mechanisms 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 k . We show, in fact, that Inverse Proportional Allocation achieves a near-optimal tradeoff between fairness and social welfare for *any* stability constraint satisfying a mild assumption.

We describe the fairness guarantee of our algorithm in the form of a “fairness constraint” f that specifies how different the allocations for two users can be as a function of the similarity between the users. We use a parameter $\ell \in (0, \infty)$ to denote the strength of this fairness constraint: the smaller that ℓ is, the stronger is the fairness guarantee. $\ell = 0$ corresponds to perfect fairness or the function that is 0 at all distances: all users receive identical allocations; and $\ell = \infty$ corresponds to unconstrained allocations or the function that is 1 at all distances: users can receive arbitrary allocations. For intermediate ℓ , we consider a polynomial family f_ℓ of fairness constraints that best captures the fairness profile of IPA (parameterized appropriately).

► **Theorem 1 (Informal).** *For any $\ell \in (0, \infty)$, Inverse Proportional Allocation with parameter ℓ is value stable with respect to $f_\ell(\lambda) = 1 - \lambda^{-2\ell}$. Its worst case approximation ratio for social welfare, α_ℓ , is $3/4$ at $\ell = 1$ and approaches 1 as $\ell \rightarrow \infty$. No prior free allocation algorithm that is value stable with respect to f_ℓ can achieve a better approximation ratio than $\alpha_\ell + 1/k$.*

When we further expand the class of allocation algorithms to include IPA mixed in with the uniform allocation, we achieve near-optimality against a broad class of constraints f .

► **Theorem 2 (Informal).** *Let $f : [1, \infty) \rightarrow [0, 1]$ be any non-decreasing function such that $f(x)/\log x$ is non-increasing. Suppose there exists an allocation algorithm that is value stable with respect to f and achieves a worst case approximation factor of α . Then there is an IPA algorithm that is value stable with respect to f and achieves a worst case approximation factor of $\Omega(\alpha/(1 + \log(1/\alpha)))$.*

Empirical evaluation. In addition to theoretical analysis, we perform experiments on sponsored search auction data from Yahoo [21] to corroborate our theoretical bounds with respect to both the social welfare and fairness. One practical challenge for our approach,

⁴ The strength of the fairness guarantee can be changed by allocating each ad with probability proportional to an appropriate function of the corresponding value.

⁵ This class includes, for example, the exponential algorithm from differential privacy literature [18].

which relies on values to infer similarity between users, is that not all advertisers bid on all users: some advertisers may run out of budget, and some choose to bid on only a subpopulation. Interpreting missing bids as zero values can cause the algorithm to form an incorrect estimate of the similarity between two users, potentially violating fairness. For our experiments, we define two users as similar as long as there is significant overlap between the sets of advertisers that bid on both (in terms of Jaccard similarity) and the values of those common advertisers are multiplicatively close.

Our main observations are twofold. On the one hand, the highest-bid-wins auction is far from value stable, exhibiting widely varying allocations even for pairs of users with value ratios very close to 1. On the other hand, IPA remains fair in the face of missing bids and, as predicted by the theoretical bounds, the degree of value stability weakens as ℓ increases. Relative to the optimal social welfare, IPA obtains approximation ratios of over 90% for a large range of parameter values and up to 98% in some cases. IPA therefore offers a vast improvement over highest-bid-wins in terms of fairness while suffering a minimal loss in performance. We defer these empirical results to the full version of the paper.

Comparison between IPA and PA. While our primary focus is on comparing the IPA with the highest-bid-wins auction, a.k.a. the unfair optimum, our results provide insight into a comparison between IPA and PA as well. In order to obtain an accurate picture, it is important to consider parameterizations of both algorithms that satisfy the same fairness constraint. In our theoretical analysis, we use the fact that for any fixed ℓ , the fairness constraint f_ℓ is satisfied by IPA_ℓ and $\text{PA}_{2\ell}$. When we take a limit as $k \rightarrow \infty$, we show that IPA_ℓ achieves a constant factor approximation for social welfare, while the worst case social welfare of PA_ℓ approaches 0. Moreover, IPA_ℓ achieves a better worst-case social welfare than $\text{PA}_{2\ell}$ as long as $k \geq 6$. In this sense, IPA greatly surpasses PA in terms of the fairness/social-welfare tradeoffs.

In our empirical evaluation, we must consider the performance of these algorithms for small constant values of k . This is tricky because the theoretical analyses of fairness for both algorithms are loose for smaller values of k . To resolve this, we determine a parameterization for PA by empirically matching its fairness profile with that of IPA_ℓ as closely as possible. For the Yahoo dataset [21] that we use, a close match occurs between IPA_ℓ and $\text{PA}_{\frac{4}{3}\ell}$. We then compare the performance of the two algorithms and find that IPA consistently performs better on social welfare (with the exception of when the number of advertisers in the auction is 2). In fact, as the number of advertisers increases, the performance of IPA remains unchanged, while that of PA degrades rapidly. We conclude that in most settings, IPA provides a better tradeoff between fairness and social welfare.

1.2 Further extensions

We also show our results can be extended to the following two settings.

Fairness across different categories of ads. So far in our discussion we have assumed that there is a single notion of similarity across users that all advertisers are required to respect. In reality, platforms often service many different categories of advertisers (for example job advertisers and credit advertisers) that may be subject to different notions of similarity over the users. These advertisers all compete against each other for the same users. Chawla et al. [7] proposed fairness definitions for the multi-category setting that combine individual fairness with envy-freeness. Our algorithms from this work can be utilized within the multi-category framework in [7] to achieve significantly improved tradeoffs between fairness and social welfare for the multi-category setting. We discuss this extension in the full version of the paper.

Total-variation and subset fairness. In [7], we showed that PA satisfies a stronger fairness guarantee called *total variation fairness*: the distributions over ads assigned to similar users are not only close under ℓ_∞ distance, but also close under ℓ_1 or total variation distance. Equivalently, total-variation fairness provides guarantees over any subset of ads, rather than just a single ad. However, we observe that our family of IPA auctions does not always satisfy total variation fairness. Indeed it appears to be challenging to satisfy this stronger fairness property while also guaranteeing a constant factor approximation in social welfare. In the full version of the paper, we show, however, that IPA can be adapted to achieve an intermediate fairness guarantee that we call *subset fairness*: fairness is guaranteed over a restricted collection of subsets of advertisers.

1.3 Other related work

While algorithmic fairness is by now an established area of study, theoretical work on fairness in ad auctions is relatively new. There are two complementary sides to the fairness problem in the context of ad auctions. As in our work, Celis et al. [6] study the problem from the viewpoint of the platform – redesigning the ad auction so as to achieve fair outcomes. They provide an optimization framework within which, given an explicit population of users and bids over them, a platform can compute the optimal fair assignment of ads to users. Celis et al. focus on algorithmic techniques for solving the optimization problem, comparing the performance of their algorithm to the optimal fair solution; whereas our focus in this work is on comparing the performance of our algorithm against the unfair optimum and therefore characterizing the cost of fairness.

The complementary viewpoint is that of the advertisers – designing bidding and targeting strategies that preemptively correct for unfairness introduced by the platform mechanism. Nasr and Tschantz [19] design bidding strategies for advertisers aimed at obtaining parity in impressions across fixed categories (such as gender). Gelauff et al. [12] likewise design targeting strategies aimed at obtaining parity in outcomes or conversions across different demographic groups. A key point of difference between these works and ours is that they focus on the notion of group fairness whereas our work considers individual fairness guarantees.

Besides these works, the tradeoff between fairness and social welfare has also been studied in other contexts such as classification problems (see, e.g., Hu and Chen [13]).

Complementary to algorithmic fairness, there is an extensive literature on fair division that also considers the tension between fairness and social welfare. A major difference between the algorithmic fairness and envy-freeness is that algorithmic fairness generally focuses on individual qualifications, while envy-freeness generally focuses on individual preferences, though recent work [7, 15, 4, 22] has proposed definitions which combine aspects of envy-freeness and algorithmic fairness notions.

1.4 Outline for the rest of the paper

In Section 2, we present our model for fair ad auctions and describe the connection between value stability and individual fairness. In Section 3, we describe Inverse Proportional Allocation algorithms and state our main results on social welfare and value stability. In Section 4, we prove the value stability results. In Section 5 we consider extensions of the IPA mixed in with uniform allocation, and show that this family of allocation algorithms achieves near optimal tradeoffs between social welfare and value stability.

In the full version of the paper, we also provide an empirical evaluation of inverse proportional allocation on the Yahoo sponsored search auction dataset. Moreover, we consider the stronger notion of subset fairness and develop fair allocation algorithms with social welfare that degrade with the complexity of the subset collection.

2 Ad Auctions and Value Stability

We use a simple stylized model for ad auctions. Every time a user arrives on the platform, k advertisers bid for the chance to display their ad in the *single* ad slot available. We assume that the platform employs a truthful direct auction to determine which ad to display. In particular, when the user arrives, each advertiser i reports a value v_i for the user. The auction takes the value vector $\mathbf{v} = [v_1, \dots, v_k]$ as input and returns a distribution $\mathbf{a}(\mathbf{v}) = [a_1(\mathbf{v}), \dots, a_k(\mathbf{v})]$ over the advertisers. We drop the argument \mathbf{v} when it is clear from the context. With probability a_i , advertiser i is picked and the corresponding ad is displayed for the user. We call the function $\mathbf{a}(\mathbf{v})$ the allocation function.

An allocation function \mathbf{a} is truthful if there exists an accompanying payment function $\mathbf{p}(\mathbf{v})$ that incentivizes advertisers to bid their true values. Such a payment function exists if and only if for all i , $a_i(v_i, v_{-i})$ is a monotone non-decreasing function of the argument v_i ; payments can be computed using the standard payment identity. Henceforth we focus on weakly monotone allocation functions. We briefly discuss the computation of payments using the payment identity in the full version of the paper.

Value stability. In order to prevent the allocation function from introducing unfairness into the system, we require that our allocation is *value-stable*. If two users have similar values from all advertisers, then they must receive similar allocations. That is, we require that for two value vectors \mathbf{v} and \mathbf{v}' that are component-wise close in a multiplicative sense, the allocations $\mathbf{a}(\mathbf{v})$ and $\mathbf{a}(\mathbf{v}')$ must be close in an additive sense. Formally:

► **Definition 3.** An allocation is *value-stable* with respect to function $f : [1, \infty] \rightarrow [0, 1]$ if the following condition is satisfied for every pair of value vectors \mathbf{v} and \mathbf{v}' :

$$|a_i(\mathbf{v}) - a_i(\mathbf{v}')| \leq f(\lambda) \text{ for all } i \in [k], \text{ where } \lambda \text{ is defined as } \max_{i \in [k]} \left(\max \left\{ \frac{v_i}{v_i'}, \frac{v_i'}{v_i} \right\} \right).$$

The function f governs the strength of the value stability constraint. When two value vectors are identical, the allocation should be the same, so $f(1) = 0$. When two value vectors are arbitrarily different, the allocation should be permitted to be arbitrarily different, so $f(\infty) = 1$. For intermediate values, observe that the smaller f is, the tighter is the constraint on the allocation, and the harder it should be for the auction to achieve value stability. Following [7], we characterize the strength of the fairness constraint on values by considering an explicit family \mathcal{F} of functions parameterized by $\ell \in (0, \infty)$: $f_\ell(x) = 1 - x^{-2\ell}$.⁶

Simplicity through prior-free design. Observe that our notion of value stability is defined over all possible pairs of value vectors \mathbf{v} and \mathbf{v}' , and not just those that arise in a real-world context. This is intentional: as discussed in the introduction, we require the allocation algorithm to be prior-free in the sense of not knowing or even assuming a prior over value vectors. While utilizing a prior can potentially improve the algorithm's performance, there are several reasons for preferring a prior-free design: (1) prior-free algorithms are generally simpler than prior-dependent ones; (2) a single prior-free algorithm can be used in multiple different contexts without modification, and is therefore robust to changes in the market;

⁶ The precise form of the value-stability in [7] is actually $\frac{1-x^{-2\ell}}{1+x^{-2\ell}}$. The two functions are related to each other within small constant factors, and the form we use is mathematically more convenient for our proofs.

(3) in the context of fair ad allocation, parts of the prior may be sensitive information that is inappropriate or illegal to use; and, most importantly, (4) we would like to ensure fairness guarantees even if the prior is misspecified, or if new value vectors arise that were not expected to be in the support.

Social welfare. In this work we focus on the objective of social welfare maximization. The social welfare achieved by the allocation vector \mathbf{a} on value vector \mathbf{v} is given by $\mathbf{a} \cdot \mathbf{v}$. The maximum social welfare achievable for any value vector \mathbf{v} is therefore $\max_i v_i$. As discussed previously, welfare maximizing allocation functions generally do not satisfy value stability, so we look towards approximation. In keeping with the prior-free design of the allocation function, we measure its performance in the worst case.

► **Definition 4.** An allocation function \mathbf{a} achieves an α -approximation to social welfare for $\alpha \leq 1$ if for all value vectors \mathbf{v} we have $\mathbf{a}(\mathbf{v}) \cdot \mathbf{v} \geq \alpha \cdot \max_i v_i$.

2.1 Fairness and value stability

We now discuss how fairness connects to value stability. [7] introduces a model for ad allocation where users are drawn from a universe U endowed with a distance metric d that captures similarities between different users: the shorter the distance between two users, the more similar they are to each other. The model employs the notion of individual fairness from [10] which requires, informally, that similar users should receive similar outcomes. When the outcomes are probability distributions, Dwork and Ilvento [11] propose using the ℓ_∞ distance between difference distributions as a measure of (dis-)similarity in outcomes.

► **Definition 5** (Paraphrased from [10] and [11]). A function $\mathcal{A} : U \rightarrow \mathcal{O}$ assigning users to outcomes from a set \mathcal{O} is said to be **individually fair** with respect to distance metrics d over U and D over \mathcal{O} , if for all $u_1, u_2 \in U$ we have $D(\mathcal{A}(u_1), \mathcal{A}(u_2)) \leq d(u_1, u_2)$.

An allocation function $\mathbf{a} : U \rightarrow \Delta([k])$ satisfies **multiple-task fairness** with respect to distance metric d if for all $u_1, u_2 \in U$ and $i \in [k]$, we have $|a_i(u_1) - a_i(u_2)| \leq d(u_1, u_2)$.

Like in this work, [7] studies the design of allocation algorithms that return fair outcomes when the advertisers bid fairly. [7] connects fairness in advertisers' values and allocations to the distance metric over users through a constraint on advertisers' values called the "bid ratio constraint". Informally, if two users are close to each other under the distance metric d , then *all* advertisers are required to bid multiplicatively similar amounts on the users.

► **Definition 6** ([7]). A **bid ratio constraint** is a continuous function $h : [0, 1] \rightarrow [1, \infty]$ with $h(0) = 1$ and $h(1) = \infty$. We say that the value function $v^i : U \rightarrow [0, \infty)$ of advertiser i satisfies the bid ratio constraint h with respect to metric d if for all $u_1, u_2 \in U$: $\frac{1}{h(d(u_1, u_2))} \leq \frac{v^i(u_1)}{v^i(u_2)} \leq h(d(u_1, u_2))$.

[7] then proceed to design allocation algorithms that achieve multiple task fairness with respect to the underlying metric d as long as the values reported to the algorithm satisfy an appropriate bid ratio constraint. There is a close connection between our approach and that of [7]. An allocation algorithm is value stable if and only if it satisfies multiple task fairness whenever paired with value vectors satisfying an appropriate bid ratio constraint.

► **Fact 7.** Let $h : [0, 1] \rightarrow [1, \infty]$ be a strictly increasing function with $h(0) = 1$ and $h(1) = \infty$. Let $f : [1, \infty] \rightarrow [0, 1]$ be the inverse of h . Then an allocation function \mathbf{a} is value stable with respect to function f if and only if over any universe U endowed with distance metric d and

any set of value vectors assigned to users in the universe that satisfy the bid ratio constraint h , the function \mathbf{a} applied to the value vectors produces a multiple task fair allocation with respect to d .

3 The Inverse Proportional Allocation algorithm and our results

The main contribution of this paper is a new class of allocation functions that we call inverse-proportional (IPA) allocation. In this section we introduce the algorithm and outline its properties. Let us begin with an informal description of our mechanism. Recall that in a proportional allocation algorithm, the allocations a_i are proportional to (some function of) the values v_i . In the inverse proportional allocation algorithm, we set the *unallocated amounts* $1 - a_i$ to be proportional to *some decreasing function* of the values v_i . Specifically, the algorithm is parameterized by a decreasing function $g : \mathbb{R}^{\geq 0} \rightarrow (0, \infty]$, where $g(0) = \infty$ and $\lim_{x \rightarrow \infty} g(x) = 0$. We then deduct allocation from advertiser i in proportion to $g(v_i)$.

We will now describe the algorithm more precisely in three different but equivalent ways, which will be useful for our analysis.

Formulation 1: We start with the most direct formulation. Let \mathbf{v} be a value vector. For $t \geq 0$ and $i \in [k]$, let $y_i(t) = \max(0, 1 - tg(v_i))$, and $y(t) = \sum_{i \in [k]} y_i(t)$. Observe that $y(t)$ is a decreasing function with $y(0) = k$. Then, there is a unique value of t for which $y(t) = 1$. Let $t^* > 0$ be this value. Define $\mathbf{a}(\mathbf{v}) = (y_1(t^*), \dots, y_k(t^*))$. (Intuitively, we start with an allocation of 1 to every advertiser, corresponding to $t = 0$. We gradually raise t , decreasing the allocations in proportion to $g(v_i)$. When an allocation reaches 0, it remains 0. The process stops when the total allocated mass is precisely equal to 1.)

Formulation 2: We next describe an algorithmic version that describes how to find t^* . For $i \in [k]$ let w_i denote the unallocated amounts to be assigned to each advertiser. Observe that over any set S of advertisers, the total allocated amount is at most 1 and therefore the total unallocated amount is at least $|S| - 1$. The w_i 's split this amount in proportion to the $g(v_i)$'s. However, there is a problem: when the set S is large and v_i is very small, it may turn out to be the case that $w_i > 1$, and $1 - w_i$ is negative. In this case, we must assign an allocation of 0 to the corresponding advertiser. Formally, we start with S being the set of all advertisers. At each step we determine the weights $w_i(S)$ as below; we remove from S any advertisers j with $w_j(S) > 1$; we recurse on the remaining set of advertisers until all $w_i(S)$ values are ≤ 1 .

$$w_i(S) := (|S| - 1) \frac{g(v_i)}{\sum_{j \in S} g(v_j)}$$

Formulation 3: Our final formulation is a computationally efficient algorithm implementing the above idea. See Algorithm 1. The algorithm runs in time $O(k \log k)$.

► **Remark 8.** In Algorithm 1, when $k = 2$, the condition for the while loop is never satisfied, so both advertisers will receive nonzero allocations if both of their values are nonzero. Similarly, whenever $k \geq 2$, at least two advertisers will be assigned nonzero allocation probabilities when there are at least two advertisers with nonzero values.

3.1 Incentive Compatibility

Observe that IPA is a prior free algorithm. We will now formally prove incentive compatibility. The following lemma shows that the allocation assigned to any advertiser is weakly monotone increasing in her own value and weakly monotone decreasing in other advertisers' values.

■ **Algorithm 1** Inverse Proportional Allocation (IPA) parameterized by function g .

Input: Function $g : \mathbb{R}^{\geq 0} \rightarrow (0, \infty]$ with $g(0) = \infty$ and $\lim_{x \rightarrow \infty} g(x) = 0$. Values v_1, \dots, v_k .
Initialize $a_i = 0$ for $1 \leq i \leq k$
Sort the values so that $v_1 \leq v_2 \leq \dots \leq v_k$.
if $v_k = 0$ **then**
 Set $a_i = 1/k$ for all $1 \leq i \leq k$.
 return a.
end if
Initialize $s = \min(\{i \in [k] \mid v_i > 0\})$.
while $(k - s)g(v_s) \geq \sum_{j=s}^k g(v_j)$ **do**
 Increment s .
end while
for $i \geq s$ **do**
 Set $a_i = 1 - (k - s) \frac{g(v_i)}{\sum_{j=s}^k g(v_j)}$.
end for
return a.

► **Lemma 9.** For any value vector \mathbf{v} and any $i \in [k]$, let $\mathbf{v}' = (v_i', v_{-i})$ be another value vector that differs from \mathbf{v} only in coordinate i with $v_i > v_i'$. Then it holds that $a_i(\mathbf{v}) \geq a_i(\mathbf{v}')$ and $a_j(\mathbf{v}) \leq a_j(\mathbf{v}')$ for all $j \neq i$.

Proof. Consider determining the allocations $\mathbf{a}(\mathbf{v})$ and $\mathbf{a}(\mathbf{v}')$ using Formulation 1 of the IPA described above. Let $y(t)$ and $y'(t)$ denote the sums of allocations at a particular value of t for \mathbf{v} and \mathbf{v}' respectively. Recall that g is a decreasing function, and so, $g(v_i) \leq g(v_i')$. Then, at any t , we have $y(t) \geq y'(t)$. Therefore, the value of t^* that defines the final allocations is smaller under \mathbf{v} than under \mathbf{v}' . This immediately implies $a_j(\mathbf{v}) \leq a_j(\mathbf{v}')$ for all $j \neq i$. Then $a_i(\mathbf{v}) \geq a_i(\mathbf{v}')$ follows by recalling that the allocations sum up to 1. ◀

Lemma 9, by guaranteeing monotonicity, implies that the algorithm can be implemented as an incentive-compatible mechanism using an appropriately designed payment rule. For $g(x) = 1/x$, the payments can be specified explicitly; see the full version of the paper.

3.2 Choosing the function g

An important choice in the design space for IPA algorithms is the choice of the function g . This choice depends on the tradeoff between the stability condition we want to satisfy and the social welfare approximation we desire. However, there is an additional property that should be satisfied. Since we focus on prior free design and worst case analysis, the algorithm should be scale free in that when all values are scaled by a common factor, the allocation remains the same.

Consider, for example, the value vectors $\mathbf{v} = (1, x)$ and $\mathbf{v}' = (y, xy)$. The allocation produced by IPA with parameter g in the two cases is, respectively,

$$\mathbf{a} = \left(1 - \frac{g(1)}{g(1) + g(x)}, 1 - \frac{g(x)}{g(1) + g(x)} \right) \quad \text{and} \quad \mathbf{a}' = \left(1 - \frac{g(y)}{g(y) + g(xy)}, 1 - \frac{g(xy)}{g(y) + g(xy)} \right)$$

In order for these allocations to be the same, modulo normalization, it must hold that $g(xy) = g(x)g(y)$ for all $x, y > 0$. Under very mild assumptions on g (e.g. continuity), the only solution to this functional equation is $g(x) = x^{-\ell}$ where $\ell \in (0, \infty)$. Henceforth we focus on IPA parameterized by this family of polynomial functions.

► **Definition 10.** For $\ell \in (0, \infty)$, the *Inverse Proportional Allocation Algorithm with parameter ℓ* is Algorithm 1 with g defined as $g(x) = x^{-\ell}$ for $x \in [0, \infty)$.

This family of polynomial functions encompasses the entire range of tradeoffs between value stability and social welfare. To see this, consider IPA at the extremes. As $\ell \rightarrow 0$, IPA with $g(x) = x^{-\ell}$ ignores the value vector entirely and always assigns an equal allocation of $1/k$ to each advertiser. This algorithm achieves the strongest possible value stability condition, with $f(\lambda) = 0$ for all $\lambda \in [1, \infty]$. However, this value stability guarantee comes at the expense of social welfare: the approximation ratio approaches 0 as $k \rightarrow \infty$. As $\ell \rightarrow \infty$, IPA with $g(x) = x^{-\ell}$ becomes a highest-bid-wins allocation rule. This algorithm achieves an approximation ratio of 1 at the expense of value stability: $f(\lambda)$ approaches 1 for every $\lambda > 1$. When $\ell \in (0, \infty)$, IPA with $g(x) = x^{-\ell}$ interpolates between these two extremes, achieving different tradeoffs between social welfare and value-stability.

3.3 The value stability of IPA

We begin our analysis of the IPA by studying its value stability properties. The following theorem shows that the IPA with an appropriate choice of parameter $\ell > 0$ achieves value stability with respect to the class of functions $f_\ell(\lambda) = 1 - \lambda^{-2\ell}$.

► **Theorem 11.** For any $\ell > 0$ and any number $k > 0$ of advertisers, the inverse proportional allocation algorithm with parameter ℓ is value stable with respect to any function f that satisfies $f(\lambda) \geq f_\ell(\lambda) = 1 - \lambda^{-2\ell}$ for all $\lambda \in [1, \infty)$.

As briefly discussed above, the family of functions f_ℓ essentially represents the large spectrum of possible value stability conditions. See Figure 1 for a depiction of f_ℓ at different values of ℓ . In particular, f_0 is 0 everywhere and represents one extreme – a fixed allocation; f_∞ is 1 everywhere and represents the other extreme – an unconstrained allocation. Furthermore, for any non-decreasing function $f : [1, \infty) \rightarrow [0, 1]$ with $\lim_{x \rightarrow \infty} f(x) = 1$, there exists an ℓ with $f(x) \geq f_\ell(x)$ for all $x \in [1, \infty)$. Therefore, for any such stability constraint f , there exists a value for ℓ so that IPA with parameter ℓ is stable with respect to f .⁷

We also show a converse of Theorem 11. In particular, we observe that for stability constraints f that do not fall into the class of functions $\{f_\ell\}$, the best way to parameterize IPA is to choose the parameter ℓ corresponding to the function f_ℓ that most closely approximates f from below.

► **Theorem 12.** For any function f and parameter $\ell > 0$ such that over any number of advertisers $k > 0$, the IPA with parameter ℓ is value stable with respect to f , it holds that $f(x) \geq f_\ell(x)$ for all $x \geq 1$.

The proofs of Theorem 11 and Theorem 12 are deferred to Section 4.

3.4 The social welfare of IPA

Next we analyze the worst case performance of the IPA for social welfare.

► **Theorem 13.** For any $\ell \in (0, \infty)$, the inverse proportional allocation algorithm with parameter ℓ obtains a worst case approximation ratio for social welfare of at least:

$$\alpha_\ell := \min_{x \in (0,1)} (1 - x^\ell + x^{\ell+1}) = 1 - \frac{1}{\ell+1} \left(\frac{\ell}{\ell+1} \right)^\ell$$

⁷ In Section 5.2, we construct a variant of inverse proportional allocation that can better handle when $\lim_{x \rightarrow \infty} f(x) \neq 1$ and/or when the convergence to this limit is slow.

42:12 Individual Fairness in Advertising Auctions Through Inverse Proportionality

The approximation ratio of IPA is easy to analyze and we present the argument here in its entirety. The proof essentially shows that value vectors of the form $[1, x, x, \dots, x]$ for $x < 1$ are the worst case with respect to social welfare for IPA. At such a value vector, each advertiser with value x gets some non-zero allocation, which limits the mass that is placed on the highest value advertiser.

Proof of Theorem 13. Consider any value vector \mathbf{v} and order the values so that $v_1 \geq v_2 \geq \dots \geq v_k$. Since multiplicative scaling does not affect the allocation or the approximation ratio, we may assume without loss of generality that $v_1 = 1$. Then, the optimal social welfare over this value vector is also 1. We may also assume $v_2 > 0$, since otherwise the algorithm puts the entire allocation mass on the highest value.

Consider the set of non-zero allocations. We know that this set is of the form $\{1, \dots, m\}$ for some $1 \leq m \leq k$. The social welfare achieved by IPA can be written as:

$$\sum_{i=1}^m v_i a_i \geq v_1 a_1 + v_m (1 - a_1) = a_1 + v_m (1 - a_1).$$

To analyze this expression, we first recall that $a_1 = 1 - (m - 1) / \sum_{i=1}^m v_i^{-\ell}$. Observe that since advertiser m receives a nonzero allocation, we know that

$$1 \geq 1 - a_m = (m - 1) \frac{v_m^{-\ell}}{\sum_{i=1}^m v_i^{-\ell}}, \quad \text{so,} \quad \frac{m - 1}{\sum_{i=1}^m v_i^{-\ell}} \leq v_m^{\ell}.$$



This means that $a_1 \geq 1 - v_m^{\ell}$, and the approximation ratio can be lower bounded by $1 - v_m^{\ell} + v_m^{\ell+1}$. The lowest value of this expression is achieved at $v_m = \frac{\ell}{\ell+1}$, implying the bound claimed in the theorem. \blacktriangleleft

Observe that the approximation achieved by IPA is independent of the number k of advertisers. This is one of the main features that distinguishes IPA from the family of proportional allocation (PA) algorithms. PA necessarily exhibit worse and worse performance as the number of advertisers grows because every advertiser gets some non-zero fraction of the total allocation; IPA, on the hand, achieves good performance even when k is large by better handling advertisers that have small values relative to the largest value.

Figure 1 shows our bound from Theorem 13 for IPA as well as the bounds in Chawla et al. [7] for PA. Notice that when $k \geq 6$, IPA consistently outperforms PA across most of the parameter regime for ℓ . Moreover, the gap between these bounds grows as k becomes larger, eventually reaching ∞ as $k \rightarrow \infty$. When $k = 2$, however, the bound for PA beats the bound for IPA. This turns out to be an artifact of our theoretical analysis.⁸

Finally, we remark that IPA with parameter $\ell = 1$ arises as a full-information Nash equilibrium of an algorithm that allocates proportionally to bids with an all-pay payment rule. These mechanisms are considered in [14, 5, 8], and, in this context, the approximation ratio was previously shown to be at least $3/4$, the same bound we obtain. Our analysis is different and much simpler than those other works because, in particular, it does not require determining equilibrium strategies. Moreover, our analysis generalizes to $\ell \neq 1$.

⁸ For $k = 2$, the proof of Theorem 11 actually shows that the fairness guarantees of IPA turn out to be stronger, and indeed IPA with parameter ℓ is bid stable with respect to $f_{\ell/2}$. Thus, it is more appropriate to compare IPA with parameter ℓ and PA with ℓ , and the bound in Theorem 13 can be improved to show that these two algorithms achieve the same approximation ratio.

(a) Theoretical analysis: value-stability conditions f_ℓ .

(b) Empirical analysis: Bid stability profile of IPA.

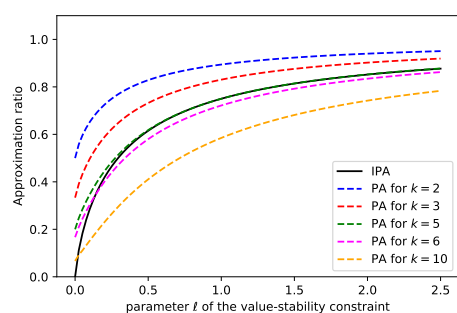

(c) Theoretical analysis: approx ratio of $\text{IPA}(\ell)$, $\text{PA}(2\ell)$.(d) Empirical analysis: approx ratio of $\text{IPA}(\ell)$, $\text{PA}(\frac{4}{3}\ell)$.

Figure 1 Properties of IPA across different settings of ℓ . The vertical axes are different in the first row and the second row. Refer to the full version for a discussion of the empirical analysis on the Yahoo dataset. (*Bid-stability* in the empirical analysis is closely related to *value-stability* in the theoretical analysis.)

4 An analysis of Inverse Proportional Allocation

We will now prove the value stability of Inverse Proportional Allocation, restated below.

► **Theorem 11.** *For any $\ell > 0$ and any number $k > 0$ of advertisers, the inverse proportional allocation algorithm with parameter ℓ is value stable with respect to any function f that satisfies $f(\lambda) \geq f_\ell(\lambda) = 1 - \lambda^{-2\ell}$ for all $\lambda \in [1, \infty)$.*

► **Theorem 12.** *For any function f and parameter $\ell > 0$ such that over any number of advertisers $k > 0$, the IPA with parameter ℓ is value stable with respect to f , it holds that $f(x) \geq f_\ell(x)$ for all $x \geq 1$.*

To prove Theorem 11, we need to understand how the allocation changes as the value vector changes by a small multiplicative amount. There are two main components to this analysis. First, given a value vector \mathbf{v} , we identify among all value vectors multiplicatively close to \mathbf{v} the one that changes a particular advertiser i 's allocation to the greatest extent. The second step is to then bound the corresponding change to the allocation. The main challenge in doing so is that as the value vector changes, the set of advertisers that get non-zero allocation also changes. This makes it difficult to directly compare the two allocations.

42:14 Individual Fairness in Advertising Auctions Through Inverse Proportionality

Before jumping into the main proof, let us consider the special case of $k = 2$ advertisers. This setting will illustrate the first component of our analysis as well as provide intuition for how the stability constraint f_ℓ relates to the parameter setting of IPA. In this setting, both advertisers always receive non-zero allocation and it becomes possible to write the allocation explicitly:

$$\mathbf{a}(\mathbf{v}) = \left[1 - \frac{v_1^{-\ell}}{v_1^{-\ell} + v_2^{-\ell}}, 1 - \frac{v_2^{-\ell}}{v_1^{-\ell} + v_2^{-\ell}} \right] = \left[\frac{v_1^\ell}{v_1^\ell + v_2^\ell}, \frac{v_2^\ell}{v_1^\ell + v_2^\ell} \right].$$

Now fix a value $\lambda \in [1, \infty)$ and consider any two value vectors \mathbf{v} and \mathbf{v}' with $v'_1 \in [1/\lambda, \lambda]v_1$ and $v'_2 \in [1/\lambda, \lambda]v_2$. We need to show that $|a_1 - a'_1| \leq f_\ell(\lambda) = 1 - \lambda^{-2\ell}$, which implies the same bound over $|a_2 - a'_2|$. Let us fix the vector \mathbf{v} and construct a worst-case value vector \mathbf{v}' such that a'_1 is as small as possible while respecting the multiplicative condition on the values. The monotonicity properties given in Lemma 9 enable us to identify such a value vector: $\mathbf{v}' = [v_1/\lambda, \lambda v_2]$. For this “worst-case” choice, we have:

$$\begin{aligned} a_1 - a'_1 &= \frac{v_1^\ell}{v_1^\ell + v_2^\ell} - \frac{(v'_1)^\ell}{(v'_1)^\ell + (v'_2)^\ell} \\ &= \frac{1}{1 + (v_2/v_1)^\ell} - \frac{1}{1 + \lambda^{2\ell}(v_2/v_1)^\ell}. \end{aligned}$$

Treating v_2/v_1 as a free variable, we see that the expression is maximized at $v_2/v_1 = 1/\lambda$ and we get,

$$a_1 - a'_1 \leq \frac{1}{1 + \lambda^{-\ell}} - \frac{1}{1 + \lambda^\ell} = \frac{\lambda^\ell - 1}{\lambda^\ell + 1} \leq 1 - \lambda^{-\ell}.$$

Recall that $f_\ell(\lambda) = 1 - \lambda^{-2\ell}$ and so the theorem follows for the case of $k = 2$.

For the general case, we carry out a similar analysis. We fix the value vector \mathbf{v} and a parameter $\lambda \in [1, \infty)$, and find a worst-case value vector \mathbf{v}' that is multiplicatively close to \mathbf{v} but changes the allocation by the greatest amount. The challenge is that in order to write down an explicit expression for some advertiser’s allocation we need to fix the set of advertisers that obtains non-zero allocation. This set changes as the value vector changes.

Our analysis is based on two ideas. The first idea is to examine the formulation of IPA in Algorithm 1 and consider how the allocations change if we terminate the while loop in Step 7 early. Let S be the set of agents that receive non-zero allocations in the algorithm. We argue that if we terminate the while loop early, “serving” a larger set of agents than S , then agents in S receive higher allocations. The second idea is when comparing allocations for the two vectors \mathbf{v} and \mathbf{v}' , if we pretend that the same set of advertisers are “served” under both the value vectors (with some of these allocations potentially being negative), then these allocations are additively close. We formalize these properties in the following lemmas.

► **Lemma 14.** *Let \mathbf{v} be any value vector with the advertisers reordered so that $v_1 \geq v_2 \geq \dots \geq v_k$. Let $S = [m]$ be the set of advertisers with nonzero allocation returned by the Inverse Proportional Allocation algorithm parameterized by function g on \mathbf{v} . Then for any m' with $m \leq m' \leq k$ and any $i \in [m]$,*

$$1 - (m - 1) \frac{g(v_i)}{\sum_{j \in [m]} g(v_j)} \leq 1 - (m' - 1) \frac{g(v_i)}{\sum_{j \in [m']} g(v_j)}.$$

► **Lemma 15.** *Let $\lambda \in [1, \infty)$ and $g(x) = x^{-\ell}$ for all $x \geq 0$. Let $S \subseteq [k]$ be an arbitrary set of advertisers, and i be an advertiser in S . Suppose that value vectors \mathbf{v} and \mathbf{v}' satisfy $v'_i = v_i/\lambda^2$ and $v'_j = v_j$ for all $j \in S$ with $j \neq i$. Then we have:*

$$\max \left(0, 1 - (|S| - 1) \frac{g(v_i)}{\sum_{j \in S} g(v_j)} \right) - \max \left(0, 1 - (|S| - 1) \frac{g(v'_i)}{\sum_{j \in S} g(v'_j)} \right) \leq f_\ell(\lambda).$$

Lemmas 14 and 15 along with Lemma 9 provide us with the necessary ingredients to prove Theorem 11.

Proof of Theorem 11. Let $\lambda \in [1, \infty)$, and let \mathbf{v} be an arbitrary value vector. We will focus on the allocation a_1 to advertiser 1, and consider a value vector \mathbf{v}' with $v'_i \in [1/\lambda, \lambda]v_i$ for all $i \in [k]$ that assigns the minimum possible allocation to 1. Since the IPA is a scale-free allocation algorithm, we may equivalently focus on value vectors \mathbf{v}' with $v'_i \in [1/\lambda^2, 1]v_i$ for all $i \in [k]$. Lemma 9 then implies that the allocation to 1 is minimized at $v'_1 = v_1/\lambda^2$ and $v'_{-1} = v_{-1}$.

Let S be the set of advertisers receiving nonzero allocation from IPA on \mathbf{v} , and let S' be the corresponding set for \mathbf{v}' . Let \mathbf{a} and \mathbf{a}' be the corresponding allocations. By Lemma 9, $a'_i \geq a_i$ for all $i \neq 1$ and $a'_1 \leq a_1$. Therefore, $S \subseteq S' \cup \{1\}$. We want to show: $a_1 - a'_1 \leq f_\ell(\lambda)$. We now break up our analysis into three cases.

Case 1: $1 \notin S$. In this case, $a_1 = 0$ and $a_1 - a'_1 \leq 0$, so the statement trivially holds.

Case 2: $1 \in S$ and $1 \in S'$. In this case we have $S \subseteq S'$, and $a_1, a'_1 > 0$. So, by Lemma 14 we obtain:

$$a_1 = 1 - (|S| - 1) \frac{g(v_1)}{\sum_{j \in S} g(v_j)} \leq 1 - (|S'| - 1) \frac{g(v_1)}{\sum_{j \in S'} g(v_j)}.$$

Furthermore, Lemma 15 implies:

$$1 - (|S'| - 1) \frac{g(v_1)}{\sum_{j \in S'} g(v_j)} \leq \max \left(0, 1 - (|S'| - 1) \frac{g(v'_1)}{\sum_{j \in S'} g(v'_j)} \right) + f_\ell(\lambda) = a'_1 + f_\ell(\lambda),$$

so $a_1 - a'_1 \leq f_\ell(\lambda)$ as desired.

Case 3: $1 \in S$ and $1 \notin S'$. In this case we cannot directly apply Lemma 15 to S' . Instead, let $R = \{i \mid v'_i \geq v'_1\} = \{i \mid v_i \geq v_1/\lambda^2\}$. We claim that $S \subseteq R$. This is because $1 \notin S'$, and therefore all $i \in S'$ have $v'_i \geq v'_1$, and therefore $S \subseteq S' \cup \{1\} \subseteq R$.

Thus, we can apply Lemma 14 to obtain:

$$a_1 = 1 - (|S| - 1) \frac{g(v_1)}{\sum_{j \in S} g(v_j)} \leq 1 - (|R| - 1) \frac{g(v_1)}{\sum_{j \in R} g(v_j)}.$$

Next, we can apply Lemma 15 to R to see that

$$1 - (|R| - 1) \frac{g(v_1)}{\sum_{j \in R} g(v_j)} \leq \max \left(0, 1 - (|R| - 1) \frac{g(v'_1)}{\sum_{j \in R} g(v'_j)} \right) + f_\ell(\lambda).$$

Finally, since advertiser 1 receives a zero allocation on value vector \mathbf{v}' and since $v'_1 = \min(\{v'_i \mid i \in R\})$, we know by Step 7 of Algorithm 1 that $(|R| - 1)g(v'_1) \geq \sum_{j \in R} g(v'_j)$. Thus, $1 - (|R| - 1) \frac{g(v'_1)}{\sum_{j \in R} g(v'_j)} \leq 0$. This shows that $a_1 \leq f_\ell(\lambda) = a'_1 + f_\ell(\lambda)$ as desired. \blacktriangleleft

It remains to prove the lemmas.

Proof of Lemma 14. For $i \in [m]$, let $a_i = 1 - (m - 1) \frac{(1/v_i)}{\sum_{j \in [m]} (1/v_j)}$ and $a'_i = 1 - (m' - 1) \frac{(1/v_i)}{\sum_{j \in [m']} (1/v_j)}$. We wish to prove that for $i \in [m]$, it holds that $a_i \leq a'_i$. In order to compare these values, we consider the following expressions for a_i and a'_i using Formulation 1 of IPA in Section 3. First, we see that $a_j = \max(0, 1 - t \cdot (1/v_j))$, where t is the unique value such that $\sum_{j=1}^k a_j = 1$. Equivalently, since $[m']$ contains the set of advertisers that receive

42:16 Individual Fairness in Advertising Auctions Through Inverse Proportionality

nonzero allocations, t is also the unique value such that $\sum_{j \in [m']} a_j = 1$. Similarly, we see that $a'_j = 1 - t' \cdot (1/v_j)$, where t' is the unique value such that $\sum_{j \in [m']} a'_j = 1$. In other words, the only difference between t and t' is that the allocations are not constrained to be nonnegative for t' . This implies that $t' \leq t$, and so for $i \in [m]$:

$$a_i = \max(0, 1 - t \cdot (1/v_i)) = 1 - t \cdot (1/v_i) \leq 1 - t' \cdot (1/v_i) = a'_i$$

as desired. ◀

Proof of Lemma 15. Let us consider the expression we need to bound:

$$\begin{aligned} d &:= \max \left(0, 1 - (|S| - 1) \frac{g(v_i)}{\sum_{j \in S} g(v_j)} \right) - \max \left(0, 1 - (|S| - 1) \frac{g(v'_i)}{\sum_{j \in S} g(v'_j)} \right) \\ &= \max \left(0, 1 - \frac{(|S| - 1) \cdot g(v_i)}{g(v_i) + \sum_{j \in S; j \neq i} g(v_j)} \right) - \max \left(0, 1 - \frac{(|S| - 1) \cdot g\left(\frac{v_i}{\lambda^2}\right)}{g\left(\frac{v_i}{\lambda^2}\right) + \sum_{j \in S; j \neq i} g(v_j)} \right) \end{aligned}$$

We want to show that $d \leq f_\ell(\lambda)$. We first observe that the statement is vacuously true when the maximum in the first term of d is equal to 0 as well as when $|S| = 1$. So henceforth we assume that $|S| \geq 2$ and $1 - (|S| - 1) \frac{g(v_i)}{g(v_i) + \sum_{j \in S; j \neq i} g(v_j)} > 0$.

Let r denote the ratio $(\sum_{j \in S; j \neq i} g(v_j))/g(v_i)$. Recalling that $g(v_i/\lambda^2) = \lambda^{2\ell} g(v_i)$, we get,

$$\begin{aligned} d(r) &= 1 - \frac{|S| - 1}{1 + r} - \max \left(0, 1 - \frac{|S| - 1}{1 + \lambda^{-2\ell} b} \right) \\ &= \frac{|S| - 1}{\max(|S| - 1, 1 + \lambda^{-2\ell} b)} - \frac{|S| - 1}{1 + r} \end{aligned}$$

Now we will consider three cases. First, if $|S| = 2$, the expression becomes

$$d(r) = \frac{1}{1 + \lambda^{-2\ell} b} - \frac{1}{1 + r}$$

This is identical to the expression we analysed for the 2 advertisers setting at the beginning of this section. The expression maximized at $b = \lambda^\ell$, and we obtain

$$d(r) \leq \frac{\lambda^\ell - 1}{\lambda^\ell + 1} \leq 1 - \lambda^{-\ell}.$$

Second, suppose that $|S| > 2$ and $\max(|S| - 1, 1 + \lambda^{-2\ell} r) = |S| - 1$. This means that $r \leq \lambda^{2\ell}(|S| - 2)$. Then we have, $d(r) = 1 - (|S| - 1)/(1 + r)$. Note that $d(r)$ is increasing as a function of r . So in this case it is maximized at $r = \lambda^{2\ell}(|S| - 2)$. We get

$$\max_{r \leq \lambda^{2\ell}(|S| - 2)} d(r) \leq 1 - \frac{|S| - 1}{1 + \lambda^{2\ell}(|S| - 2)} \leq 1 - \lambda^{-2\ell} = f_\ell(\lambda)$$

Finally, suppose that $|S| > 2$ and $\max(|S| - 1, 1 + \lambda^{-2\ell} r) = 1 + \lambda^{-2\ell} r$. This means that $r \geq \lambda^{2\ell}(|S| - 2)$. Then we have

$$d(r) = (|S| - 1) \left(\frac{1}{1 + \lambda^{-2\ell} r} - \frac{1}{1 + r} \right)$$

In this case, the function $d(r)$ is non-increasing on $r \in [\lambda^{2\ell}(|S| - 2), \infty)$ (see Proposition 32 in the full version of the paper.) Therefore, once again we have,

$$\max_{r \geq \lambda^{2\ell}(|S| - 2)} d(r) \leq 1 - \frac{|S| - 1}{1 + \lambda^{2\ell}(|S| - 2)} \leq 1 - \lambda^{-2\ell} = f_\ell(\lambda)$$

This completes the proof. ◀

We now prove Theorem 12.

Proof of Theorem 12. Let x be any value in $[1, \infty]$. We consider the value vectors $\mathbf{v} = [x, \dots, x]$ and $\mathbf{v}' = [x^2, 1, \dots, 1]$. Notice that $\max_{i \in [k]} \left(\max \left(\frac{v_i}{v'_i}, \frac{v'_i}{v_i} \right) \right) = x$. Observe that the allocation a'_1 is $1 - (k-1) \frac{g(x^2)}{g(x^2) + (k-1)g(1)} = 1 - \frac{g(x^2)}{\frac{g(x^2)}{k-1} + g(1)}$. As $k \rightarrow \infty$, this becomes $1 - \frac{g(x^2)}{g(1)} \rightarrow 1 - x^{-2\ell}$. In comparison, the allocation $a_1 \rightarrow 0$. Thus, we must have that $f(x) \geq 1 - x^{-2\ell}$. ◀

5 Near-optimality of Inverse Proportional Allocation

In this section, we prove that inverse proportional allocation is near-optimal in comparison to any prior-free allocation algorithm. First, we show in Section 5.1 that IPA is optimal for the family of value-stability constraints f_ℓ : that is, no prior-free allocation algorithm can beat the performance of IPA within this family of value-stability constraints. Then, we consider the performance of IPA for value-stability constraints f that are not of the form f_ℓ . We construct an extension of IPA called *capped IPA* that performs near-optimally for any “reasonable” function f satisfying a mild constraint. In Section 5.2, we discuss this extension of IPA, and in Section 5.3, we prove our near-optimality result: for any function f satisfying the constraint that admits a prior-free value-stable allocation algorithm with approximation ratio α , there exists a capped IPA algorithm that achieves an approximation ratio of $\Omega\left(\frac{\alpha}{\ln(1/\alpha)+1}\right)$.

5.1 Optimality of IPA

Recall that Theorems 11 and 13 together imply that for any $\ell > 0$, it is possible to achieve value stability with respect to the function f_ℓ while obtaining an $\alpha_\ell = 1 - \ell^\ell / (\ell + 1)^{\ell+1}$ approximation to social welfare. We will now show that no prior-free allocation algorithm that is value stable with respect to f_ℓ can obtain a better approximation as the number of advertisers k tends to infinity.

► **Theorem 16.** *For every $\ell > 0$, no prior-free allocation algorithm that is value stable with respect to f_ℓ can obtain a worst case approximation ratio for social welfare over k advertisers that is better than $\alpha_\ell - 1/k$.*

Thus, within the family f_ℓ , the approximation ratio obtained by IPA is optimal within an additive factor of $1/k$. In the limit as $k \rightarrow \infty$, this theorem implies that IPA achieves the optimal approximation ratio in comparison to any prior-free allocation algorithm that is value-stable for f_ℓ .

The proof of this theorem is based on the following lemma from [7] that we restate and reprove for completeness.

► **Lemma 17** ([7], rephrased). *Let $f : [1, \infty] \rightarrow [0, 1]$ be any weakly increasing function, and $\lambda > 1$ be an arbitrary real number. Then, no prior free allocation algorithm over k advertisers that is value-stable with respect to f can obtain an approximation to social welfare better than: $\left(\frac{1}{k} + f(\lambda) + \lambda^{-2} \left(1 - \frac{1}{k} - f(\lambda)\right)\right)$.*

Proof. Consider the value vectors $\mathbf{v} = [1, \dots, 1]$. Since the total allocation is 1, there must exist an advertiser $i \in [k]$ such that $a_i \leq 1/k$. We define a value vector \mathbf{v}' that’s within a multiplicative factor of λ on every coordinate and place a lower bound on its approximation ratio for social welfare. More specifically, we consider \mathbf{v}' so $v'_i = \lambda$ and $v'_j = \lambda^{-1}$ for $j \neq i$.

Since the allocation is value-stable, it must satisfy $a'_i \leq f(\lambda) + a_i \leq \frac{1}{k} + f(\lambda)$. Thus, the approximation ratio for social welfare is $a'_i + \lambda^{-2}(1 - a'_i)$. Using the above bound on a'_i , we see that the approximation ratio is at most $\frac{1}{k} + f(\lambda) + \lambda^{-2}(1 - \frac{1}{k} - f(\lambda))$ as desired. ◀

The proof of Theorem 16 now follows directly.

Proof of Theorem 16. The theorem follows from Lemma 17 by substituting $f(\lambda) = 1 - \lambda^{-2\ell}$; minimizing the expression over λ ; and recalling that $\alpha_\ell = \min_{x \in (0,1)}(1 - x^\ell + x^{\ell+1})$, which can be written as $\min_{\lambda > 1}(1 - \lambda^{-2\ell} + \lambda^{-2(\ell+1)}) = \min_{\lambda > 1}(f(\lambda) + \lambda^{-2}(1 - f(\lambda)))$. ◀

5.2 A simple extension of IPA

We now consider general value-stability constraints f , and consider the design of algorithms that are value-stable for general f . As mentioned earlier, for every function f with $\lim_{x \rightarrow \infty} f(x) = 1$, there exists a parameter $\ell > 0$ such that $f(x) \leq f_\ell(x)$ for all $x \in [1, \infty)$, and so value stability for such functions can be achieved with an IPA algorithm. However, this family excludes stability constraints f with $\lim_{x \rightarrow \infty} f(x) = c < 1$. With this shortcoming in mind, we propose an extension of IPA that we call *capped IPA*. The idea is to run IPA with probability c , so that the maximum difference between any two advertiser's allocations is always capped by c .

► **Algorithm 18.** *The Capped Inverse Proportional Allocation algorithm with parameters ℓ and c applies the Inverse Proportional Algorithm with parameter ℓ with probability c and uniformly assigns allocation across all advertisers with probability $1 - c$.*

The value stability and social welfare properties of capped inverse proportional allocation follow immediately from our analysis of IPA in the previous sections.

► **Corollary 19.** *Let $\ell > 0$ and $c \in [0, 1]$ be parameters, and let $k > 0$ be any number advertisers. The capped inverse proportional allocation algorithm with parameters ℓ and c is value stable with respect to any function f that satisfies $f(\lambda) \geq cf_\ell(\lambda) = c(1 - \lambda^{-2\ell})$ for all $\lambda \in [1, \infty)$. Moreover, this algorithm obtains a worst case approximation ratio for social welfare of at least*

$$c \cdot \alpha_\ell = c \left(1 - \frac{1}{\ell + 1} \left(\frac{\ell}{\ell + 1} \right)^\ell \right).$$

Proof. For any value vector \mathbf{v} , we compare the allocation \mathbf{a} of this capped IPA algorithm to the allocation \mathbf{a}' of IPA with parameter ℓ . For any $i \in [k]$, it holds that $a_i = ca'_i + (1 - c)/k$. The value-stability and social welfare properties then follow from Theorem 11 and Theorem 13. ◀

5.3 Near-optimality with respect to general value-stability functions

We now show that capped inverse proportional allocation can perform near-optimally for any “reasonable” value stability constraint $f : [1, \infty) \rightarrow [0, 1]$. First, we argue that any “reasonable” value-stability constraint f should at minimum satisfy the following conditions: (a) f is weakly increasing, and (b) $f(x)/\ln(x)$ is weakly decreasing.

For constraint (a), observe that if a prior-free allocation algorithm is value-stable with respect to f , then it is also value-stable with respect to $g(x) = \min_{y \geq x} g(y)$ which is weakly increasing.

For constraint (b), observe that if a prior-free allocation algorithm is value-stable with respect to f , then it is also value-stable with respect to $g(x) = \inf_{n \in \mathbb{N}} n f(x^{1/n})$. Then for any $x \in [1, \infty]$ and $n \in \mathbb{N}$, it holds that $g(x) \cdot n = n \cdot \inf_{m \in \mathbb{N}} m f(x^{1/m}) \geq \inf_{m' \in \mathbb{N}} m' f(x^{n/m'}) = g(x^n)$. This implies that $\frac{g(x)}{\ln(x)} \geq \frac{g(x^n)}{\ln(x^n)}$. After some “smoothing”, this essentially guarantees that $g(x)/\ln(x)$ is weakly decreasing.

With these conditions in place, we consider value-stability constraints f where f is weakly increasing and $f(x)/\ln(x)$ is weakly decreasing. We now show that for any f of this form, the family of capped inverse proportional allocation algorithms achieves value-stability while obtaining an approximation ratio that is only logarithmically worse than optimal.

► **Theorem 20.** *Let $f : [1, \infty] \rightarrow [0, 1]$ be any function where $f(x)/\ln(x)$ is weakly decreasing and f is weakly increasing. Suppose that there exists a prior-free allocation algorithm that is value-stable for f that achieves a worst-case approximation ratio of α . Then, there exist parameters $\ell \in (0, \infty)$ and $c \in [0, 1]$ such that the capped IPA algorithm with parameters ℓ and c is value-stable for f and achieves a worst-case approximation ratio of at least: $\left(\frac{\alpha/2}{2\ln(1/\alpha)+1}\right)$.*

The proof of Theorem 20 leverages Lemma 17 to lower bound $f(x)$ as a function of the approximation ratio. This bound is applied to a carefully chosen point, and the basic properties of f can be used to extend this to a pointwise lower bound by a “well-behaved” function F . This function F can be approximated from below by a function of the form $c \cdot f_\ell$. We then show that capped IPA with parameters ℓ and c achieves a significant fraction of the optimal approximation ratio for the value-stability constraint f .

Proof of Theorem 20. Let $c = \frac{\alpha - \alpha^2}{1 - \alpha^2} = \frac{\alpha}{1 + \alpha}$, and let $F(x) = \min(c, \ln(x) \cdot c / \ln(1/\alpha))$. We begin by showing that $f(x) \geq F(x)$ for all $x \in [1, \infty]$. By Lemma 17, we know that $f(x) \geq (\alpha - x^{-2}) / (1 - x^{-2})$ for every $x \in [1, \infty)$. Now, we apply this fact at $x = 1/\alpha$ to conclude that $f(1/\alpha) \geq \frac{\alpha - \alpha^2}{1 - \alpha^2} = c$. Using the basic properties of f , we can extend this to obtain a pointwise lower bound by F . Since $\frac{f(x)}{\ln(x)}$ is decreasing, this means that for $x \leq 1/\alpha$, we have that $f(x) \geq \ln(x) \cdot c / \ln(1/\alpha) = F(x)$. Moreover, since f is weakly increasing, we know that for $x \geq 1/\alpha$, we have that $f(x) \geq c = F(x)$. Putting these two conditions together, we have that $f(x) \geq \min(c, \ln(x) \cdot c / \ln(1/\alpha)) = F(x)$.

Now, we relate $F(x)$ to the value-stability condition for a capped IPA algorithm. By Proposition 33 in the full version of the paper, we know that $\min(c, \ln(x) \cdot c / \ln(1/\alpha)) \geq c(1 - x^{-1/\ln(1/\alpha)})$. This implies that $f(x) \geq c(1 - x^{-1/\ln(1/\alpha)})$ as desired. Let $\ell = 1/(2\ln(1/\alpha))$. We claim that the capped IPA with parameters ℓ and c is value-stable for f . By Corollary 19, we know that this capped IPA algorithm is value-stable for $f(x) \geq c(1 - x^{-2\ell}) = \frac{\alpha - \alpha^2}{1 - \alpha^2}(1 - x^{-1/\ln(1/\alpha)})$, as desired.

Now, we show that this capped IPA algorithm achieves an approximation ratio of at least $\left(\frac{\alpha/2}{2\ln(1/\alpha)+1}\right)$. By Corollary 19, we know that the algorithm achieves social welfare of at least

$$c \left(1 - \frac{1}{\ell + 1} \left(\frac{\ell}{\ell + 1}\right)^\ell\right) \geq \frac{\alpha - \alpha^2}{1 - \alpha^2} \left(1 - \frac{1}{\ell + 1}\right) = \frac{\alpha}{1 + \alpha} \left(\frac{1}{2\ln(1/\alpha) + 1}\right).$$

This completes the proof. ◀

Theorem 20 shows that capped IPA achieves an approximation ratio of at least $\frac{\alpha/2}{2\ln(1/\alpha)+1}$, where α is the best possible approximation ratio achieved by any prior-free allocation algorithm that is value-stable for f . In fact, this result also shows that when $\alpha \geq 0.5$ (i.e. the best possible approximation ratio achieves a 0.5 fraction of the social welfare), then capped IPA achieves a approximation ratio of at least 0.2α . In other words, in this regime, capped

IPA is able to achieve a *constant fraction* of the optimal approximation ratio. When the approximation ratio is small, the guarantees provided in Theorem 20 do degrade with $\ln(1/\alpha)$. An interesting direction for future work would be to try to close this gap and maintain a constant fraction of the optimal social welfare even when the approximation ratio is small.

6 Future work

In this work we address the tradeoffs between fairness and social welfare through a simplistic model that assumes that advertisers bid to maximize their utility in every auction. In reality, advertisers have budgets and optimize for their long term returns. Incorporating fairness constraints into these more general models of ad auctions is an interesting avenue for future work. For the model we study, our work presents an alternative to the proportional allocation mechanism of [7]. The two classes of allocation algorithms perform well under different contexts, and it would be interesting to understand whether one can interpolate between the two to obtain fairness guarantees as strong as those of proportional allocation while at the same time achieving performance commensurate with that of inverse proportional allocation. Finally, a very interesting open direction is to examine ad auction design under a group fairness constraint.

References

- 1 Muhammad Ali, Piotr Sapiezynski, Miranda Bogen, Aleksandra Korolova, Alan Mislove, and Aaron Rieke. Discrimination through optimization: How facebook’s ad delivery can lead to biased outcomes. *PACMHCI*, 3(CSCW):199:1–199:30, 2019.
- 2 Julia Angwin and Terry Perris. Facebook lets advertisers exclude users by race. *ProPublica*, 2016.
- 3 Julia Angwin, Noam Scheiber, and Ariana Tobin. Facebook job ads raise concerns about age discrimination. *The New York Times* in collaboration with *ProPublica*, 2017.
- 4 Maria-Florina Balcan, Travis Dick, Ritesh Noothigattu, and Ariel D. Procaccia. Envy-free classification. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, pages 1238–1248, 2019.
- 5 Ioannis Caragiannis and Alexandros A. Voudouris. Welfare guarantees for proportional allocations. *Theory Comput. Syst.*, 59(4):581–599, 2016.
- 6 L. Elisa Celis, Anay Mehrotra, and Nisheeth K. Vishnoi. Toward controlling discrimination in online ad auctions. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, pages 4456–4465, 2019.
- 7 Shuchi Chawla, Christina Ilvento, and Meena Jagadeesan. Multi-category fairness in sponsored search auctions. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, FAT* ’20*, pages 348–358, New York, NY, USA, 2020. Association for Computing Machinery.
- 8 George Christodoulou, Alkmini Sgouritsa, and Bo Tang. On the efficiency of the proportional allocation mechanism for divisible resources. *Theory Comput. Syst.*, 59(4):600–618, 2016.
- 9 Amit Datta, Michael Carl Tschantz, and Anupam Datta. Automated experiments on ad privacy settings. *PoPETs*, 2015(1):92–112, 2015.
- 10 Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard S. Zemel. Fairness through awareness. In *3rd Innovations in Theoretical Computer Science, ITCS 2012, Cambridge, MA, USA, January 8-10, 2012*, pages 214–226, 2012.
- 11 Cynthia Dwork and Christina Ilvento. Fairness under composition. In *10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10-12, 2019, San Diego, California, USA*, pages 33:1–33:20, 2019.

- 12 Lodewijk Gelauff, Ashish Goel, Kamesh Munagala, and Sravya Yandamuri. Advertising for demographically fair outcomes. *CoRR*, abs/2006.03983, 2020.
- 13 Lily Hu and Yiling Chen. Fair classification and social welfare. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, FAT* '20, pages 535–545, New York, NY, USA, 2020. Association for Computing Machinery.
- 14 Ramesh Johari and John N. Tsitsiklis. Efficiency loss in a network resource allocation game. *Math. Oper. Res.*, 29(3):407–435, 2004.
- 15 Michael P. Kim, Aleksandra Korolova, Guy N. Rothblum, and Gal Yona. Preference-informed fairness. In *Proceedings of the 11th Innovations in Theoretical Computer Science, ITCS 2020, Seattle, Washington, USA, January 12-14, 2020*, page to appear, 2020.
- 16 Ava Kofman and Ariana Tobin. Facebook ads can still discriminate against women and older workers, despite a civil rights settlement. ProPublica, 2019.
- 17 Anja Lambrecht and Catherine 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.
- 18 Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2007), October 20-23, 2007, Providence, RI, USA, Proceedings*, pages 94–103. IEEE Computer Society, 2007.
- 19 Milad Nasr and Michael Carl Tschantz. Bidding strategies with gender nondiscrimination constraints for online ad auctions. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, FAT* '20, pages 337–347, New York, NY, USA, 2020. Association for Computing Machinery.
- 20 Latanya Sweeney. Discrimination in online ad delivery. *Commun. ACM*, 56(5):44–54, 2013.
- 21 Yahoo. Yahoo. a1 - yahoo! search marketing advertising bidding data. URL: <https://webscope.sandbox.yahoo.com/catalog.php?datatype=a>.
- 22 Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez-Rodriguez, Krishna P. Gummadi, and Adrian Weller. From parity to preference-based notions of fairness in classification. In *Advances in Neural Information Processing Systems 30: NeurIPS 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 229–239, 2017.