

Individually-Fair Auctions for Multi-Slot Sponsored Search*

Shuchi Chawla
UT-Austin
shuchi@cs.utexas.edu

Rojin Rezvan
UT-Austin
rojinrezvan@utexas.edu

Nathaniel Sauerberg
UT-Austin
njs@cs.utexas.edu

Abstract

We design fair sponsored search auctions that achieve a near-optimal tradeoff between fairness and quality. Our work builds upon the model and auction design of Chawla and Jagadeesan [4], who considered the special case of a single slot. We consider sponsored search settings with multiple slots and the standard model of click through rates that are multiplicatively separable into an advertiser-specific component and a slot-specific component. When similar users have similar advertiser-specific click through rates, our auctions achieve the same near-optimal tradeoff between fairness and quality as in [4]. When similar users can have different advertiser-specific preferences, we show that a preference-based fairness guarantee holds. Finally, we provide a computationally efficient algorithm for computing payments for our auctions as well as those in previous work, resolving another open direction from [4].

*This work was supported in part by NSF Award CCF-2008006.

1 Introduction

We study the design of ad auctions under a fairness constraint. Fairness in the context of sponsored content has received considerable attention in recent years. It has been observed, for example, that ads on platforms such as Facebook and Google disproportionately target certain demographics, discriminating across users on the basis of race and gender. Furthermore, standard auction formats such as highest-bids-win can lead to discrimination even when the input to these algorithms, namely bids, CTRs, and relevance scores are themselves non-discriminatory.

Chawla et al. [5] initiated the study of optimal auction design under the constraint that the auction does not add any unfairness beyond what is already present in bids, and proposed a class of *proportional allocation* algorithms as a solution that achieves fairness while also providing an approximation to the optimal social welfare. In a followup work, Chawla and Jagadeesan [4] designed a class of *inverse proportional allocation* algorithms and showed that this class of mechanisms achieves an optimal tradeoff between social welfare and fairness. Both of these works focused on the simple case of a single item auction and left open the problem of designing a fair and efficient multi-slot position auction.

In this paper we extend the design of fair auctions from the single item setting to arbitrary position auction settings. We show that both the proportional allocation and inverse proportional allocation algorithms can be adapted to the setting of a position auction while inheriting their single-unit fairness properties as well as their approximation to social welfare. As in [5, 4] our auctions provide fair solutions when the advertisers' bids are themselves non-discriminatory. Auctions for multi-slot settings must take into account both the advertisers' preferences over users as captured by per-click values, as well as the users' preferences over advertisers as captured by click through rates. We consider two different models for formalizing fairness in these settings. In the first, we consider differences of allocation across users that are close both in terms of the values advertisers assign to them as well as in terms of their own click through rates; we require that such users receive similar allocations. In the second setting, we consider pairs of users that are similarly qualified as per advertisers' values, but have different preferences (i.e. CTRs). In this case, while the users may receive different allocations, we require that allocations are suitably aligned with users' preferences. We elaborate on the details of these models below. Finally, we address another open question in [5, 4] and show how to efficiently compute supporting prices for both proportional and inverse proportional allocation.

Formalizing fairness across users. Consider two users Alice and Bob who are similar in most respects but differ in a sensitive demographic such as gender or race. Individual fairness then posits that Alice and Bob should see similar ad allocations. For example, it would be unfair to show more employment ads to Bob and more online retail ads to Alice. One potential source of unfairness in ad allocations is the use of discriminatory targeting by advertisers. However, empirical studies as well as theoretical analysis shows that unfairness in allocations can persist even in the absence of discriminatory targeting. The culprit is allocation algorithms that turn minor differences in advertisers' bids into large swings in allocation. Suppose, for example, that an employment agency places a slightly higher value on Bob than on Alice whereas an online retail store places a slightly higher value on Alice because of minor differences in the users' profiles. Then the highest-bid-wins auction would show entirely different ads to the two users.

To combat this problem, Chawla and Jagadeesan [4] formalize the notion of fairness in auctions

as a “value stability” constraint. Informally speaking, value stability requires that whenever two users receive multiplicatively similar values from all advertisers (such as Alice and Bob in the example above) they must receive close allocations (as measured in terms of the ℓ_∞ distance between the respective probability distributions over the ad displayed). Previous work shows that while optimal auctions do not satisfy value stability, there are simple auction formats that do. In the *Proportional Allocation* (PA) mechanism, allocations are proportional to (some increasing function of) the advertisers’ reported values. In the *Inverse Proportional Allocation* (IPA) mechanism, the unallocated amounts, i.e., one minus the probability of allocation, are inversely proportional to (some increasing function of) the advertisers’ reported values. In both mechanisms, the allocation is a sufficiently smooth function of the advertisers’ values and therefore satisfies some form of value stability. We mostly focus on the IPA mechanism in this paper as it provides better tradeoffs between fairness and welfare.

Multi-slot extensions. As a simple extension of the single slot setting, consider a setting with k slots, where each ad and each slot are equally likely to be clicked by the user, so the relative placement of ads in slots does not matter. In this case, one straightforward way to extend the single-slot allocations is to simply multiply them by k ; if this provides a valid allocation, the fairness and welfare guarantees follow immediately from the single-slot case. The problem is that some ads may receive a total allocation greater than 1 and simply capping allocations at 1 breaks the fairness guarantee. We propose a different extension of the IPA. As in the single slot case, we ensure that the unallocated amounts to advertisers are inversely proportional to (some function of) the reported values, subject to the total allocation equaling k . The fairness a.k.a. value stability of this extension follows easily from the single-slot special case. We further show that the social welfare approximation of multi-slot IPA matches its approximation for the single-item case by characterizing worst case instances for the approximation factor.

While the above discussion provides a complete story for the case of a multi-unit auction, in the case of online advertising, we also need to take click through rates into account. Throughout this paper, we assume that click through rates are multiplicatively separable into ad-specific and slot-specific components. In other words, the click through rate of an ad i placed in slot j is given by $\alpha_i \times \beta_j$ for some parameters α and β specific to each user that are known to the platform/auctioneer. We further assume that all users weakly prefer earlier slots to later slots. Under these assumptions, we present an extension of the IPA to the ad auction setting that exactly maintains the social welfare guarantees of their single- and multi-unit counterparts. In particular, the social welfare approximation is independent of the number of slots.

Fairness in the context of click through rates is tricky to define, however. As before, we may assume that if two users are similarly qualified for all ads but differ in their sensitive attributes, then the two users receive multiplicatively similar per-click values from all advertisers. However, click through rates capture the users’ own preferences and similar users may not have similar click through rates. What sort of fairness guarantees can we then provide?

We first show that differences in slot-specific CTRs do not impact fairness guarantees.¹ In particular, two users with similar values and similar ad-specific CTRs α receive allocations that are close in ℓ_∞ distance. In particular, the probability of assigning any particular slot to any particular

¹In fact, the allocations produced by our algorithms do not depend on the slot-specific CTRs, although the payments made by advertisers necessarily must.

ad is additively close for the two users. In fact, this additive closeness holds also for the probability that any particular ad is assigned to slot j or better for any j .

We then consider settings with similarly qualified users that have arbitrarily different ad-specific and slot-specific CTRs. Observe that in order to achieve any reasonable guarantee for social welfare, our allocation algorithms must take ad-specific CTRs into account. As a result, it is impossible to provide a value-stability guarantee in this setting while also providing an approximation to social welfare. Nevertheless, we show that a form of preference-aligned fairness holds. Specifically, let Alice and Bob be two users with multiplicatively similar values and let α and α' denote their ad-specific CTR vectors. Then we show that although the two users' allocations can be quite far from each other, Alice receives a higher allocation than Bob for precisely the ads that she is more likely to click on, and vice versa. Formally, if we sort the advertisers in decreasing order of the ratio α_i/α'_i , then for every i , the probability that Alice gets to see an ad with index $\leq i$ is at least as large as Bob's probability of seeing the same set of ads.

Computing payments. We conclude our study with a discussion of payments. It is easy to observe that both generalized IPA and generalized PA have monotone allocation rules in the advertisers' reported values. However, computing the supporting prices is not straightforward and was left open in previous work. Let $x_i(v_i)$ denote the net allocation (expected probability of click) to advertiser i for a particular user, when the advertiser reports a per-click value of v_i . We show that $x_i(v_i)$ is a piecewise rational function with polynomially many pieces and that it is possible to compute the functional form of each piece in polynomial time. Computing payments using Myerson's lemma then boils down to computing polynomially many integrals over rational functions.

Organization of the paper. We present our extension of the IPA in Section 3 and prove its social welfare and fairness guarantees for the setting of similarly qualified users with similar preferences. In Section 4 we discuss fairness for users that are similarly qualified but have different preferences. Section 5 presents our algorithm for computing payments. We extend our results to the PA in Section 6. Most proofs are deferred to the appendix or removed due to space limitations.²

Related Work

Journalism and empirical work have revealed the myriad ways in which existing ad auction systems lead to unfairness and discrimination [2, 10, 11, 12, 14]. One approach to addressing these issues develops advertiser strategies for bidding in existing auction formats while ensuring statistical parity between groups [9, 15].

More related to our approach is theoretical work on designing auctions and, more generally, algorithms that guarantee fairness properties. These fairness properties typically differ in two dimensions: 1) whether they apply to individuals or only to groups as a whole, and 2) whether they enforce fairness by similarity of treatment or outcome, satisfaction of preferences (e.g., in the form of envy-freeness), or something bridging the two.

These notions of fairness grew out of the fair classification literature, where Dwork *et al.* [7] were the first to propose an individual fairness notion requiring agents who are similar under some task-specific metric to receive similar classifications. Dwork and Ilvento investigate in [6]

²For the full version, visit <https://arxiv.org/abs/2204.04136>.

whether compositions of such classification algorithms that are fair in isolation maintain their fairness properties.

Kim *et al.* [13] introduce individual preference-informed fairness by augmenting this notion of individual fairness with envy-freeness, allowing the allocations of similar users to differ in accordance with their preferences. Similarly, Zafar *et al.* in [18] develop notions of preference-informed group fairness by allowing deviations from parity in treatment and impact if the deviations are envy-free.

Our work employs and expands upon a model of individual fairness in sponsored search first developed by Chawla *et al.* [5] and based on the multi-category fairness work of Dwork and Ilvento [6]. An alternate model, also based on [6], was presented by Watts [16], albeit in a Bayesian setting. A main difference between our work and [16] is that we study the design of auctions that achieve an optimal tradeoff between fairness and welfare, whereas [16] analyzes the fairness and welfare of two specific mechanisms. Another relevant work is that of Essaidi and Weinberg [8] who study the fairness-welfare tradeoff in a Bayesian setting. [8] draws a connection between individual fairness in this context and multi-item auctions with an item symmetry constraint, giving simple mechanisms that achieve a constant-approximation to the revenue-optimal fair mechanism.

There is also some recent work on group-fair ad auctions, such as [17], which shows that constraints on advertiser behavior which enforce group fairness notations can actually increase the profit of the platform. In a Bayesian setting, [3] augments generalized second price auctions with fair division schemes to achieve good social welfare guarantees while satisfying envy-freeness properties among advertiser groups.

As far as we know, ours is the first work addressing fairness specifically in the positional auctions setting where different users have different click through rates.

2 Models and Definitions

We consider the following stylized model for online advertising auctions. Let U be the set of users, n the number of advertisers, and k the number of slots. We use index u for users, i for advertisers and j for slots. At each point in time, a user $u \in U$ arrives. Each advertiser $i \in [n]$ bids a per-click value v_i^u on that user. This is the value the advertiser receives if the user clicks on their ad. Let $\text{CTR}_{i,j}^u$ denote the click through rate of advertiser i in slot j , that is, the probability that the user u will click on the ad i if it is placed in slot j .

A truthful auction decides which ads to display in each of the k slots. The auction receives the vector $v = (v_1^u, \dots, v_n^u)$ as well as the click through rates CTR^u and returns an allocation matrix $\mathbf{a}(v) = [\mathbf{a}_{ij}]_{i \in [n], j \in [k]}$ where $\mathbf{a}_{ij}$ denotes the probability that ad i is displayed in slot j .³ We omit the superscript u whenever it is clear from the context that we are discussing a certain user.

Truthfulness. Given an allocation $\mathbf{a}(v)$ (where the user u is implicit), advertiser i receives a net allocation (expected number of clicks) of $\sum_j \text{CTR}_{i,j}^u \mathbf{a}_{ij}$ and a net expected value of $v_i \cdot \sum_j \text{CTR}_{i,j}^u \mathbf{a}_{ij}$ from the allocation. To ensure truthfulness, there should exist a supporting pricing function $p_i(v)$ for every advertiser i such that bidding truthfully maximizes the advertiser's net expected utility. For such a payment function to exist, it is sufficient and necessary that the allocation probability $\sum_j \text{CTR}_{i,j}^u \mathbf{a}_{ij}$ is monotone non-decreasing in the per-click value v_i . All of the mechanisms we discuss

³We require $\sum_i \mathbf{a}_{ij} = 1$ for all j and $\sum_j \mathbf{a}_{ij} \leq 1$ for all i . Every matrix $\mathbf{a}(\cdot)$ satisfying these matching constraints can be expressed as a distribution over deterministic assignments of ads to slots.

in this paper satisfy monotonicity. In Section 5 we discuss how to compute supporting payments efficiently.

Separable click through rates. Throughout this paper we assume that the click through rates $\text{CTR}_{i,j}^u$ are multiplicatively separable into an advertiser-specific component and a slot-specific component. This is a standard model (see, for example, [1]).

Definition 2.1 (Separable Click Through Rates). *Click through rates are separable if, for every user u , there exists a advertiser dependent vector $\alpha_u = (\alpha_1, \dots, \alpha_n)$ and a slot dependent vector $\beta_u = (\beta_1, \dots, \beta_k)$ in which $\alpha_1, \dots, \alpha_n > 0$ and $1 \geq \beta_1 \geq \beta_2 \geq \dots \geq \beta_k \geq 0$ such that $\text{CTR}_{i,j}^u = \alpha_i \beta_j$ for all $i \in [n]$ and $j \in [k]$.*

Observe that in the separable model the value an advertiser i obtains from slot j is $\alpha_i \beta_j v_i$. Since the slot specific components β_j are common to all advertisers, the relative values of advertisers are given by $\alpha_i v_i$. These relative values are important in the mechanisms we design. We call them the “effective values” of the advertisers:

Definition 2.2 (Effective Value). *The effective value of advertiser i is given by $\hat{v}_i = v_i \alpha_i$.*

We call the above model of online advertising auctions with separable CTRs the **Position Auction Setting**.

Prior-free design. As in previous works, the mechanisms we design and analyze in this paper are prior-free, meaning that the allocation to a user does not depend on the distribution of users or advertisers’ value vectors or the history of users already served. Besides the well-documented benefits of prior-free mechanism design, in the context of fairness we get the added benefit that fairness guarantees hold for all users that are served by the mechanism regardless of whether or not the auctioneer’s model accounts for them.

Definition 2.3 (Scale-Free). *A mechanism is scale-free if it has the property that multiplying the input values by a uniform constant does not change the resulting allocation.*

2.1 Social Welfare

The goal of this work, as in [4, 5], is to achieve a tradeoff between fairness and social welfare for the mechanisms we design. The social welfare of an allocation $\mathbf{a}(v)$ is defined to be the sum of all of the advertisers’ net expected values:

$$SW(\mathbf{a}(v)) = \sum_{i \in [n], j \in [k]} v_i \text{CTR}_{i,j}^u \mathbf{a}_{i,j}$$

We compare this social welfare to the maximum achievable by any feasible allocation. When click through rates are separable, the maximum social welfare is achieved by the allocation that assigns advertisers to slots in decreasing order of $\hat{v}_i$, the effective values. We call the allocation sorted by effective values the UNFAIR-OPT and also use the same term to denote the social welfare of this allocation.

Formally, if π is the order of advertisers where $\hat{v}_{\pi_1} \geq \hat{v}_{\pi_2} \geq \dots \geq \hat{v}_{\pi_n}$, then the (unfair) optimal social welfare is given by:

$$\text{UNFAIR-OPT}(v, \alpha, \beta) = \sum_{j=1}^k \alpha_{\pi_j} v_{\pi_j} \beta_j$$

Since it is generally impossible to achieve optimal social welfare and fairness simultaneously, we look for mechanisms that guarantee our fairness notions while giving a good approximation to the optimal social welfare.

Definition 2.4 (Social Welfare Approximation). *We say mechanism $\mathcal{A}(\cdot)$ achieves an η -approximation to social welfare for $\eta \leq 1$, if for all instances (v, α, β) , we have $SW(\mathcal{A}(v, \alpha, \beta)) \geq \eta \cdot \text{UNFAIR-OPT}(v, \alpha, \beta)$.*

2.2 Fairness

[4] formalized fairness in ad auctions as a value stability condition based on the notion of individual fairness. Individual fairness requires that the auction assign similar allocations to similar users. [4] defined similarity between two users on the basis of closeness between the value vectors assigned to them by the advertisers. Informally speaking, if two users receive similar values from all advertisers, then they should also receive similar allocations. In order for the definition to be scale-free with respect to values, similarity between values is defined in multiplicative terms.

In the context of a single item auction, allocations are probability vectors. Similarity in allocations is therefore defined based on some notion of distance between probability vectors. [4] formalized similarity in terms of the ℓ_∞ distance between the probability vectors whereas [5] used total variation or ℓ_1 distance. We state the value stability definition from [4] below.

Definition 2.5 (Definition 2.1 from [4], Value Stability). *An allocation mechanism $\mathbf{a}(\cdot)$ 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 v and v' :*

$$|\mathbf{a}_i(v) - \mathbf{a}_i(v')| \leq f(\lambda) \text{ for all } i \in [n], \text{ where } \lambda = \max_{i \in [n]} \left(\max \left\{ \frac{v_i}{v'_i}, \frac{v'_i}{v_i} \right\} \right).$$

In this definition, the function f , called the value stability constraint, governs the strength of the value stability condition. We assume f to be non-decreasing, with $f(0) = 0$ and $f(\infty) = 1$. Following [4], we focus on the family of constraints $f_\ell(\lambda) = 1 - \lambda^{-2\ell}$. [4] argue that this family of stability constraints captures the entire spectrum of possible fairness conditions in the context of allocation algorithms.

In order to extend these fairness definitions to the position auctions setting, we need to extend the notion of closeness in allocations to multi-dimensional allocation matrices M as well as extend the notion of closeness in values to click through rates.

Let us consider the latter issue first. A straightforward manner of extending closeness over value vectors to the separable setting is to require that two similar users are assigned similar values, as well as have similar click through rates. But this notion of closeness is too restrictive. Values capture how advertisers perceive users as potential customers; whereas click through rates capture how users perceive the relevance of ads to their needs and how users behave in perusing ads on a search page. Two users that are similarly qualified for a set of ads may nevertheless exhibit very different behavior in responding to ads on a search page. Ideally the fairness guarantees an allocation

algorithm provides should hinge only on the closeness between values v_i and not on the closeness between click through rates $\text{CTR}_{i,j}$. However, in order to obtain good social welfare, allocations necessarily need to depend on the advertiser specific click through rates α_i . We accordingly define closeness between users in terms of their effective values $\alpha_i v_i$ (while ignoring dissimilarity in slot specific CTRs, β). In Section 4 we extend our fairness definitions and guarantees to settings where closeness is defined only in terms of the values v_i , ignoring dissimilarity in α and β .

Let us now consider closeness over probability matrices. We consider three notions. The first is ℓ_∞ distance, the maximum difference of allocations in any one entry (i, j) of the corresponding matrices.

Definition 2.6 (Value Stability for Position Auctions). *An allocation mechanism $\mathcal{A}(\cdot)$ is value stable with respect to function $f : [1, \infty] \rightarrow [0, 1]$ if the following condition is satisfied for every set of value and CTR vectors v, v', α, α' and β :*

$$|M_{i,j} - M'_{i,j}| \leq 2f_\ell(\lambda) \text{ for all } i \in [n], j \in [k] \text{ where } \lambda \text{ is defined as } \max_{i \in [n]} \left(\max \left\{ \frac{\alpha_i v_i}{\alpha'_i v'_i}, \frac{\alpha'_i v'_i}{\alpha_i v_i} \right\} \right)$$

and $M = \mathcal{A}(v, \alpha, \beta)$ and $M' = \mathcal{A}(v', \alpha', \beta)$.

Suppose, as an example, for a particular advertiser i , user u has an allocation of $a = (.1, .1, .1, .1)$. Consider two possible allocation vectors for some v close to u : $a' = (.15, .15, .15, .15)$ and $a'' = (.15, .05, .15, .05)$. In some sense, allocation a' is much more unfair than a'' because in a' the entry-wise differences from a compound while in a'' they offset each other. Weak value stability cannot distinguish these two cases because it is concerned only with the absolute differences. Our next definition, ordered value stability is intended to allow a'' but not a' .

To do this, we bound the absolute differences in the total allocation of an advertiser across all columns, weighted by a vector $h_{i,j}$. This vector represents the utility the first user receives from seeing advertisement i in slot j . Since we assume the slots are in decreasing order of salience, this should be weakly decreasing in j .

Definition 2.7 (Ordered Value Stability for Position Auctions). *An allocation mechanism $\mathcal{A}(\cdot)$ is ordered value stable with respect to function $f : [1, \infty] \rightarrow [0, 1]$ if the following condition is satisfied for every set of value and CTR vectors v, v', α, α' and β , as well as for any advertiser i and any decreasing vector h_i with $1 \geq h_{i,1} \geq \dots \geq h_{i,k} \geq 0$:*

$$\left| \sum_{j=1}^k h_{i,j} (M_{i,j} - M'_{i,j}) \right| \leq f_\ell(\lambda) \text{ where } \lambda \text{ is defined as } \max_{i \in [n]} \left(\max \left\{ \frac{\alpha_i v_i}{\alpha'_i v'_i}, \frac{\alpha'_i v'_i}{\alpha_i v_i} \right\} \right)$$

where $M = \mathcal{A}(v, \alpha, \beta)$ and $M' = \mathcal{A}(v', \alpha', \beta)$.

The previous two definitions are concerned only with a single advertiser. In some instances, however, there are meaningful subsets of advertisers and bounding the differences of the allocations each advertiser individually may not be sufficient to ensure fairness overall. For example, if there are several different ads giving information about registering to vote, the total volume of voter registration ads a user sees is more important from a fairness perspective than the amount they see any particular voter registration ad. Therefore, the last notion we consider is a combination of ℓ_1 and ℓ_∞ distance: we consider, for any subset of advertisers, the total variation distance between the allocations of these advertisers to one slot, and bound the maximum over all slots of this distance.

Definition 2.8 (Total Variation Value Stability for Position Auctions). *A mechanism $\mathcal{A}(\cdot)$ with satisfies total variation value stability with respect to a function $f : [1, \infty] \rightarrow [0, 1]$ if the following condition is satisfied for every set of value and CTR vectors v, v', α, α' and β , as well as every subset of advertisers $S \subseteq [n]$ and for every column j :*

$$\left| \sum_{s \in S} \mathcal{A}(\hat{v})_{s,j} - \sum_{s \in S} \mathcal{A}(\hat{v}')_{s,j} \right| \leq f(\lambda) \text{ where } \lambda \text{ is defined as } \max_{i \in [n]} \left(\max \left\{ \frac{\alpha_i v_i}{\alpha'_i v'_i}, \frac{\alpha'_i v'_i}{\alpha_i v_i} \right\} \right)$$

and where $M = \mathcal{A}(v, \alpha, \beta)$ and $M' = \mathcal{A}(v', \alpha', \beta)$.

3 Inverse Proportional Allocation

In this section, we present a generalization of the mechanism first introduced in [4] as IPA to the position auction setting. We show that the generalization retains a constant approximation to the optimal social welfare and an appropriate generalization of the value stability condition. In Section 3.1 we describe the generalization of the mechanism from $k = 1$ to general k . In Section 3.2 we show that two different value stability conditions hold and in Section 3.3 we show that the exact same guarantee in [4] holds for the generalization as well. Some of the proofs in this section are deferred to Appendix A.

3.1 Generalized IPA

In [4], IPA was presented as a mechanism for the single item auction. An interpretation of this mechanism is as follows: start with an infeasible allocation of 1 unit to each advertiser (for a total allocation of n) and then gradually decrease the allocations until the total allocation reaches 1. The rate of this decrease is determined by a function g of the reported values. The IPA with parameter ℓ uses $g(x) = x^{-\ell}$. [4] also presents an algorithmic interpretation of the mechanism. The following is the generalization of this mechanism to the position auction setting.

First, as a warm-up, we generalize IPA to a special case of the position auction setting where $\beta = \vec{1}$. Our algorithm allocates a total of k units to the advertisers, with each advertiser receiving an allocation $a_i \in [0, 1]$ such that $\sum_i a_i = k$.

We follow the same intuition as for the case of $k = 1$. The mechanism first allocates 1 to each advertiser, then decreases the allocations until the total allocation reaches k rather than 1. See Appendix A for an algorithmic interpretation of this mechanism. Note that setting $k = 1$ gives the exact same mechanism as in [4]. Algorithm 3 is scale free and produces allocations that are non-decreasing in k . Furthermore, the allocation to advertiser i , namely a_i , is non-decreasing in $\hat{v}_i$ and non-increasing in $\hat{v}_{-i}$.

We now extend the k -unit setting to the position auction setting. The resulting allocation algorithm is called Generalized IPA. The algorithm assigns to every slot j a distribution over advertisers given by the difference in the j -unit and $j - 1$ -unit allocations produced by k -unit IPA.

Feasibility. We observe that the allocation produced by the generalized IPA algorithm is feasible. That is, there exists a distribution over matchings from advertisers to slots, for which the total probability that advertiser i is allocated a slot is equal to M .

ALGORITHM 1: Generalized IPA

Input: Vector v of non-negative advertiser bids for user u ; CTRs $\alpha_1, \dots, \alpha_n$ and $\beta_1, \dots, \beta_k$; number of slots k ; function $g : \mathbb{R}^{\geq 0} \rightarrow (0, \infty]$ with $g(0) = \infty$ and $\lim_{x \rightarrow \infty} g(x) = 0$;
for $h \in [k]$ **do**
 | Set $\mathbf{a}^{(h)} \leftarrow$ the output of the IPA k -unit algorithm on input (v, α, h, g)
end
for $j \in [k]$ **do**
 | Set $M_{\cdot, j} = \mathbf{a}^{(j)} - \mathbf{a}^{(j-1)}$
end
return M

Note that the generalized IPA algorithm is scale-free and independent of β .

3.2 Fairness

We now prove the value stability of the Generalized IPA mechanism.

Theorem 3.1. *The Generalized IPA mechanism with parameter $\ell > 0$ and for any number of advertisers n is value stable with respect to any function f satisfying $f(\lambda) \geq f_\ell(\lambda) = 1 - \lambda^{-2\ell}$ for all $\lambda \in [1, \infty)$, as in Definition 2.6.*

Our proof has two parts. First, give a bound on the deviation between allocations given by the k -unit IPA mechanism to similar users. Then, we use the bound to show that Generalized IPA achieves value stability.

Lemma 3.2. *For the k -unit IPA mechanism with parameter ℓ run on any k and any bid vectors v and v' with $\lambda = \max_{i \in [n]} \{\hat{v}_i/\hat{v}'_i, \hat{v}'_i/\hat{v}_i\}$, for all indices i , $|\mathbf{a}_i(v) - \mathbf{a}_i(v')| \leq f_\ell(\lambda)$.*

Next, we show that Generalized IPA satisfies ordered value stability.

Theorem 3.3. *Generalized IPA with parameter ℓ satisfies ordered value stability with respect to $f_\ell(\lambda)$. That is, for every set of value and CTR vectors v, v', α, α' and β , as well as for any advertiser i and any decreasing vector h with $1 \geq h_1 \geq \dots \geq h_k \geq 0$:*

$$\left| \sum_{j=1}^k h_j (M_{i,j} - M'_{i,j}) \right| \leq f_\ell(\lambda) \text{ where } \lambda \text{ is defined as } \max_{i \in [n]} \left(\max \left\{ \frac{\alpha_i v_i}{\alpha'_i v'_i}, \frac{\alpha'_i v'_i}{\alpha_i v_i} \right\} \right)$$

where $M = \mathcal{A}(v, \alpha, \beta)$ and $M' = \mathcal{A}(v', \alpha', \beta)$.

3.3 Social Welfare

We now show that Generalized IPA achieves a good approximation to the optimal social welfare UNFAIR-OPT.

Theorem 3.4. *The IPA algorithm for the separable case, Algorithm 1, run with parameter $\ell > 0$ and any number of advertisers n achieves a $\left(1 - \frac{\ell^\ell}{(1+\ell)^{\ell+1}}\right)$ -approximation the social welfare of the unfair optimum.*

To do so, we first show an approximation result for the special case of $\vec{\beta} = 1$, the k -unit algorithm.

Lemma 3.5. *The IPA algorithm for the k -unit case, Algorithm 3, run with parameter ℓ and any number of advertisers n achieves a $\left(1 - \frac{\ell^\ell}{(1+\ell)^{\ell+1}}\right)$ -approximation to the social welfare of the unfair optimum.*

We use Lemma 3.5 and extend definition of Generalized IPA allocation vector based on k -unit vectors to show Theorem 3.4. The proof is deferred to Appendix A. The approximation factor is $\frac{3}{4}$ at $\ell = 1$ and as $\ell \rightarrow \infty$, the approximation factor goes to 1.

Remark 3.6. *The approximation factor in Lemma 3.5 is tight for IPA mechanism.*

Proof. Consider the following example. Fix a user u and let the bidding vector of the advertisers be:

$$\underbrace{(1, \dots, 1)}_k, \underbrace{(\epsilon, \dots, \epsilon)}_{n-k}$$

where $1 > \epsilon = \frac{-5k + \sqrt{25k^2 - 16(n-k)\frac{k^2}{n-k} - 4 - 4(n-k)}}{8(n-k)} > 0$. Let $\ell = 1$ and $n > 2k$. We get:

$$SW(\text{ALG}) = k\left(1 - \frac{n-k}{(n-k)\epsilon^{-1} + k}\right) + (n-k)\epsilon\left(1 - (n-k)\frac{\epsilon^{-1}}{(n-k)\epsilon^{-1} + k}\right), \quad \text{UNFAIR-OPT} = k$$

For the aforementioned value of ϵ , we will have $\frac{SW(\text{ALG})}{\text{UNFAIR-OPT}} = \frac{3}{4}$. Note that this example fits the maxima point we found in the proof of Lemma 3.5. $\square$

4 Fairness for users with different preferences

So far we have assumed that similar users are similar in all aspects – the values advertisers assign to them as well as the rates at which the users click on different ads. However, these two sets of parameters are asymmetric. Values capture advertisers’ preferences over users whereas CTRs capture users’ preferences over advertisers. We will now distinguish between similarity in *qualification* (i.e. values) from similarity in *user preferences* (i.e. CTR).

A myopic viewpoint might suggest that two users that are similarly qualified should be treated similarly by the auction no matter their preferences. However, this is fundamentally at odds with the objective of maximizing the social welfare⁴ a.k.a. the collective value of the advertisers, as the latter are contingent upon clicks. Consequently, the outcome of the auction cannot be completely independent of user preferences and we look towards a notion of fairness that is appropriately preference aligned.

To motivate our definitions, consider the following example. We have two users Alice and Bob, two advertisers A and B, and a single slot to display an ad. The users look identical to the advertisers: A places a value of \$1 on a click from either user and B places a value of \$10 from either click. However the users behave differently when they view ads. Bob clicks both ads with certainty. Alice clicks A’s ad with certainty but B’s ad with probability only 1%. The platform should clearly display ad A for Alice and ad B for Bob. Although these outcomes are different,

⁴Social welfare is a misnomer in this context, as it does not take into account the benefit or value users derive from viewing the ad.

both users are happy: Bob is essentially indifferent between A and B, while Alice greatly prefers A. In this case, any differences in allocation are aligned with user preferences.

Can we always expect this to be the case? Formally, consider a single slot auction with n advertisers, and two users with identical value vectors $v = v'$. Let $\mathbf{a}$ and $\mathbf{a}'$ denote their respective allocation vectors. Can we ensure that any allocation mass that is moved between advertisers in $\mathbf{a}'$ relative to $\mathbf{a}$ is moved from low CTR advertisers to high CTR advertisers?

Unfortunately, we cannot ensure this property while also maintaining a reasonable approximation for social welfare. To see this, consider the above example with Alice and Bob once again and suppose that Bob's CTR for advertiser B changes to 20%. In order to obtain a good social welfare, the auction must continue to display ad B for Bob. However, now Bob gets to see much more of ad B and much less of ad A than Alice even though he greatly prefers ad A to ad B. The key observation here is that the allocation mass in B's allocation shifts to an advertiser with high *relative* CTR, when measured relative to the CTRs of Alice.

Motivated by this example, we propose the following new preference-aligned definition of fairness for identically valued users. Underlying this definition is a relative ordering of advertisers for two users u and v with advertiser specific CTR vectors $\alpha_u = (\alpha_1^u, \dots, \alpha_n^u)$ and $\alpha_v = (\alpha_1^v, \dots, \alpha_n^v)$. We will assume that advertisers are ordered in (weakly) decreasing order of the ratio α_i^v/α_i^u , and require that allocation mass for user v is shifted from advertisers that appear later in the ordering to those that appear earlier in the ordering.

Definition 4.1 (Value Stability for Identically-Valued Users with Heterogeneous Preferences). *An allocation mechanism $\mathcal{A}(\cdot)$ is value-stable for identical users with heterogeneous preferences if for every pair of users with identical value vectors v ; CTR vectors α, α', β , and β' ; any ordering over advertisers that is weakly decreasing in α/α' ; and for every advertiser $i \in [n]$ and slot $j \in [k]$:*

$$\sum_{t=1}^i \sum_{s=1}^j M_{t,s} \geq \sum_{t=1}^i \sum_{s=1}^j M'_{t,s}, \quad \text{where } M = \mathcal{A}(v, \alpha, \beta) \text{ and } M' = \mathcal{A}(v, \alpha', \beta').$$

Similar users: The above definition extends in a straightforward manner to pairs of users that are similarly rather than identically qualified, and again have different preferences over advertisers as expressed through CTRs. Once again we require that allocation mass shifts from advertisers with low relative CTR to those with higher relative CTR, but we allow for additive errors in allocation that grow with the dissimilarity in the users' values.

Definition 4.2 (Value Stability for Similarly-Valued Users with Heterogeneous Preferences). *An allocation mechanism $\mathcal{A}(\cdot)$ is value-stable for users with heterogeneous preferences with respect to function $f_\ell : [1, \infty] \rightarrow [0, 1]$ if for every pair of users with value vectors v and v' ; CTR vectors α, α', β , and β' ; any ordering over advertisers that is weakly decreasing in α/α' ; and for every advertiser $i \in [n]$ and slot $j \in [k]$:*

$$\sum_{t=1}^i \sum_{s=1}^j M_{t,s} \geq \sum_{t=1}^i \sum_{s=1}^j M'_{t,s} - i f_\ell(\lambda)$$

where $M = \mathcal{A}(v, \alpha, \beta)$, $M' = \mathcal{A}(v', \alpha', \beta')$ and $\lambda = \max_{i \in [n]} \left\{ \max \left\{ \frac{v_i}{v'_i}, \frac{v'_i}{v_i} \right\} \right\}$.

Comparing Definition 4.1 and Definition 4.2, note that if $v = v'$ then $\lambda = 1$ and, as discussed in [4], a proper f function has the property of $f(1) = 0$. Therefore, Definition 4.1 is exactly Definition 4.2 in the special case of $v = v'$.

4.1 Fairness of IPA and PA for heterogeneous users

We show that both the Generalized IPA and Generalized PA mechanisms satisfy Definition 4.1 and more generally Definition 4.2.

To begin, we show that any mechanism for the k -unit case satisfying certain mild conditions also satisfies Definition 4.1. Both k -unit IPA and k -unit PA satisfy these conditions and hence are value-stable for identically qualified users with heterogeneous preferences.

Lemma 4.1. *Let $\mathbf{a}(v)$ be a scale-free k -unit allocation algorithm such that $\mathbf{a}_i(v)$ is weakly increasing in v_i . Suppose further that for all $t \neq i$, $\mathbf{a}_i(v)$ is weakly decreasing in v_t . Then $\mathbf{a}(v)$ satisfies Definition 4.1.*

Proof. Fix i and scale α' so that $\alpha_i = \alpha'_i$. Since the advertisers are sorted, we now know that for all $t < i$, $\alpha_t \geq \alpha'_t$ and for all $t > i$, $\alpha_t \leq \alpha'_t$.

We proceed by two cases and then use a transitivity argument to show the theorem holds in general.

Consider the case where for all $t \leq i$, $\alpha_t = \alpha'_t$. Then $\alpha v \begin{cases} = \alpha' v \text{ for all } t \leq i \\ \leq \alpha' v \text{ for all } t > i \end{cases}$.

Therefore, since the allocation a_t is weakly decreasing in v_s for all $s \neq t$, we have that for all $t \leq i$, $a(\alpha v) \geq a(\alpha' v)$. Hence, $\sum_{t=1}^i a_t(\alpha v) \geq \sum_{t=1}^i a_t(\alpha' v)$, as desired.

Now, consider the case where for all $t \geq i$, $\alpha_t = \alpha'_t$. Then $\alpha v \begin{cases} \geq \alpha' v \text{ for all } t < i \\ = \alpha' v \text{ for all } t \geq i \end{cases}$.

Therefore, since the allocation a_t is weakly decreasing in v_s for all $s \neq t$, we have that for all $t > i$, $a(\alpha v) \leq a(\alpha' v)$ and hence $\sum_{t=i+1}^n a_t(\alpha v) \leq \sum_{t=i+1}^n a_t(\alpha' v)$. But $\sum_{t=1}^i a_t(\alpha v) = k - \sum_{t=i+1}^n a_t(\alpha v)$ and likewise $\sum_{t=1}^i a_t(\alpha' v) = k - \sum_{t=i+1}^n a_t(\alpha' v)$. Therefore, $\sum_{t=i+1}^n a_t(\alpha v) \leq \sum_{t=i+1}^n a_t(\alpha' v)$ implies $\sum_{t=1}^i a_t(\alpha v) \geq \sum_{t=1}^i a_t(\alpha' v)$, as desired.

We now argue that the theorem holds in general. Let $\alpha''_t := \begin{cases} \alpha_t & \text{if } t \leq i \\ \alpha'_t & \text{if } t > i \end{cases}$. By the first case, $\sum_{t=1}^i a_t(\alpha v) \geq \sum_{t=1}^i a_t(\alpha'' v)$, and by the second case $\sum_{t=i+1}^n a_t(\alpha'' v) \geq \sum_{t=i+1}^n a_t(\alpha' v)$. Hence, $\sum_{t=1}^i a_t(\alpha v) \geq \sum_{t=1}^i a_t(\alpha' v)$, as desired. $\square$

Corollary 4.2. *The k -unit IPA and k -unit PA mechanisms satisfy Definition 4.1.*

Because our generalized mechanisms are defined in terms of telescoping differences of the k -unit allocations, Theorem 4.3 follows directly from Corollary 4.2.

Theorem 4.3. *The Generalized IPA and Generalized PA mechanisms satisfy Definition 4.1.*

Next, we show Generalized IPA and Generalized PA are value-stable for similarly-valued users with heterogeneous preferences. The only thing changing from Definition 4.1 to Definition 4.2 is that we need to keep track of small changes between the two allocations, which leads to the following theorem. The proof is deferred to Appendix B.

Theorem 4.4. *The Generalized IPA and Generalized PA mechanisms $\mathcal{A}(\cdot)$ with parameter ℓ are value-stable for similarly-valued users with heterogeneous preferences.*

5 Computing payments

In this section we develop an algorithm for computing supporting payments for the generalized IPA and generalized PA allocation rules. Our main observation is that the allocation functions of IPA and PA are piecewise rational functions with polynomially many pieces where each piece can be computed in polynomial time. With these pieces in hand, and using Myerson's lemma, computing payments amounts to computing polynomially many integrals of rational functions.

We focus on the generalized IPA; the argument for generalized PA is similar. Formally, for a fixed and implicit user u , and a fixed and implicit advertiser i , let $x_i(v)$ denote the net allocation to the advertiser, a.k.a. the expected number of clicks the advertiser receives from the user. If the user is assigned allocation $M = \mathcal{A}(v, \alpha, \beta)$ then we have $x_i(v) = \sum_j M_{i,j} \alpha_i \beta_j$. Let $a^{(j)}$ denote the cumulative allocation to the user in the first j slots as in the description of Algorithm 2 and recall that $M_{i,j} = a_i^{(j)} - a_i^{(j-1)}$. Accordingly we get:

$$x_i(v) = \alpha_i \sum_j a_i^{(j)} (\beta_j - \beta_{j+1}) \quad (1)$$

In other words, $x_i(v)$ is a linear combination of the functions $a_i^{(j)}(v)$.

We will now argue that for all i, j , the function $a_i^{(j)}(v)$, as defined in Algorithm 1, is piecewise rational in v_i . Consider the following equivalent formulation of Algorithm 1. Given the values $v_1, \dots, v_n$, ad-specific CTRs $\alpha_1, \alpha_2, \dots, \alpha_n$, and decreasing function g , we find a parameter t such that

$$\sum_{i'} \min(1, t \cdot g(\alpha_{i'} v_{i'})) = n - j \quad (2)$$

The allocation $a_i^{(j)}$ is then given by $1 - \min(1, t \cdot g(\alpha_i v_i))$.

Suppose without loss of generality that i receives a non-zero allocation at value v_i (otherwise $a_i^{(j)}$ is trivially piecewise rational at values $\leq v_i$). We can then rewrite Equation (2) as:

$$t \cdot g(\alpha_i v_i) + \sum_{i' \neq i} \min(1, t \cdot g(\alpha_{i'} v_{i'})) = n - j \quad (3)$$

Now, the expression $\sum_{i' \neq i} \min(1, t g(\alpha_{i'} v_{i'}))$ is independent of v_i and piecewise linear in t with at most n pieces. Given the values v_{-i} and CTRs α_{-i} , we can efficiently compute the linear pieces in this function. Substituting any particular linear piece with t in the range $[t_1, t_2]$ in Equation (3) then gives us an equation of the following form with appropriate parameters x and y :

$$t \cdot g(\alpha_i v_i) + xt = y$$

leading to the solution

$$a_i^{(j)}(v_i) = 1 - g(\alpha_i v_i) \cdot \frac{y}{g(\alpha_i v_i) + x} \quad \text{for } v_i \in \left[\frac{1}{\alpha_i} g^{-1} \left(\frac{y - xt_2}{t_2} \right), \frac{1}{\alpha_i} g^{-1} \left(\frac{y - xt_1}{t_1} \right) \right].$$

Observe that the RHS in the above equation is a rational function as the function g in the definition of IPA is also rational.

Summarizing, we first compute the piecewise rational form of the function $a_i^{(j)}(v_i)$ for all slots j . Each of these functions has at most n pieces. We then use Equation (1) to express $x_i(v_i)$ as a piecewise rational function with at most nk pieces. Finally, we use Myerson's lemma and compute per-impression payments as

$$p_i(v_i) = v_i x_i(v_i) - \int_{z=0}^{v_i} x_i(z) dz.$$

6 Proportional Allocation

In this section, we present a generalization of the mechanism first introduced in [5] as Proportional Allocation (PA) to the position auction setting. We show that the generalization retains the same approximation ratio to the optimal social welfare and an appropriate generalization of the total variation value stability condition. This is a stronger fairness guarantee than that of Generalized IPA, but comes at the cost of a weaker approximation to the optimal social welfare. For a detailed discussion of the trade-offs between the single-unit versions these methods, see [4]. Some of the proofs in this section are deferred to Appendix C.

6.1 Generalized PA

In contrast to IPA, PA can be thought of as initially assigning each advertiser an allocation of 0 and then increasing the allocations in proportion to (some function of) the bid amounts until the total allocation reaches 1. [5] analyzes this mechanism for the single unit case. In particular, they prove value stability with respect to the total variation distance on the allocations, rather than with respect to the ℓ^∞ distance as with IPA. However, in exchange, the social welfare approximation achieved by PA degrades as the number of advertisers increases.

Just like the previous section, we start with a warm-up case in which we consider a special case of position auction where $\beta = \vec{1}$. For this case, we will attempt to allocate proportionally, assigning $k \cdot \frac{g(v_i)}{\sum_t g(v_t)}$ to each bidder i . If this allocation is more than 1 for any advertiser, we cap their allocation at 1 and divide the additional mass proportionally among the remaining advertisers. See Algorithm 4 in Appendix C for an algorithmic interpretation of this mechanism. Note that the function g in this mechanism is different than the one in Section 3, as it is a continuous, super-additive and increasing function.

The extension of this algorithm to the position auction case is similar to the extension we saw in Section 3 for IPA, and works as follows:

ALGORITHM 2: Generalized PA

Input: Vector v of non-negative advertiser bids for user u ; CTRs $\alpha_1, \dots, \alpha_n$ and $\beta_1, \dots, \beta_k$; number of slots k ; function $g : \mathbb{R}^{\geq 0} \rightarrow [0, \infty]$ with g a continuous, super-additive, increasing function and $g(0) = 0$;
for $h \in [k]$ **do**
 | Set $p^{(h)} \leftarrow$ the output of the PA k -unit algorithm on input (v, α, h, g)
end
for $j \in [k]$ **do**
 | Set $P_{\cdot, j} = p^{(j)} - p^{(j-1)}$
end
return P

Observe that Generalized PA is scale-free, independent of β , and produces feasible allocations by essentially the same argument as in ??.

6.2 Fairness

First, we prove the fairness guarantees of our mechanism. We begin by showing the total variation value stability of PA, which as we've discussed is the main advantage of PA over IPA.

Theorem 6.1. *The Generalized PA mechanism with parameter $g(x) = x^\ell$ satisfies definition 2.8 Total Variation Value Stability for Position Auctions with respect to $f_\ell(\lambda)$. That is, for all pairs of effective value vectors $\hat{v}, \hat{v}'$, subsets of advertisers $S \subseteq [n]$, and slots j ,*

$$\left| \sum_{s \in S} P_{s,j}(\hat{v}) - \sum_{s \in S} P_{s,j}(\hat{v}') \right| \leq 2f_\ell(\lambda)$$

The proof of Theorem 6.1 uses the following key lemma, which shows a similar property holds for k -unit PA mechanism.

Lemma 6.2. *The k -unit PA mechanism with parameter $g(x) = x^\ell$ satisfies the property that, for all pairs of effective value vectors $\hat{v}, \hat{v}'$ and subsets of advertisers $S \subseteq [n]$,*

$$\left| \sum_{s \in S} a_s(\hat{v}) - \sum_{s \in S} a_s(\hat{v}') \right| \leq \frac{\lambda^\ell - 1}{\lambda^\ell + 1} \leq f_\ell(\lambda).$$

We now show that Generalized PA also satisfies the same ordered value stability property as IPA. The proof is essentially identical as the proof of Theorem 3.3 except in that it uses the total variation value stability of PA instead the value stability of IPA. For the full proof, see Appendix C.

Theorem 6.3. *Generalized PA with parameter ℓ satisfies ordered value stability with respect to $f_\ell(\lambda)$. That is, for every set of value and CTR vectors v, v', α, α' and β , as well as for any advertiser i and any decreasing vector h with $1 \geq h_1 \geq \dots \geq h_k \geq 0$:*

$$\left| \sum_{j=1}^k h_j (P_{i,j} - P'_{i,j}) \right| \leq f_\ell(\lambda) \text{ where } \lambda \text{ is defined as } \max_{i \in [n]} \left(\max \left\{ \frac{\alpha_i v_i}{\alpha'_i v'_i}, \frac{\alpha'_i v'_i}{\alpha_i v_i} \right\} \right)$$

where $P = \mathcal{A}(v, \alpha, \beta)$ and $P' = \mathcal{A}(v', \alpha', \beta)$.

6.3 Social Welfare

Finally, we give our guarantee on the social welfare approximation ratio achieved by Generalized PA relative to UNFAIR-OPT. The proof relies on a lemma showing the same approximation result for the special case of $\vec{\beta} = 1$, k -unit PA.

Theorem 6.4. *The Generalized PA mechanism with parameter ℓ achieves a $\left(\frac{n-k}{n}(n-k)^{-1/\ell} + 1/n\right)$ -approximation to the optimal social welfare for any instance with n advertisers and k slots.*

Lemma 6.5. *The k -unit PA subroutine with parameter ℓ achieves a $\left(\frac{n-k}{n}(n-k)^{-1/\ell} + 1/n\right)$ -approximation to the optimal social welfare for any instance with n advertisers and k slots.*

References

- [1] Gagan Aggarwal, Ashish Goel, and Rajeev Motwani. Truthful auctions for pricing search keywords. In *Proceedings of the 7th ACM Conference on Electronic Commerce, EC '06*, page 1–7, New York, NY, USA, 2006. Association for Computing Machinery. ISBN 1595932364. doi: 10.1145/1134707.1134708. URL <https://doi.org/10.1145/1134707.1134708>. 5
- [2] 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. *Proc. ACM Hum.-Comput. Interact.*, 3(CSCW), nov 2019. doi: 10.1145/3359301. URL <https://doi.org/10.1145/3359301>. 3
- [3] Andrea Celli, Riccardo Colini-Baldeschi, and Stefano Leonardi. Learning fair equilibria in sponsored search auctions. *arXiv preprint arXiv:2107.08271*, 2021. 4
- [4] Shuchi Chawla and Meena Jagadeesan. Individual fairness in advertising auctions through inverse proportionality. In *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2022. 1, 5, 6, 8, 11, 14
- [5] 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*, page 348–358, New York, NY, USA, 2020. Association for Computing Machinery. 1, 4, 5, 6, 14, 21
- [6] 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. 3, 4
- [7] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference, ITCS '12*, page 214–226, New York, NY, USA, 2012. Association for Computing Machinery. ISBN 9781450311151. doi: 10.1145/2090236.2090255. URL <https://doi.org/10.1145/2090236.2090255>. 3
- [8] Meryem Essaidi and S Matthew Weinberg. On symmetries in multi-dimensional mechanism design. In *International Conference on Web and Internet Economics*, pages 59–75. Springer, 2021. 4
- [9] Lodewijk Gelauff, Ashish Goel, Kamesh Munagala, and Sravya Yandamuri. Advertising for demographically fair outcomes. *arXiv preprint arXiv:2006.03983*, 2020. 3
- [10] Julia Angwin and Terry Perris. Facebook lets advertisers exclude users by race. 2016. URL <https://www.propublica.org/article/facebook-lets-advertisers-exclude-users-by-race>. 3
- [11] Julia Angwin, Noam Scheiber, and Ariana Tobin. Facebook job ads raise concerns about age discrimination. 2017. URL <https://www.nytimes.com/2017/12/20/business/facebook-job-ads.html>. 3

- [12] Katie Benner, Glenn Thrush, and Mike Isaac. Facebook engages in housing discrimination with its ad practices, u.s. says. 2019. URL <https://www.nytimes.com/2019/03/28/us/politics/facebook-housing-discrimination.html>. 3
- [13] 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. 4
- [14] A Lambrecht and C E Tucker. Algorithmic bias? an empirical study of apparent gender-based discrimination in the display of stem career ads. *Management Science*, 65(7):2966–2981, July 2019. URL <https://lbsresearch.london.edu/id/eprint/967/>. © 2019 INFORMS This manuscript has been accepted for publication in Management Science. The version of record can be found at <https://doi.org/10.1287/mnsc.2018.3093>. 3
- [15] 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*, page 337–347, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450369367. doi: 10.1145/3351095.3375783. URL <https://doi.org/10.1145/3351095.3375783>. 3
- [16] Alison Watts. Fairness and efficiency in online advertising mechanisms. *Games*, 12(2):36, 2021. 4
- [17] Di Yuan, Manmohan Aseri, and Tridas Mukhopadhyay. Is fair advertising good for platforms? *Available at SSRN 3913739*, 2021. 4
- [18] Muhammad Bilal Zafar, Isabel Valera, Manuel Rodriguez, Krishna Gummadi, and Adrian Weller. From parity to preference-based notions of fairness in classification. *Advances in Neural Information Processing Systems*, 30, 2017. 4

A Deferred Proofs from Section 3

Below is the algorithmic description of position auction in the case of $\vec{\beta} = 1$:

ALGORITHM 3: k -unit IPA

Input: Vector v of non-negative advertiser bids for user u ; ad-specific CTRs $\alpha_1, \dots, \alpha_n$; number of slots k ; function $g : \mathbb{R}^{\geq 0} \rightarrow (0, \infty]$ with $g(0) = \infty$ and $\lim_{x \rightarrow \infty} g(x) = 0$;

Initialization: Determine effective values, $\hat{v}_i = v_i \alpha_i$ for all i ;

WLOG assume $\hat{v}_1 \geq \dots \geq \hat{v}_n$;

if $k = 0$ **then**

return $\mathbf{a}(v) = \vec{0}$

end

if $\hat{v}_1 \leq 0$ **then**

 Set $\mathbf{a}_i = \frac{k}{n}$ for all $i \in [n]$, **return** $\mathbf{a}(v)$;

end

Set $s \leftarrow \max\{i \in [n] : \hat{v}_i > 0\}$;

while $(s - k)g(\hat{v}_s) \geq \sum_{i=1}^s g(\hat{v}_i)$ **do**

$s \leftarrow s - 1$;

end

For $i > s$: set $\mathbf{a}_i = 0$;

For $i \leq s$ set $\mathbf{a}_i = 1 - (s - k) \frac{g(\hat{v}_i)}{\sum_{t=1}^s g(\hat{v}_t)}$;

return $\mathbf{a}(v)$

Theorem 3.3. Generalized IPA with parameter ℓ satisfies ordered value stability with respect to $f_\ell(\lambda)$. That is, for every set of value and CTR vectors v, v', α, α' and β , as well as for any advertiser i and any decreasing vector h with $1 \geq h_1 \geq \dots \geq h_k \geq 0$:

$$\left| \sum_{j=1}^k h_j (M_{i,j} - M'_{i,j}) \right| \leq f_\ell(\lambda) \text{ where } \lambda \text{ is defined as } \max_{i \in [n]} \left(\max \left\{ \frac{\alpha_i v_i}{\alpha'_i v'_i}, \frac{\alpha'_i v'_i}{\alpha_i v_i} \right\} \right)$$

where $M = \mathcal{A}(v, \alpha, \beta)$ and $M' = \mathcal{A}(v', \alpha', \beta)$.

Proof. Fix some vectors v, v', α, α' , and β , and the corresponding allocation matrices M and M' . Consider some advertiser i . We begin by using the definition Generalized IPA and then rearranging terms. Note that we define $h_{k+1} := 0$ for notational simplicity.

$$\begin{aligned} \left| \sum_{j=1}^k h_j (M_{i,j} - M'_{i,j}) \right| &= \left| \sum_{j=1}^k h_j \left((\mathbf{a}_i^{(j)} - \mathbf{a}_i^{(j-1)}) - (\mathbf{a}_{i'}^{(j)} - \mathbf{a}_{i'}^{(j-1)}) \right) \right| \\ &= \left| \sum_{j=1}^k \left(h_j (\mathbf{a}_i^{(j)} - \mathbf{a}_i^{(j-1)}) - h_j (\mathbf{a}_{i'}^{(j)} - \mathbf{a}_{i'}^{(j-1)}) \right) \right| = \left| \sum_{j=1}^k \left(\mathbf{a}_i^{(j)} - \mathbf{a}_{i'}^{(j-1)} \right) (h_j - h_{j+1}) \right| \end{aligned}$$

Now, observe that because $h_1 \leq 1$ and the coefficients $(h_j - h_{j+1})$ telescope, the sum of these coefficients is at most 1. Since the expression is a weighted sum over columns of the differences in allocation at that column, the expression is bounded by the maximum difference in any column.

But because Generalized IPA satisfies value stability (by Lemma 3.2), this is bounded by $f_\ell(\lambda)$, as desired.

$$\left| \sum_{j=1}^k h_j (M_{i,j} - M'_{i,j}) \right| = \left| \sum_{j=1}^k \left(\mathbf{a}_i^{(j)} - \mathbf{a}_{i'}^{(j-1)} \right) (h_j - h_{j+1}) \right| = \left| \max_j \left(\mathbf{a}_i^{(j)} - \mathbf{a}_{i'}^{(j-1)} \right) \right| \leq f_\ell(\lambda)$$

□

Theorem 3.4. The IPA algorithm for the separable case, Algorithm 1, run with parameter $\ell > 0$ and any number of advertisers n achieves a $\left(1 - \frac{\ell^\ell}{(1+\ell)^{\ell+1}}\right)$ -approximation the social welfare of the unfair optimum.

Proof. Suppose the k -unit IPA mechanism attains an η approximation to the optimal social welfare in the k -unit setting. Then the Generalized IPA mechanism attains the same approximation factor η in the position auction when run with the k -unit IPA mechanism as a subroutine. In order to prove this, we consider the social welfare attained by the Generalized IPA mechanism. Since $\beta_{k+1} = 0$ and $\mathbf{a}_i^{(0)} = \vec{0}$,

$$SW(\text{ALG}) = \sum_{i=1}^n \sum_{j=1}^k \alpha_i v_i \beta_j M_{ij} = \sum_{i=1}^n \sum_{j=1}^k \hat{v}_i \beta_j \left[\mathbf{a}_i^{(j)} - \mathbf{a}_i^{(j-1)} \right] = \sum_{i=1}^n \sum_{j=1}^k \hat{v}_i (\beta_j - \beta_{j+1}) \mathbf{a}_i^{(j)}$$

Since for all $j \in [k]$, $\sum_i \hat{v}_i \mathbf{a}_i^{(j)} \geq \eta(\hat{v}_1 + \dots + \hat{v}_j)$, then:

$$SW(\text{ALG}) = \sum_{j=1}^k (\beta_j - \beta_{j+1}) \left(\sum_{i=1}^n \hat{v}_i \mathbf{a}_i^{(j)} \right) \geq \eta \sum_{j=1}^k (\beta_j - \beta_{j+1}) (\hat{v}_1 + \dots + \hat{v}_j) = \eta \sum_{j=1}^k \hat{v}_j \beta_j = \eta \text{UNFAIR-OPT}$$

Finally, we know by Lemma 3.5 that the k -unit IPA mechanism is an $\eta = \left(1 - \frac{\ell^\ell}{(1+\ell)^{\ell+1}}\right)$ -approximation to the optimal k -unit social welfare. Replacing η by $\left(1 - \frac{\ell^\ell}{(1+\ell)^{\ell+1}}\right)$ concludes the proof. □

B Deferred Proofs from Section 4

Theorem 4.4 The Generalized IPA and Generalized PA mechanisms $\mathcal{A}(\cdot)$ with parameter ℓ are value-stable for similar users with heterogeneous preferences.

Proof. Fix users with user-dependent CTR vectors α and α' and value vectors v and v' . Also fix slot-dependent CTR vector β , advertiser i , and column j . Let $M = \mathcal{A}(v, \alpha, \beta)$ and $M' = \mathcal{A}(v, \alpha', \beta)$, where $\mathbf{a}_t = \sum_{s=1}^j M_{t,s}$, and $\mathbf{a}'_t = \sum_{s=1}^j M'_{t,s}$. Finally, fix a permutation π on advertisers for which $\frac{\alpha_{\pi_1}}{\alpha'_{\pi_1}} \geq \dots \geq \frac{\alpha_{\pi_n}}{\alpha'_{\pi_n}}$.

Since $\mathcal{A}(\cdot)$ is envy-free, we know that $\sum_{s=1}^i \sum_{t=1}^j M_{st}(\alpha v) \geq \sum_{s=1}^i \sum_{t=1}^j M_{st}(\alpha' v)$. Therefore, it suffices to show $\sum_{s=1}^i \sum_{t=1}^j M_{st}(\alpha' v) \geq \sum_{s=1}^i \sum_{t=1}^j M_{st}(\alpha' v') - if(\lambda)$. Consider the difference $\sum_{s=1}^i \sum_{t=1}^j M_{st}(\alpha' v') - \sum_{s=1}^i \sum_{t=1}^j M_{st}(\alpha' v)$. Since $\sum_{t=1}^j M_{st}(\alpha v) = a_s^j(\alpha v)$, we can simplify this to:

$$\begin{aligned}
\sum_{s=1}^i a_s^j(\alpha'v') - \sum_{s=1}^i a_s^j(\alpha'v) &\leq \left| \sum_{s=1}^i a_s^j(\alpha'v') - \sum_{s=1}^i a_s^j(\alpha'v) \right| = \left| \sum_{s=1}^i a_s^j(\alpha'v') - a_s^j(\alpha'v) \right| \\
&\leq \sum_{s=1}^i |a_s^j(\alpha'v') - a_s^j(\alpha'v)| \leq \sum_{s=1}^i f(\lambda) = i * f(\lambda)
\end{aligned}$$

Simply combining this with the previous inequality gives the desired result. $\square$

C Deferred Proofs from Section 6

Below is the algorithmic description of position auction in the case of $\vec{\beta} = 1$:

ALGORITHM 4: k -unit PA

Input: Vector v of non-negative advertiser bids for user u ; ad-specific CTRs $\alpha_1, \dots, \alpha_n$; number of slots k ; function $g : \mathbb{R}^{\geq 0} \rightarrow [0, \infty]$ with g a continuous, super-additive, increasing function and $g(0) = 0$;

Initialization: Determine effective values, WLOG assume $\hat{v}_1 \geq \dots \geq \hat{v}_n$;

if $k = 0$ **then**

return $p(v) = \vec{0}$

end

if $\hat{v}_1 \leq 0$ **then**

 Set $\mathbf{p}_i = \frac{k}{n}$ for all $i \in [n]$, **return** $p(v)$;

end

Set $s \leftarrow \max\{i \in [n] : \hat{v}_i > 0\}$;

Set $r = 1$;

while $\frac{k \cdot g(\hat{v}_r)}{\sum_{t=r}^s g(\hat{v}_t)} \geq 1$ **do**

$\mathbf{p}_r = 1$;

$r \leftarrow r + 1$;

end

For $i \geq r$: set $\mathbf{p}_i = \frac{(k-r) \cdot g(\hat{v}_i)}{\sum_{t=r}^s g(\hat{v}_t)}$;

return $p(v)$;

Theorem 6.3. Generalized PA with parameter ℓ satisfies ordered value stability with respect to $f_\ell(\lambda)$. That is, for every set of value and CTR vectors v, v', α, α' and β , as well as for any advertiser i and any decreasing vector h with $1 \geq h_1 \geq \dots \geq h_k \geq 0$:

$$\left| \sum_{j=1}^k h_j (P_{i,j} - P'_{i,j}) \right| \leq f_\ell(\lambda) \text{ where } \lambda \text{ is defined as } \max_{i \in [n]} \left(\max \left\{ \frac{\alpha_i v_i}{\alpha'_i v'_i}, \frac{\alpha'_i v'_i}{\alpha_i v_i} \right\} \right)$$

where $P = \mathcal{A}(v, \alpha, \beta)$ and $P' = \mathcal{A}(v', \alpha', \beta)$.

Proof. Fix some vectors v, v', α, α' , and β , and the corresponding allocation matrices P and P' . Consider some advertiser i . We begin by using the definition of Generalized PA and then rearranging terms. Note that we define $h_{k+1} := 0$ for notational simplicity.

$$\begin{aligned} & \left| \sum_{j=1}^k h_j (P_{i,j} - P'_{i,j}) \right| = \left| \sum_{j=1}^k h_j \left((p_i^{(j)}(\hat{v}) - p_i^{(j-1)}(\hat{v})) - (p_i^{(j)}(\hat{v}') - p_i^{(j-1)}(\hat{v}')) \right) \right| \\ &= \left| \sum_{j=1}^k \left(h_j (p_i^{(j)}(\hat{v}) - p_i^{(j-1)}(\hat{v})) - h_j (p_i^{(j)}(\hat{v}') - p_i^{(j-1)}(\hat{v}')) \right) \right| = \left| \sum_{j=1}^k \left(p_i^{(j)}(\hat{v}) - p_i^{(j-1)}(\hat{v}') \right) (h_j - h_{j+1}) \right| \end{aligned}$$

Now, observe that because $h_1 \leq 1$ and the coefficients $(h_j - h_{j+1})$ telescope, the sum of these coefficients is at most 1. Since the expression is a weighted sum over columns of the differences in allocation at that column, the expression is bounded by the maximum difference in any column. But because Generalized PA satisfies total variation value stability (by Lemma 6.2), this is bounded by $f_\ell(\lambda)$ for all subsets of advertisers, including the singleton i , as desired.

$$\left| \sum_{j=1}^k h_j (P_{i,j} - P'_{i,j}) \right| = \left| \sum_{j=1}^k \left(p_i^{(j)}(\hat{v}) - p_i^{(j-1)}(\hat{v}') \right) (h_j - h_{j+1}) \right| = \left| \max_j \left(p_i^{(j)}(\hat{v}) - p_i^{(j-1)}(\hat{v}') \right) \right| \leq f_\ell(\lambda)$$

□

Lemma 6.2 The k -unit PA mechanism with parameter $g(x) = x^\ell$ satisfies the property that, for all pairs of effective value vectors $\hat{v}, \hat{v}'$ and subsets of advertisers $S \subseteq [n]$,

$$\left| \sum_{s \in S} a_s(\hat{v}) - \sum_{s \in S} a_s(\hat{v}') \right| \leq \frac{\lambda^\ell - 1}{\lambda^\ell + 1} \leq f_\ell(\lambda).$$

Proof. Fix some pairs of effective value vectors $\hat{v}, \hat{v}'$ and a subset of advertisers $S \subseteq [n]$. Define E to be $\sum_{s \in S} a_s(\hat{v}) - \sum_{s \in S} a_s(\hat{v}')$ and assume without loss of generality that $E \geq 0$. We want to upper bound E by $f_\ell(\lambda)$.

First, we reduce the general case to that where the while loop never executes. That is, we modify the given instance so that the while loop never executes while only increasing E and decreasing λ . First, we can assume that $i \in S$ if $a_i(\hat{v}) > a_i(\hat{v}')$ and $i \notin S$ if $a_i(\hat{v}) < a_i(\hat{v}')$, since that those choices maximize E (and do not effect λ). We also assume that for all i , $\hat{v}_i \geq \hat{v}'_i$ and therefore $\lambda = \max_i \{\hat{v}_i / \hat{v}'_i\}$. If this is violated for $i \in S$, then raising $\hat{v}_i$ to $\hat{v}'_i$ cannot decrease E (it can only increase $\sum_{s \in S} a_s(\hat{v})$) and cannot increase λ . Similarly, if the assumption violated for $i \notin S$, then lowering $\hat{v}'_i$ to $\hat{v}_i$ cannot decrease E (it can only decrease $\sum_{s \in S} a_s(\hat{v}')$) and cannot increase λ .

Now, suppose there exists some $i \in S$ such that $\frac{k \cdot g(\hat{v}_i)}{\sum_{t=s}^n g(\hat{v}_t)} > 1$. Then we can reduce $\hat{v}_i$ so that $\frac{k \cdot g(\hat{v}_i)}{\sum_{t=s}^n g(\hat{v}_t)} = 1$ since this doesn't change E but potentially decreases λ . Finally, suppose there exists some $i \in S$ such that $\frac{k \cdot g(\hat{v}'_i)}{\sum_{t=s}^n g(\hat{v}'_t)} > 1$. Then consider lowering $\hat{v}'_i$ so that $\frac{k \cdot g(\hat{v}'_i)}{\sum_{t=s}^n g(\hat{v}'_t)} = 1$ and then scaling $\hat{v}''$ so that $\hat{v}'_i$ has its original value. This does not change E and potentially decreases λ . Therefore, we've successfully reduced to an instance in which the while loop never executes.

We now assume without loss of generality that the while loop never executes. The remaining argument follows closely from [5].

Define $\alpha := \sum_{s \in S} g(\hat{v}_s)$ and $\beta := \sum_{s \notin S} g(\hat{v}_s)$, and define α' and β' analogously. Note now that the while loop never executes, we have that for all i , $a_i(\hat{v}) = g(\hat{v}_i) / \sum_{i=1}^n g(\hat{v}_s)$, and similarly for $a_i(\hat{v}')$. Therefore we can write

$$E = \frac{\alpha}{\alpha + \beta} - \frac{\alpha'}{\alpha' + \beta'} = 1 - \frac{\beta}{\alpha + \beta} - \frac{\alpha'}{\alpha' + \beta'}.$$

Let $R_\alpha := \alpha/\alpha'$ and $R_\beta := \beta'/\beta$. Note that $R_\alpha \leq g(\lambda)$ because for any $s \in S$, $\hat{v}_s/\hat{v}'_s \leq \lambda$ so $g(\hat{v}_s)/g(\hat{v}'_s) \leq g(\lambda)$. Similarly, $R_\beta \leq g(\lambda)$. Observe also that our expression for E can be upper bounded by the case that these inequalities for R_α and R_β are tight.

$$\begin{aligned} E &\leq 1 - \frac{\alpha \cdot g(\lambda)}{\alpha \cdot g(\lambda) + \beta'} - \frac{\alpha}{\alpha + \beta \cdot g(\lambda)} = \frac{\alpha\beta'(g(\lambda)^2 - 1)}{(\alpha + \beta'g(\lambda))(g(\lambda)\alpha + \beta')} = \frac{\alpha\beta'(g(\lambda)^2 - 1)}{g(\lambda)\alpha^2 + g(\lambda)\beta'^2 + \alpha\beta'(g(\lambda)^2 + 1)} \\ &\leq \frac{\alpha\beta'(g(\lambda)^2 - 1)}{2g(\lambda)\alpha\beta' + \alpha\beta'(g(\lambda)^2 + 1)} = \frac{g(\lambda)^2 - 1}{2g(\lambda) + g(\lambda)^2 + 1} = \frac{g(\lambda) - 1}{g(\lambda) + 1} \end{aligned}$$

Finally, we observe that $\frac{g(\lambda)-1}{g(\lambda)+1} \leq f_\ell(\lambda)$, as desired:

$$E \leq \frac{\lambda^\ell - 1}{\lambda^\ell + 1} = 1 - 2(\lambda^\ell + 1)^{-1} \leq 1 - 2(\lambda^\ell + \lambda^\ell)^{-1} = 1 - \lambda^{-\ell} \leq 1 - \lambda^{-2\ell} = f_\ell(\lambda)$$

□

Theorem 6.4. The Generalized PA mechanism with parameter ℓ achieves a $(\frac{n-k}{n}(n-k)^{-1/\ell} + 1/n)$ -approximation to the optimal social welfare for any instance with n advertisers and k slots.

Proof. First, we consider the social welfare attained by the Generalized PA mechanism. Since $\beta_{k+1} = 0$ and $p_i^{(0)} = \vec{0}$,

$$\begin{aligned} SW(\text{ALG}) &= \sum_{i=1}^n \sum_{j=1}^k \alpha_i v_i \beta_j M_{ij} = \sum_{i=1}^n \sum_{j=1}^k \hat{v}_i \beta_j [p_i^{(j)} - p_i^{(j-1)}] \\ &= \sum_{i=1}^n \sum_{j=1}^k \hat{v}_i (\beta_j - \beta_{j+1}) p_i^{(j)} = \sum_{j=1}^k (\beta_j - \beta_{j+1}) \left(\sum_{i=1}^n \hat{v}_i p_i^{(j)} \right) \end{aligned}$$

Lemma 6.5 proves the approximation ratio of the k -unit PA mechanism. Observe that this ratio is decreasing in k . Therefore, for any j , $(\sum_{i=1}^n \hat{v}_i p_i^{(j)})$ is at least an $\eta = (\frac{n-k}{n}(n-k)^{-1/\ell} + 1/n)$ fraction of UNFAIR-OPT. Therefore, we have

$$SW(\text{ALG}) = \sum_{j=1}^k (\beta_j - \beta_{j+1}) \left(\sum_{i=1}^n \hat{v}_i a_i^{(j)} \right) \geq \eta \sum_{j=1}^k (\beta_j - \beta_{j+1}) (\hat{v}_1 + \dots + \hat{v}_j) = \eta \sum_{j=1}^k \hat{v}_j \beta_j = \eta \text{UNFAIR-OPT}$$

□