Fair Equilibria in Sponsored Search Auctions: The Advertisers' Perspective

Georgios Birmpas[1] , **Andrea Celli**[2] , **Riccardo Colini-Baldeschi**[3] , **Stefano Leonardi**[1]

[1]Department of Computer, Control and Management Engineering, Sapienza University, Rome, Italy
[2]Department of Computing Sciences, Bocconi University, Milan, Italy
[3]Core Data Science, Meta, London, UK
{birbas, leonardi}@diag.uniroma1.it, andrea.celli2@unibocconi.it, rickuz@fb.com

Abstract

In this work we introduce a new class of mechanisms composed of a traditional Generalized Second Price (GSP) auction and a fair division scheme, in order to achieve some desired level of fairness between groups of Bayesian strategic advertisers. We propose two mechanisms, β-Fair GSP and GSP-EFX, that compose GSP with, respectively, an *envy-free up to one item*, and an *envy-free up to any item* fair division scheme. The payments of GSP are adjusted in order to compensate advertisers that suffer a loss of efficiency due the fair division stage. We investigate the strategic learning implications of the deployment of sponsored search auction mechanisms that obey to such fairness criteria. We prove that, for both mechanisms, if bidders play so as to minimize their external regret they are guaranteed to reach an equilibrium with good social welfare. We also prove that the mechanisms are budget balanced, so that the payments charged by the traditional GSP mechanism are a good proxy of the total compensation offered to the advertisers. Finally, we evaluate the quality of the allocations through experiments on real-world data.

1 Introduction

Over the last decades, online advertising has been one of the main tools for small and medium business (SMB) to grow. Online advertising allows SMBs to reach potential customers without geographical or demographic barriers. Moreover, it offers better return-on-investments than other advertising mediums thanks to its highly personalized system. Given its crucial role in the growth of businesses (and, in its turn, society), it is natural to study how the mechanisms implemented in online advertising platforms can be improved to obey to different fairness criteria for advertisers. There are indeed various settings where one may care about balancing ads allocations, even just between a majority and a minority group of advertisers. For example, one setting is when large companies and small businesses are competing for the same set of users, and the platform may want to ensure that small companies get reasonable visibility despite smaller budgets. As another example, there may be businesses based in different geographical locations but offering similar products/services. These businesses may have unbalanced budgets due to their location, but one may want to guarantee to each business some visibility in a target set of users regardless of geographical attributes. More in general, one could think of incorporating fairness in sponsored search auctions whenever two different parties are advertising to increase users' awareness about some sensitive topic.

We study how fairness notions borrowed from the fair division literature can be used to model fairness with respect to advertisers. Notions from fair division already found applications in online advertising [Chawla and Jagadeesan, 2020; Ilvento *et al.*, 2020]. In the fair division literature, the dominant notion of fairness aims at providing guarantees for individual agents (see, e.g., [Brandt *et al.*, 2016; Moulin, 2003]). However, we argue that a *group approach* would be more practical, better aligned with societal expectations, and easier to implement. Following recent works in group fair division literature [Kyropoulou *et al.*, 2020; Manurangsi and Suksompong, 2017; Manurangsi and Suksompong, 2019; Segal-Halevi and Suksompong, 2019], we propose to use *envy-freeness* [Foley, 1967] to study *group fairness*. *Group envy-freeness* guarantees that no group of agents envies the allocation obtained by any other group. Unfortunately, envy-freeness cannot be guaranteed for indivisible items even in simple settings with two agents and one item. Thus, we focus on two natural relaxations: envy-freeness up to *one* good (EF1) [Budish, 2011; Lipton *et al.*, 2004], and envy-freeness up to *any* good (EFX) [Caragiannis *et al.*, 2019; Gourvès *et al.*, 2014]. [1]

In practice, any attempt to guarantee such properties in real advertising platforms will inevitably collide with real-world engineering constraints. Therefore, a credible solution should be a mechanism that can be easily integrated with a pre-existing auction framework, without requiring substantial changes to it. We focus on the *generalised second price* (GSP) auction framework [Edelman *et al.*, 2007], which is one of the most frequently adopted mechanisms for the allocation of advertising opportunities in large Internet advertising companies. In this setting, we show the existence of simple mechanisms that guarantee some notion of group EF1 (resp., group EFX) for advertisers. Such mechanisms can be

[1]See the full paper for additional discussion of related work.

implemented as a post-auction layer to be run after a standard GSP mechanism. In the spirit of the work by [Dwork and Ilvento, 2019], we study the properties of such composite mechanisms.

Original contributions We focus on a Bayesian setting with incomplete information (i.e., the valuations for advertising opportunities are stochastic, and each bidder does not observe the realized valuations of the other bidders). For each auction, bidders are divided in two groups (a majority group and a minority group) based on their characteristics and competitiveness. Given the different characteristics, the users interact in different ways with the ads from the two groups. This is modeled through group specifics click-through rates and quality factors.[2] First, in Section 3, we introduce two notions of group envy-freeness (*group* β-EF1 and *group* β-EFX), parameterized on a factor β which allows the platform to tune the strength of the fairness requirement. We show that group β-EF1 and group β-EFX allocations always exist for two groups with monotonic click-through rates. Moreover, they can be computed efficiently with two fair division schemes which are, respectively, "group versions" of the round robin procedure, and the envy-cycle elimination algorithm by Lipton *et al.* [2004]. We define fairness with respect to reported bidder's valuations. Indeed, we argue that an auction is perceived to be fair if the slots are divided between the majority/minority group in a fair way, given that prices are not publicly known to all advertisers in real-world sponsored search auctions. Considering the social welfare of the allocation is also the customary approach in the fair division literature, in which notions such as EF1 and EFX are defined on valuations (see, e.g., [Budish, 2011; Amanatidis *et al.*, 2019]). Then, in Section 4, we study the efficiency and budget-balance of the mechanisms resulting from the composition of GSP with the two fair division schemes. We show that the social welfare of the composite mechanisms is a good approximation of the optimum. Moreover, we prove that the welfare loss experienced by the advertisers due to the fairness constraints can be partly compensated through monetary incentives via redistribution of the GSP payments. Finally, Section 5 studies the behavior of the composite mechanisms when bidders behave as no-regret learning agents (i.e., they take decisions so as to minimize their external regret) and they are conservative (i.e., they do not overbid).[3] We show that the learning dynamic originating from the interaction with our composite mechanisms converges to a *Bayesian coarse correlated equilibrium* with good social welfare properties characterized through their *price of composition*. We complement our theoretical results by evaluating the quality of our fair sponsored search mechanisms on real-world data (Section 6).

2 Preliminaries

Throughout the paper, bold case letters denote column vectors. Given a vector $\boldsymbol{y}$, its i-th component is denoted by y_i. The set $\{1, \ldots, x\}$ is denoted by $[x]$, and $\Delta_{\mathcal{X}}$ is the $|\mathcal{X}|$-dimensional simplex over the discrete set $\mathcal{X}$.

Sponsored search framework There is a set I of n bidders and a set J of m slots. An *outcome* is an assignment of bidders to slots. Each bidder i has a private type v_i, representing their valuation on the item which is being sold. The vector of types is denoted as $\boldsymbol{v} = (v_1, \ldots, v_n)$. Each bidder i belongs to a group from a finite set of possible groups $\mathcal{G}$. The function $g : [n] \to \mathcal{G}$ maps bidders to their group, that is, we write $g(i)$ to denote the group to which bidder i belongs. We assume that bidders may belong to one of two groups $\mathcal{G} = \{h, \ell\}$ (e.g., a majority group and a minority group) and the two groups may have different sizes. We denote the set of bidders belonging to the two groups by I_h and by I_ℓ. We make the assumption that $|I_h| \geq m$ and $|I_\ell| \geq m$. This assumption is reasonable in the context of large Internet advertising markets. As it is customary in the literature, we use the model of separable click probabilities (see, e.g., [Edelman *et al.*, 2007; Varian, 2007]), in which each slot j is associated with a *click-through rate* $\alpha_{j,g(i)}$ for group $g(i)$. We assume that, for each group $w \in \mathcal{G}$, $\alpha_{1,w} \geq \alpha_{2,w} \geq \ldots \geq \alpha_{m,w}$, and without loss of generality we take $n = m$. Group-specific click-through rates can model, for example, cases in which different groups of advertisers mainly resort to different media types. Finally, each group $w \in \mathcal{G}$ is associated with a *quality factor* $\gamma_w \in [0, 1]$, which reflects the *clickability* of ads from bidders belonging to group w. For example, ads coming from advertisers in the majority group may have higher clickability than ads from advertisers in the minority groups because of differences in the available budgets to develop the campaigns. Quality factors are private knowledge of the advertising platform, and not known by the bidders. [4]

GSP auction A *mechanism* elicits a bid $b_i \in \mathbb{R}_{\geq 0}$ for each bidder i, which is interpreted as a type declaration, and computes an outcome as well as a price $p_i(\boldsymbol{b}, \boldsymbol{\gamma})$ for each bidder i. We denote by $\pi(\boldsymbol{b}, \boldsymbol{\gamma}, j)$ the bidder assigned to slot j when the mechanism observes the bid vector $\boldsymbol{b}$ and vector of quality factors $\boldsymbol{\gamma}$. We also denote by $\nu(\boldsymbol{b}, \boldsymbol{\gamma}, i)$ the slot assigned to bidder i when the mechanism observes the bid vector $\boldsymbol{b}$ and vector of quality factors $\boldsymbol{\gamma}$. When the vectors of bids and quality factors are clear from the context we simplify the notation by writing $\pi(j)$, $\nu(i)$, and p_i in place of $\pi(\boldsymbol{b}, \boldsymbol{\gamma}, j)$, $\nu(\boldsymbol{b}, \boldsymbol{\gamma}, i)$, and $p_i(\boldsymbol{b}, \boldsymbol{\gamma})$, respectively. The value perceived by bidder i when they are allocated $\nu(i)$ is $\alpha_{\nu(i),g(i)}\gamma_{g(i)}v_i$, and their utility is $u_i(\boldsymbol{b}, \boldsymbol{v}, \boldsymbol{\gamma}) := \alpha_{\nu(i),g(i)}\gamma_{g(i)}v_i - p_i$. We focus on a family of mechanisms derived from the *Generalized Second Price* (GSP) auction [Varian, 2007]. In a GSP auction the mechanism assigns the slots in order from 1 to m and sets $\pi(\boldsymbol{b}, \boldsymbol{\gamma}, j)$ to be the bidder with the highest *effective bid* $\gamma_{g(i)}\alpha_{j,g(i)}b_i$ not yet assigned (breaking ties arbitrarily).

[2] The online advertising problem with group specifics click-through rates has already been formalized in literature in the Ad-Types setting [Colini-Baldeschi *et al.*, 2020; Elzayn *et al.*, 2021].

[3] The assumption of having conservative bidders has often been used in the analysis of ad auction mechanisms (see, e.g., [Leme and Tardos, 2010; Christodoulou *et al.*, 2016; Feldman *et al.*, 2013]). The rationale for this is usually that, in practice, bidders are only partially informed and, thus, they may be more inclined to avoid risks (i.e., to submit bids that might yield negative payoffs).

[4] Our results hold also in the case of advertiser-dependent clickability, because that does not alter the relative ordering of items.

For any bid profile $\boldsymbol{b}$, quality factors $\boldsymbol{\gamma}$ and for each $j \in [m]$, $i = \pi(j)$, the price charged to bidder i is computed as

$$p_i^{\mathrm{G}}(\boldsymbol{b}, \boldsymbol{\gamma}) := \frac{\gamma_{g(\pi(j+1))}\alpha_{j,g(\pi(j+1))}}{\gamma_{g(i)}} b_{\pi(j+1)}, \qquad (1)$$

where we set $b_{n+1} = 0$.[5] The mechanism is *Individually Rational* (IR) if, for each bidder $i \in I$, $u_i(\boldsymbol{b}, \boldsymbol{v}, \boldsymbol{\gamma}) \geq 0$, for all $\boldsymbol{b}$, $\boldsymbol{v}$, and $\boldsymbol{\gamma}$. The mechanism is *Individually Rational at the Equilibrium* (IRE) if it is IR at the equilibrium bid vectors.

Online Bayesian framework The n bidders participate in a series of GSP auctions. At each iteration t, each bidder i observes a valuation v_i^t for the item being sold at time t. Let $\mathcal{V}_i$ be the finite set of types of bidder i. The vector of types $\boldsymbol{v}^t = (v_1^t, \ldots, v_n^t)$ is drawn, at each t, from a (possibly correlated) probability distribution $\mathcal{F}$ supported on a finite set of joint types $\mathcal{V}$, that is, $\mathcal{V} := \times_{i\in[n]}\mathcal{V}_i$. Moreover, at each t, a vector of quality factors $\boldsymbol{\gamma}^t \in [0,1]^{|\mathcal{G}|}$ is drawn from a (possibly correlated) distribution $\mathcal{G}$. Each bidder i has an arbitrary finite set of available bids $\mathcal{B}_i \subseteq \mathbb{R}_{\geq 0}$, with $\bar{B}_i := \max \mathcal{B}_i$ and $\bar{B}_i \geq \sup \mathcal{V}_i$. Moreover, we denote by $\mathcal{B} := \times_{i\in[n]}\mathcal{B}_i$ the set of all possible joint bid profiles. A bidding strategy σ_i for bidder i is a (possibly randomized) mapping from their types $\mathcal{V}_i$ to their available bids $\mathcal{B}_i$. We represent such strategies as a $|\mathcal{V}_i| \times |\mathcal{B}_i|$ right stochastic matrix in which each row specifies a well-defined probability distribution over bids: bidder i's strategy space is $\Sigma_i := \left\{\sigma_i \in \mathbb{R}_{\geq 0}^{|\mathcal{V}_i|\times|\mathcal{B}_i|} : \sigma_i \mathbf{1} = \mathbf{1}\right\}$. We observe that bidders cannot condition their bids on their quality factors $\boldsymbol{\gamma}$, since they are only known to the platform, and not to advertisers. Finally, we define the set of joint bidding strategies as

$$\Sigma := \left\{\sigma \in \Delta_{\mathcal{V}\times\mathcal{B}} \,:\, \sum_{\boldsymbol{b}\in\mathcal{B}} \sigma[\boldsymbol{v}, \boldsymbol{b}] = \mathcal{F}(\boldsymbol{v}),\ \forall \boldsymbol{v} \in \mathcal{V}\right\}.^{6}$$

At each iteration t, bidder i places bids according to a bidding strategy $\sigma_i^t \in \Sigma_i$. In particular, bidder i observes its own type v_i^t, and then submits a bid $b_i^t \sim \sigma_i^t[v_i^t]$. Then, bidder i experiences a reward which we define as a function $u_i^t : \mathcal{B}_i \to \mathbb{R}$. The utility function u_i^t observed at time t implicitly depends on the realized vector of bids $\boldsymbol{b}_{-i}^t$, and quality factors $\boldsymbol{\gamma}^t$, and it is such that, for each possible bid $b \in \mathcal{B}_i$, $u_i^t(b) := \alpha_{j,g(i)}^t \gamma_{g(i)}^t v_i^t - p_i((b, \boldsymbol{b}_{-i}^t), \boldsymbol{\gamma})$, with $j = \nu((b, \boldsymbol{b}_{-i}^t), \boldsymbol{\gamma}^t, i)$.

Regret and equilibria Given a sequence of decisions $(b_i^1, \ldots, b_i^T)$ up to time T, the *external regret* of bidder i in type v_i is how much they regret not having played the best fixed action in hindsight at each iteration in which they observed type v_i. Formally, the regret experienced by bidder i under a certain type $v_i \in \mathcal{V}_i$ is

$$R_{v_i}^T := \max_{\hat{b}\in\mathcal{B}_i}\left\{\sum_{t=1}^{T} \mathbb{1}[v_i = v_i^t]\left(u_i^t(\hat{b}) - u_i^t(b_i^t)\right)\right\}.$$

Then, the cumulative external regret of bidder i at time T is $R_i^T := \sum_{v_i\in\mathcal{V}_i} R_{v_i}^T$. Let $(\boldsymbol{b}^t)_{t=1}^T$ be the sequence of decisions made by the bidders up to time T. Then, the *empirical frequency of play* $\bar{\sigma}^T \in \Delta_{\mathcal{V}\times\mathcal{B}}$ obtained from the realized sequence of types $(\boldsymbol{v}^t)_{t=1}^T$, and from the sequence of play $(\boldsymbol{b}^t)_{t=1}^T$ is such that, for every $(\boldsymbol{v}, \boldsymbol{b}) \in \mathcal{V} \times \mathcal{B}$:

$$\bar{\sigma}^T[\boldsymbol{v}, \boldsymbol{b}] := \frac{1}{T}\left|\left\{1 \leq t \leq T : \boldsymbol{b}^t = \boldsymbol{b}, \boldsymbol{v}^t = \boldsymbol{v}\right\}\right|.$$

If each bidder i plays so as to obtain a regret R_i^T growing sublinearly in T, then, in the limit as $T \to \infty$, the empirical frequency of play $\bar{\sigma}^T$ is guaranteed to converge almost surely to a *Bayesian coarse correlated equilibrium* (BCCE) [Caragiannis *et al.*, 2015; Hartline *et al.*, 2015].We denote by $\Sigma^* \subseteq \Sigma$ the set of all Bayesian coarse correlated equilibria of the game.

3 Group Fairness in GSP Auctions

In this section we present the two group fair division schemes that will be added as a post-auction layer to GSP.

Preliminary definitions Consider an arbitrary stage t of the repeated auctions process (dependence on t will be omitted when clear from the context). For each group $w \in \mathcal{G}$, let $\mathrm{ALG}_w(\boldsymbol{b}) := \sum_{i\in I_w} \gamma_w \alpha_{\nu(i),w} b_i$ be the value obtained by group w via a generic mechanism with allocation rule ν, on bid vector $\boldsymbol{b}$. Since $\mathcal{G} = \{h, \ell\}$, the overall value is $\mathrm{ALG}(\boldsymbol{b}) = \mathrm{ALG}_h(\boldsymbol{b}) + \mathrm{ALG}_\ell(\boldsymbol{b})$. Moreover, given a set of slots $J' \subseteq [m]$ assigned to group $w \in \mathcal{G}$, we define

$$\mathrm{ALG}_w(J', \boldsymbol{b}) := \sum_{j\in[|J'|]} \gamma_w \alpha_{J'[j],w} b_{I_w[j]},$$

where $J'[j]$ denotes the slot with the j-th click-through rate among slots in J', and $I_w[j]$ is the bidder belonging to I_w with the j-th effective bid in decreasing order (e.g., $I_h(1)$ is the bidder of group h with the highest effective bid). Intuitively, $\mathrm{ALG}_w(J', \boldsymbol{b})$ is the maximum value attainable by group w when it is allocated J' and the bid vector is $\boldsymbol{b}$.

β-EF1 mechanism The first mechanism that we describe employs a fair division scheme that guarantees the following notion of group fairness.

Definition 1. *(Group β-EF1 fairness) Let $\beta := \xi_\ell/\xi_h$, with $\xi_h, \xi_\ell \in \mathbb{N}^+$, $\xi_h \geq \xi_\ell$, and $\xi_h + \xi_\ell \leq m$. We say that an allocation is* group β envy-free up to one good (β-EF1 fair*) for $\beta \leq 1$ and bid profile* **b** *if, for each pair of groups $h, \ell \in \mathcal{G}$, there exists one item $j_h \in J_\ell$ such that $\mathrm{ALG}_h(J_h, \boldsymbol{b}) \geq \beta\,\mathrm{ALG}_h(J_\ell \setminus \{j_h\}, \boldsymbol{b})$.*

We fix ξ_ℓ and ξ_h to be the *smallest* integers such that $\beta = \xi_\ell/\xi_h$. A group β-EF1 fair allocation can be obtained through a round robin procedure that assigns ξ_h slots to group h for each ξ_ℓ slots assigned to group ℓ. The proof of this result is very similar to the one for the classical EF1 fair division scheme by Markakis [2017]. As an example, if we assume group h to be the majority group (i.e., advertisers from group h are allocated the highest slots in the ranking), then the result of the application of the group β-EF1 round robin procedure is the shift of advertisers assigned to a position $j \in J_h$ to

[5] Alternatively, we could charge bidder $\pi(j)$ with the *threshold price*, which is the smallest effective bid p_i^{T} that guarantees them the same slot. Observe that $p_i \leq p_i^{\mathrm{T}}$ since the next highest bidder on j is not necessarily $\pi(j+1)$ if the bidders have different quality factors.

[6] Given a matrix M, we denote by $M[i]$ its i-th row vector, and by $M[i,j]$ the entry in position (i,j).

position at most $\lceil(1+\beta)j\rceil - 1$. For more details on how this fair division scheme is implemented see the extended version of the paper. We observe that a round robin procedure guarantees group EF1 even when the number of groups is $|\mathcal{G}| > 2$. The reason is that, for any pair $w, w' \in \mathcal{G}$, and positions $j \in [|J_w|], j' \in [|J_{w'}|]$, with $j \le j'$, it holds $\alpha_{J_w[j],w} \ge \alpha_{J_{w'}[j'],w}$. Moreover, it is possible to show that a round robin procedure also guarantees group β-EF1 with respect to valuations (proofs can be found in the extended version of the paper).

Theorem 1. *Given a bid profile $\boldsymbol{b}$, the allocation computed by the composite mechanism is β-EF1 fair with respect to the valuation profile $\boldsymbol{v}$. In particular, given the allocation of slots to the two groups J_h, J_ℓ, for each pair of groups $h, \ell \in \mathcal{G}$, there exists one item $j \in J_\ell$ such that* $\textsc{Alg}_h(J_h, \boldsymbol{v}) \ge \beta\, \textsc{Alg}_h(J_\ell \setminus \{j\}, \boldsymbol{v})$.

Intuitively, this is because the β-EF1 mechanism computes J_h, J_ℓ without employing the reported bid profile, which is used only to determine the per-group ranking.

β-EFX mechanism The second notion of fairness which we consider is group β-EFX fairness.

Definition 2. *(Group β-EFX fairness) An allocation is* group β-envy free up to any good *(*β-EFX fair*) for $\beta \le 1$ and bid profile $\boldsymbol{b}$ if, for each pair of groups $h, \ell \in \mathcal{G}$, and for each item $j_h \in J_\ell$, it holds* $\textsc{Alg}_h(J_h, \boldsymbol{b}) \ge \beta\, \textsc{Alg}_h(J_\ell \setminus \{j_h\}, \boldsymbol{b})$.

A group β-EFX fair division scheme can be obtained through a "group version" of the envy-cycle elimination algorithm by Lipton *et al.* [2004]. In particular, we propose the *Group Envy-Cycle-Elimination algorithm* (GECE). The GECE algorithm can be summarized as follows: denote by J_h and J_ℓ the set of slots assigned, respectively, to group h and ℓ. We say that group h envies group ℓ if $\textsc{Alg}_h(J_h, \boldsymbol{b}) < \beta \textsc{Alg}_h(J_\ell, \boldsymbol{b})$. Initially, all the slots are not assigned, that is, $J_h = J_\ell = \varnothing$. Then, the algorithm iterates through the slots in decreasing order of click-through rate. The first slot is assigned to group h. For each subsequent slot j, the algorithm checks if groups h and ℓ envy each other, and, if this is the case, the algorithm swaps their allocations. Otherwise, if group ℓ does not envy group h, then the next slot is assigned to group h, else, if ℓ envies h, the slot is assigned to group ℓ. In the following theorem, we prove that the GECE algorithm is guaranteed to obtain a β-EFX allocation.

Theorem 2. *The allocation computed by the group envy-cycle-elimination (GECE) algorithm is group β-EFX fair.*

It is possible to show that a variation of GECE yields a group EFX fair allocation (i.e., β-EFX with $\beta = 1$) even with more than 2 groups.

Corollary 1. *The allocation computed by the k-group envy-cycle-elimination (k-GECE) algorithm is group EFX.*

4 Efficiency and Budget Balance

In this section, we study the efficiency and budget balance of the two mechanisms obtained by combining GSP with the two fair division schemes described in Section 3. Let us denote one such composite mechanism by $\textsc{c}$. The post-auction layer of the composite mechanism $\textsc{c}$ is modifying the GSP allocation of slots to bidders. Ideally, no bidder should be penalized for this re-allocation. Therefore, we need to update the payments so that bidders' utility is not negatively affected by the composition of GSP with the fair division scheme. Interestingly, we can do so starting from the payments of GSP. In particular, denote by p_i^{G} and by p_i^{C}, the payments charged to advertiser i computed by GSP and by the composite mechanism, respectively. Moreover, let $\nu^{\text{G}}(i)$ and $\nu^{\text{C}}(i)$ be the slots assigned to advertiser i by GSP and by the composite mechanism, respectively. Then, we define the payments charged by the composite mechanism $\textsc{c}$ as:

$$p_i^{\text{C}} := \begin{cases} p_i^{\text{G}} & \text{if} \quad \nu^{\text{C}}(i) \le \nu^{\text{G}}(i) \\ p_i^{\text{G}} - 2b_i\, \gamma_{g(i)}\big(\alpha_{\nu^{\text{G}}(i),g(i)} - \alpha_{\nu^{\text{C}}(i),g(i)}\big) & \text{else.} \end{cases} \tag{2}$$

The pricing rule shows that the composite mechanism $\textsc{c}$ compensates the loss of social welfare of the advertisers that obtain a worse slot by reducing their payments. In order to ensure individual rationality, the advertisers that obtain a better slot are not asked to compensate with a higher payment. Observe that the pricing rule does not exclude positive transfers to the advertisers.

Now, let us first define an appropriate notion of budget balance for a composite mechanism which assigns a payment p_i^{C} to bidder $i \in I$, with respect to the GSP mechanism. Let $p^{\text{G}} := \sum_{i \in I} p_i^{\text{G}}$, and $p^{\text{C}} := \sum_{i \in I} p_i^{\text{C}}$.

Definition 3. *A composite mechanism is α-budget balanced, for $\alpha \ge 0$, if $p^{\text{G}} - p^{\text{C}} \le \alpha\, p^{\text{G}}$.*

An α-budget balanced mechanism is therefore able to cover via the GSP payments at least an α fraction of the compensations given to the bidders by the composite mechanisms.

Let $\textsc{Alg}^{\text{G}}(\boldsymbol{b})$ and $\textsc{Alg}^{\text{C}}(\boldsymbol{b})$ be, respectively, the value of the GSP mechanism, and of the fair composite mechanism on an arbitrary bid vector $\boldsymbol{b}$. Moreover, we use the following two assumptions in the analysis of the composite mechanisms.

Assumption 1. *The value of the minority group increases after the application of the composite mechanism, i.e.,* $\textsc{Alg}_\ell^{\text{C}}(\boldsymbol{b}) \ge \textsc{Alg}_\ell^{\text{G}}(\boldsymbol{b})$.

Assumption 2. *The first slot is assigned by GSP to $I_h(1)$, i.e., the first bidder of group h.*

Assumption 1 is natural since the basic goal of the proposed mechanisms would be to make the minority group better off with respect to the GSP case. Assumption 2 only requires that the advertiser with the highest bid belongs to the majority group. This is natural, for example, in settings where groups have different economic power. First, we provide efficiency and budget balance results for the β-Fair GSP mechanism.

Theorem 3. *The β-Fair GSP mechanism achieves a value that is at least a $1/(1+\beta)$ fraction of the value of GSP, i.e., for all bid vectors $\boldsymbol{b} \in \mathcal{B}$,* $\textsc{Alg}^{\text{C}}(\boldsymbol{b}) \ge \textsc{Alg}^{\text{G}}(\boldsymbol{b})/(1+\beta)$.

Theorem 4. *The β-Fair GSP mechanism is 2-budget balance.*

This result does not rule out that the mechanism may suffer a net loss. However, such loss is bounded by a small constant (see Section 6 for an empirical evaluation of such loss).

Second, we study efficiency and budget balance of the GSP-EFX mechanism. The allocation done by the mechanism is described in Section 3 and it is obtained by the composition of the GSP mechanism with the group EFX fair division scheme. The payments of GSP-EFX are computed as in Equation (2). Then, we have the following.

Theorem 5. *GSP-EFX achieves a value that is at least a fraction* $1/3$ *of the value of GSP, i.e., for all bid vectors* $\boldsymbol{b} \in \mathcal{B}$, $\text{ALG}^{\text{C}}(\boldsymbol{b}) \geq \frac{1}{3}\text{ALG}^{\text{G}}(\boldsymbol{b})$.

Finally, we prove that GSP-EFX is able to compensate at least $1/4$ of the total welfare loss generated by the application of the EFX fair division scheme. Formally,

Theorem 6. *The GSP-EFX mechanism is 4-budget balance.*

Remark 1. *Given a number of groups* $k > 2$ *and* $\beta = 1$, *Theorem 3 and Theorem 5 can be extended to show that the group EF1 (resp., group EFX) mechanism achieves a value that is at least a 1/2k fraction (resp., a fraction 1/3k) of the value of GSP.*

5 Price of Composition

Equipped with the results from Section 4 we can study the performance of the proposed mechanisms at equilibrium. We study the quality of the equilibria emerging as the results of the no-regret learning dynamics in which each bidder behaves as an external-regret minimizer. To do so, we propose the *price of composition* (PoC) as a natural measure to evaluate the social welfare guarantee of our mechanisms at equilibrium. For an arbitrary mechanism, the social welfare attained for bids $\boldsymbol{b}$, valuations $\boldsymbol{v} \in \mathcal{V}$, and quality factors $\boldsymbol{\gamma}$ is

$$SW(\boldsymbol{b}, \boldsymbol{v}, \boldsymbol{\gamma}) := \sum_{j\in[m]} \alpha_{j,g(\pi(j))}\gamma_{g(\pi(j))}v_{\pi(j)}. \tag{3}$$

Given an equilibrium strategy $\sigma \in \Sigma^*$ in an incomplete-information game, its social welfare is evaluated by comparing it to the expected ex-post social welfare of the GSP mechanism, which we denote by $\mathbb{E}_{\boldsymbol{v},\boldsymbol{\gamma}}[SW^{\text{G}}(\boldsymbol{v}, \boldsymbol{\gamma})]$. In particular, we can define the following worst-case ratio.

Definition 4. *The* price of composition (PoC) *of a composite mechanism* C *is defined as*

$$\text{PoC} := \inf_{\mathcal{F},\mathcal{G},\sigma\in\Sigma^*} \frac{\mathbb{E}_{\boldsymbol{v},\boldsymbol{\gamma},\boldsymbol{b}\sim\sigma}[SW^{\text{C}}(\boldsymbol{b}, \boldsymbol{v}, \boldsymbol{\gamma})]}{\mathbb{E}_{\boldsymbol{v},\boldsymbol{\gamma}}[\text{SW}^{\text{G}}(\boldsymbol{v}, \boldsymbol{\gamma})]}.$$

By the definition of the composite mechanisms, their social-welfare is at most equal to the social welfare of the GSP mechanism (see Section 4), i.e., $\text{PoC} \in [0, 1]$. Moreover, bounding the worst case PoC automatically yields a PoC/2 guarantee on the *price of total anarchy* [Blum and Mansour, 2007]. This is because $SW^{\text{G}}(\boldsymbol{v}, \boldsymbol{\gamma}) \geq SW^*(\boldsymbol{v}, \boldsymbol{\gamma})/2$, where SW^* is the optimal social welfare with valuations $\boldsymbol{v}$. [7]

In order to characterize the PoC of our mechanisms, we introduce a natural smoothness condition for composite mechanisms, which is a generalization of the notion of semi-smoothness by Lucier and Paes Leme [2011]. Let SW^{G} be the social welfare of the baseline mechanism (that is, in our setting, GSP), and let SW^{C} be the social welfare provided by the composite mechanism (in our case, β-Fair GSP or GSP-EFX) given an arbitrary vector or valuations. Then, our notion of smoothness states that social welfare gaps between the social welfare at the baseline mechanism with truthful bid vector $\boldsymbol{v}$, and the social welfare for an arbitrary joint strategy profile $\sigma \in \Sigma$ in the composite mechanism, can be captured by the marginal increases in the individual agents' utilities when unilaterally switching to a deviation strategy profile $\sigma' = (\sigma'_1, \ldots, \sigma'_n)$. Formally, a composite mechanism C is (λ, μ)*-semi-smooth with respect to a baseline mechanism* G if there exists a profile of individual bidding strategies $\sigma' = (\sigma'_1, \ldots, \sigma'_n) \in \times\Sigma_i$ such that, for any joint strategy $\sigma \in \Sigma$, $\boldsymbol{v} \in \mathcal{V}$, and $\boldsymbol{\gamma}$,

$$\mathbb{E}_{\substack{\boldsymbol{b}\sim\sigma[\boldsymbol{v}] \\ \boldsymbol{b}'\sim\sigma'[\boldsymbol{v}]}} \Big[\sum_{i\in[n]} u_i^{\text{C}}((b'_i, \boldsymbol{b}_{-i}), v_i, \boldsymbol{\gamma}) \Big] \geq \lambda SW^{\text{G}}(\boldsymbol{v}, \boldsymbol{\gamma}) - \mu \mathbb{E}_{\boldsymbol{b}\sim\sigma[\boldsymbol{v}]}[SW^{\text{C}}(\boldsymbol{b}, \boldsymbol{v}, \boldsymbol{\gamma})]. \tag{4}$$

In this setting, we make the additional assumption that bidders are conservative and, therefore, they do not overbid. This is in line with previous works studying the price of anarchy of auctions with conservative bidders (see, e.g., [Leme and Tardos, 2010; Christodoulou *et al.*, 2016; Feldman *et al.*, 2013]). [8] Under this assumption, the following holds.

Lemma 1. *The* β*-Fair GSP mechanism is* $(1/2, (1+\beta))$*-semi-smooth, and the GSP-EFX mechanism is* $(1/2, 3)$*-semi-smooth. Both mechanisms are individually rational at the equilibrium.*

By the fact that (λ, μ)-semi smoothness implies a $\frac{\lambda}{1+\mu}$ PoC, we characterize the PoC of the mechanisms as follows.

Theorem 7. *The price of composition with uncertainty of* β*-Fair GSP is* $\text{PoC} = (2(2+\beta))^{-1}$.

This result shows that, even if the bidders are self-interested (i.e., they take decisions so as to minimize their own individual regret), our composite mechanism can guarantee *group fairness*, as well as convergence to a *good* equilibrium point. Analogously, it is possible to prove the following for GSP-EFX:

Theorem 8. *The price of composition with uncertainty of GSP-EFX is* $\text{PoC} = 1/8$.

6 Experimental Evaluation

We experimentally evaluate the quality of the equilibria emerging as the results of no-regret learning dynamics in which agents interact through the β-Fair GSP mechanism.

Regret minimization for the Bayesian setting For each bidder i, we consider a discrete set of bids $\mathcal{B}_i$. We focus on the partial-information setting, in which, at each time instant t, each bidder observes only the reward $u_i^t(b_i^t)$ associated to

[7]The value of the GSP solution for the Ad-Types problem with group-specific CTRs is at least $1/2$ of the optimal solution provided by maximum weighted matching [Colini-Baldeschi *et al.*, 2020].

[8]Our bounds on PoC can be adapted with minor modifications to the case in which bidders are δ-conservative [Bhawalkar and Roughgarden, 2011], i.e., $b_i \leq \delta v_i, \forall i$. In particular, all price of composition factors are multiplied by an additional δ factor.

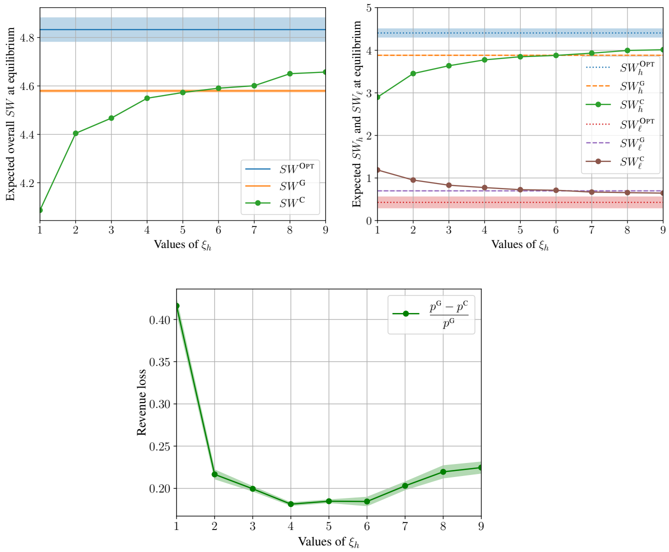

Figure 1: **(Top-Left):** Expected overall SW at equilibrium. **(Top-Right):** expected per-group SW at equilibrium. **(Bottom):** expected fraction of the GSP payments to be used for compensation in the composite mechanism. We set $\xi_\ell = 1$ and we let ξ_h vary.

its choice $b_i^t \in \mathcal{B}_i$. This is in line with what happens in real-world sponsored search auctions, where advertisers do not observe competing bids (i.e., they cannot compute a counterfactual utility $u_i^t(b)$ for each $b \in \mathcal{B}_i$), but can only observe the outcome associated to their decision b_i^t. We instantiate an external-regret minimizer $\mathcal{R}_{i,v}$ for each bidder i and $v \in \mathcal{V}_i$. In practice, we use the EXP3 algorithm by [Auer *et al.*, 2002] for each external regret minimizer. Then, we build a regret minimizer for the Bayesian bidder i as follows: at each t, bidder i observes their realized type $v^t \in \mathcal{V}_i$ and selects $b_i^t \in \mathcal{B}_i$ according to $\mathcal{R}_{i,v_i}$. Then, after utility $u_i^t(b_i^t)$ is observed, only $\mathcal{R}_{i,v_i}$ is updated. This simple procedure guarantees that $\limsup_{T\to\infty} R_i^T/T \leq 0$, which implies that the empirical frequency of play $\bar{\sigma}^T$ of the dynamic converges almost surely in the limit to a Bayesian coarse correlated equilibrium (see, e.g., [Hartline *et al.*, 2015, Lemma 10]).

Experimental setting We construct a real-world dataset through logs of a large Internet advertising company. We test our β-Fair GSP mechanism in an artificial environment where we have 20 advertisers, equally distributed among two groups, and competing for ad opportunities over a sequence of $T = 10^4$ auctions. For each i, we set $\mathcal{V}_i = \{x/100 : x \in [100] \cup \{0\}\}$. Moreover, in order to be able to create a steep unbalance between the two groups, we let, for each $\boldsymbol{v} \in \mathcal{V}$, $\mathcal{F}(\boldsymbol{v}) = \mathcal{F}_1(v_1) \cdot \ldots \cdot \mathcal{F}_n(v_n)$, with $\mathcal{F}_i \in \Delta_{\mathcal{V}_i}$ for each i. Then, for each $i \in I_h$ (i.e., advertiser i belongs to the majority group), we artificially set value distributions to be such that $\mathcal{F}_i(1) = 1$, and $\mathcal{F}_i(v) = 0$ for each $v \in \mathcal{V}_i \setminus \{1\}$. Each bidder $i \in I_\ell$ has a value distribution $\mathcal{F}_i$ built by normalizing the distribution of bids observed from real-world bidding data of a large Internet advertising company. We use bids as a proxy for true valuations of advertisers. Discount curves are computed by averaging and normalizing in $[0, 1]$ real-world discount factors. In particular, we estimate discount curves on ads optimizing for two distinct conversion types, one per group. This models different per-group preferences on the slots. Quality factors are set to be $\boldsymbol{\gamma} = (1, 1)$. Experiments are run on a 24-core machine with 57Gb of RAM.

Results Each advertiser takes decisions so as to minimize their external regret according to the procedure described above. Then, we analyse the equilibrium outcomes originating from advertisers behaving so as to minimize their regret, and interacting through GSP and β-Fair GSP, respectively. To do so, we compute the empirical frequency of play $\bar{\sigma}^T$ in the two settings. Let $\bar{\sigma}^{T,\mathrm{G}}$ and $\bar{\sigma}^{T,\mathrm{C}}$ be the empirical frequency of play obtained via the GSP mechanism, and the empirical frequency of play obtained via the β-Fair GSP mechanism, respectively. Figure 1–(Top-Left), and Figure 1–(Top-Right) report a comparison between the expected social welfare attained at $\bar{\sigma}^{T,\mathrm{G}}$ (i.e., SW^{G}), the expected social welfare attained at $\bar{\sigma}^{T,\mathrm{C}}$ (i.e., SW^{C}) for different values of ξ_h, while for simplicity we keep $\xi_\ell = 1$ (see Definition 1), and the expected social welfare obtained when advertisers are not strategic and submit bids truthfully to a GSP mechanism (i.e., SW^{OPT}). Each value is computed over 20 repetitions of the dynamics, and figures display the resulting mean and standard deviation. In particular, Figure 1–(Top-Left) reports the overall social welfare, and shows that, as expected, lower values of k increase ΔSW, that is, the gap between the social welfare provided by the GSP solution and the social welfare provided by the β-Fair GSP solution. The gap is negligible for most values of $\beta = 1/\xi_h$, and in the worst case, the empirical gap is approximately 10% of the social welfare provided by GSP. Figure 1–(Top-Right) reports the per-group social welfare at equilibrium. Finally, Figure 1–(Bottom) describes the fraction of the revenue which is lost due to compensation for advertisers as defined in Equation (2). We observe that, in practice, the fraction of the GSP prices used for compensation significantly lower than the worst case bound of Theorem 4. In particular, the composite mechanism looses nearly 40% of GSP revenue when $\xi_h = 1$ (i.e., when the group fairness constraint is as tight as possible). However, for higher values of ξ_h, the fraction of p^{G} which has to be used for compensations stabilizes around 20%. This suggests that there may be a trade-off between guarantees for advertisers and revenue losses incurred by the platform where this type of mechanisms could be viable in practice.

7 Discussion and Future Works

We proposed a class of composite mechanisms constructed through the combination of GSP with different fair division schemes. We show that these composite mechanisms yield group fairness guarantees for advertisers, and we characterized the costs which the platform has to incur. In the future, it would be interesting to study fairness guarantees in the space of valuations, and in the space of utilities. We already provided some preliminary results going in this direction (Theorem 1). Extending this to EFX presents non-trivial additional challenges, such as showing that there exists a monotone version of GECE. We leave this as a future research direction.

Acknowledgements

The Sapienza team was supported by the ERC Advanced Grant 788893 AMDROMA "Algorithmic and Mechanism Design Research in Online Markets", the MIUR PRIN project ALGADIMAR "Algorithms, Games, and Digital Markets", and the Meta Research grant on "Fairness and Mechanism Design".

References

[Amanatidis *et al.*, 2019] Georgios Amanatidis, Apostolos Ntokos, and Evangelos Markakis. Multiple birds with one stone: Beating 1/2 for EFX and GMMS via envy cycle elimination. *CoRR*, abs/1909.07650, 2019.

[Auer *et al.*, 2002] Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multiarmed bandit problem. *SIAM J COMPUT*, 32(1):48–77, 2002.

[Bhawalkar and Roughgarden, 2011] Kshipra Bhawalkar and Tim Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *SODA '11*, page 700–709, 2011.

[Blum and Mansour, 2007] Avrim Blum and Yishay Mansour. *Learning, regret minimization, and equilibria*, page 79–102. Cambridge University Press, 2007.

[Brandt *et al.*, 2016] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. *Handbook of computational social choice*. Cambridge University Press, 2016.

[Budish, 2011] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *J POLIT ECON*, 119(6):1061–1103, 2011.

[Caragiannis *et al.*, 2015] Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, Maria Kyropoulou, Brendan Lucier, Renato Paes Leme, and Éva Tardos. Bounding the inefficiency of outcomes in generalized second price auctions. *J ECON THEORY*, 156(C):343–388, 2015.

[Caragiannis *et al.*, 2019] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum Nash welfare. *ACM Trans. Economics and Comput.*, 7(3):12:1–12:32, 2019.

[Chawla and Jagadeesan, 2020] Shuchi Chawla and Meena Jagadeesan. Fairness in ad auctions through inverse proportionality. *arXiv preprint arXiv:2003.13966*, 2020.

[Christodoulou *et al.*, 2016] George Christodoulou, Annamária Kovács, and Michael Schapira. Bayesian combinatorial auctions. *J. ACM*, 63(2), April 2016.

[Colini-Baldeschi *et al.*, 2020] Riccardo Colini-Baldeschi, Julián Mestre, Okke Schrijvers, and Christopher A. Wilkens. The ad types problem. In *WINE '20*, volume 12495, pages 45–58, 2020.

[Dwork and Ilvento, 2019] Cynthia Dwork and Christina Ilvento. Fairness under composition. In *ITCS '19*, pages 33:1–33:20, 2019.

[Edelman *et al.*, 2007] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *AM ECON REV*, 97(1):242–259, 2007.

[Elzayn *et al.*, 2021] Hadi Elzayn, Riccardo Colini-Baldeschi, Brian Lan, and Okke Schrijvers. Equilibria in auctions with ad types. *arXiv preprint arXiv:2103.06177*, 2021.

[Feldman *et al.*, 2013] Michal Feldman, Hu Fu, Nick Gravin, and Brendan Lucier. Simultaneous auctions are (almost) efficient. In *STOC '13*, 2013.

[Foley, 1967] D. K. Foley. Resource allocation and the public sector. In *Yale Economics Essays*, page 7:45–98, 1967.

[Gourvès *et al.*, 2014] Laurent Gourvès, Jérôme Monnot, and Lydia Tlilane. Near fairness in matroids. In *ECAI '14*, page 393–398, 2014.

[Hartline *et al.*, 2015] Jason D Hartline, Vasilis Syrgkanis, and Éva Tardos. No-regret learning in bayesian games. *NeurIPS '15*, pages 3061–3069, 2015.

[Ilvento *et al.*, 2020] Christina Ilvento, Meena Jagadeesan, and Shuchi Chawla. Multi-category fairness in sponsored search auctions. In *ACM FAccT '20*, pages 348–358, 2020.

[Kyropoulou *et al.*, 2020] Maria Kyropoulou, Warut Suksompong, and Alexandros A Voudouris. Almost envy-freeness in group resource allocation. *Theoretical Computer Science*, 841:110–123, 2020.

[Leme and Tardos, 2010] Renato Paes Leme and Eva Tardos. Pure and bayes-nash price of anarchy for generalized second price auction. In *FOCS '10*, pages 735–744, 2010.

[Lipton *et al.*, 2004] Richard J Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *ACM EC '04*, pages 125–131, 2004.

[Lucier and Paes Leme, 2011] Brendan Lucier and Renato Paes Leme. GSP auctions with correlated types. In *ACM EC '11*, pages 71–80, 2011.

[Manurangsi and Suksompong, 2017] Pasin Manurangsi and Warut Suksompong. Asymptotic existence of fair divisions for groups. *MATH SOC SCI*, 89:100–108, 2017.

[Manurangsi and Suksompong, 2019] Pasin Manurangsi and Warut Suksompong. Computing a small agreeable set of indivisible items. *ARTIF INTELL*, 268:96–114, 2019.

[Markakis, 2017] E. Markakis. Approximation algorithms and hardness results for fair division with indivisible goods. *Trends in Computational Social Choice*, 2017.

[Moulin, 2003] Hervé Moulin. *Fair division and collective welfare*. MIT press, 2003.

[Segal-Halevi and Suksompong, 2019] Erel Segal-Halevi and Warut Suksompong. Democratic fair allocation of indivisible goods. *ARTIF INTELL*, 277:103167, 2019.

[Varian, 2007] Hal R Varian. Position auctions. *INT J IND ORGAN*, 25(6):1163–1178, 2007.