

Equilibria in Auctions with Ad Types

Hadi Elzayn
RegLab, Stanford University*
U.S.A.
hselzayn@stanford.edu

Brian Lan
Meta Platforms
U.S.A.
brianbclan@fb.com

Riccardo Colini-Baldeschi
Core Data Science, Meta Platforms
U.K.
rickuz@fb.com

Okke Schrijvers
Core Data Science, Meta Platforms
U.S.A.
okke@fb.com

ABSTRACT

This paper studies equilibrium quality of semi-separable position auctions (known as the *Ad Types* setting [9]) with *greedy* or *optimal* allocation combined with *generalized second-price* (GSP) or *Vickrey-Clarke-Groves* (VCG) pricing. We make three contributions: first, we give upper and lower bounds on the Price of Anarchy (PoA) for auctions which use greedy allocation with GSP pricing, greedy allocation with VCG pricing, and optimal allocation with GSP pricing. Second, we give Bayes-Nash equilibrium characterizations for two-player, two-slot instances (for all auction formats) and show that there exists both a revenue hierarchy and revenue equivalence across some formats. Finally, we use no-regret learning algorithms and bidding data from a large online advertising platform to evaluate the performance of the mechanisms under semi-realistic conditions. We find that the VCG mechanism tends to obtain revenue and welfare comparable to or better than that of the other mechanisms. We also find that in practice, each of the mechanisms obtains significantly better welfare than our worst-case bounds might suggest.

CCS CONCEPTS

• **Applied computing** → **Online auctions**; • **Theory of computation** → **Convergence and learning in games**; **Quality of equilibria**; **Computational pricing and auctions**.

KEYWORDS

Auctions, Game Theory, Price of Anarchy, Ad Auctions.

ACM Reference Format:

Hadi Elzayn, Riccardo Colini-Baldeschi, Brian Lan, and Okke Schrijvers. 2022. Equilibria in Auctions with Ad Types. In *Proceedings of the ACM Web Conference 2022 (WWW '22)*, April 25–29, 2022, Virtual Event, Lyon, France. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3485447.3512052>

*This work was primarily conducted while this author was an intern at Facebook and a doctoral candidate at the University of Pennsylvania.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WebConf '22, April 25–29, 2022, Lyon, France

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9096-5/22/04...\$15.00
<https://doi.org/10.1145/3485447.3512052>

1 INTRODUCTION

This paper characterizes equilibrium welfare and revenue properties of various auction formats in the *Ad Types setting*. The *Ad Types setting* [9] is a generalization of the standard position auction [10, 28], which has been a workhorse in online advertising for years. In the standard position auction setting, there are multiple positions where the auctioneer can place ads. Advertisers care about receiving clicks on their ads, and the classical model poses a *separable* click-through-rate (CTR) model, where ad slots have an associated discount $1 \geq \delta^1 \geq \delta^2 \geq \dots \geq 0$ that represents the advertiser-agnostic CTR of the slot. Importantly, the classical model assumes that while CTRs differ *across* slots, all advertisers obtain the same CTR for a *given* slot.

Because separability is required to obtain results such as the ex-post payoff equivalence to the Vickrey-Clarke-Groves (VCG) auction [10], recent work has empirically investigated its validity. On the whole, the literature (e.g. [12, 16, 23]) suggests that the separability assumption does not hold in practice. The *Ad Types setting* [9] relaxes this assumption; instead, it assumes that each ad has a publicly known type¹—such as ‘video ad’, ‘link-click ad’ or ‘impression ad’—and each ad type τ has its own associated position discount curve $1 \geq \delta_\tau^1 \geq \delta_\tau^2 \geq \dots \geq 0$. All ads from the same type share the same discount curve; as such, the model generalizes the separable position auction to be *semi-separable* while maintaining more structure than a general maximum-weight bipartite matching problem.

In this paper, we investigate what happens in the *Ad Types setting* when we perform the allocation using either the *greedy* or *optimal* allocation algorithm, and run pricing using either *Generalized Second-Price (GSP)* or *VCG* semantics. In three of the four possible combinations the resulting auction is not incentive compatible, so we investigate the revenue and welfare in equilibrium. We first study the general setting and obtain upper and lower bounds on welfare. We then turn to a special case and calculate simple Bayes-Nash equilibrium strategies, welfare, and revenue. Finally, we empirically study equilibria arrived at by repeated play using realistic data and no-regret learning algorithms.

Novelties of the *Ad Types setting*. The *Ad Types setting* is significantly more complex than the standard position auction, and key results from the position auction do not translate. For instance, in the position auction setting, the greedy and optimal

¹Type in the economics literature often refers to private information. That is not the case here: ad type refers to the conversion event that the advertiser cares about.

allocations (with respect to reported bids) are equivalent, and this allows for much quicker computation of the optimal allocation than is generally possible. The loss of this equivalence in the Ad Types setting introduces a significant computational complexity to attempting to instantiate the VCG mechanism. It also implies that position auction efficiency results, such as the existence of an equilibrium with optimal welfare [10], no longer hold.

Colini-Baldeschi et al. [9] show that in the Ad Types setting, one can compute the optimal allocation (with respect to reported bids) and associated VCG prices using an adapted version of the Kuhn-Munkres algorithm in $O(n^2(k + \log n))$ (where n is the number of slots, and k the number of ad types). However, there are two practical considerations that need to be taken into account: First, despite the auction-theoretical benefits of VCG, in practice online advertising platforms often use a GSP payment rule [2], so it is desirable to understand the impact of using GSP pricing instead of VCG pricing. Second, in content feeds there is often a large number of ads that are allocated, making even the $O(n^2(k + \log n))$ running time prohibitive, necessitating simpler non-optimal mechanisms whose strategic properties have been so far unstudied.

1.1 Contributions

This paper makes three main contributions:

- **Price of Anarchy Bounds.** In Section 3, we provide Price of Anarchy upper and lower bounds in the Ad Types setting for all combinations of greedy or optimal allocation paired with GSP and VCG pricing. In particular, greedy allocation has an upper bound for the Price of Anarchy of 4, regardless of the choice of pricing; for optimal allocation and GSP pricing, we give an upper bound that depends on the bidder types and number of bidders, but not valuations. We give lower bounds for the Price of Anarchy of 2 for greedy allocation with GSP pricing, $3/2$ for greedy allocation with VCG pricing, and $4/3$ for optimal allocation with GSP pricing.
- **Small Equilibrium Characterization.** In Section 4, we characterize the sole linear Bayes-Nash equilibrium for each mechanism in the simple case of two bidders, two slots, and valuations distributed uniformly over the unit interval.² In linear equilibrium, the greedy allocation with GSP pricing produces an equivalent amount of revenue to the optimal allocation with VCG pricing, and this revenue is larger than the revenue produced by either of the other possible mechanism (which are also equivalent to each other).
- **Evaluation on Realistic Data.** Finally, we conduct experiments to complement our theoretical results, in particular using no-regret learning algorithms to model advertisers. We first study a two-bidder two-slot simulated setting matching the special cases we analyzed, and show that the bidding strategies of no-regret learners converge to the equilibria we identified. Then in Section 5, we leverage data from a large online advertising platform to study the performance of each mechanism in larger settings with more realistic valuation

²While this may appear a very special case, explicit equilibrium characterization in auctions is notoriously complex. Most famously, in Vickrey’s original paper [29] he posed an open problem to characterize the equilibrium of a two-player first-price auction with uniform valuations in $[a_1, b_1]$ and $[a_2, b_2]$. The problem remained unsolved until nearly 50 years later [17]!

distributions. We draw bidder valuations from real (normalized and anonymized) bids to represent bidder valuations. Because the average empirical sequence of play of no-regret learners is known to converge to *coarse correlated equilibria*, the outcomes obtained represent a plausible estimate of equilibrium welfare and revenue. We find that the VCG mechanism tends to obtain revenue and welfare comparable to or better than that of the other mechanisms. We also find that while there is significant variation in revenue, each of the mechanisms tends to perform significantly better than our worst-case bounds would suggest.

1.2 Related Literature

Position Auctions. Position auctions have long been the workhorse in online advertising. The seminal works of Edelman et al. [10] and Varian [28] first proposed the separable model of the position auction—and described the GSP auction in this model—and showed that for GSP there exists an ex-post Nash equilibrium that is equivalent to the VCG outcome. Gomes and Sweeney [13] showed that GSP does not always admit a Bayes-Nash equilibrium. There is also a history of exploring alternative pricing rules for position auctions; for example Chawla and Hartline [6] study generalized first-price (GFP) semantics for position auctions and show that for independent and identically distributed (IID) valuations the equilibrium is unique and symmetric.

Price of Anarchy and Smoothness. Since explicit equilibrium computation in auction is challenging, researchers have focused on Price of Anarchy bounds, i.e. using the equilibrium conditions to give bounds on the welfare in *any* equilibrium. Paes Leme and Tardos [20] were the first to give Price of Anarchy bounds for GSP. A common approach to proving Price of Anarchy bounds is to use the smoothness framework proposed by Roughgarden [24, 26], though GSP is not smooth in this sense. Lucier and Paes Leme [21] and Caragiannis et al. [3] instead show that one can use a *semi-smoothness* condition and they give almost tight Price of Anarchy bounds for GSP. Smoothness has also been applied to other payment rules, such as GFP by Syrgkanis and Tardos [27].

Complex Ad Auctions. Advertising auctions have become significantly more complex since the initial position auction model was proposed; as such, recent work empirically tests its validity in modern practice. [23] finds that CTRs differ by advertiser size or content in the context of keyword search. [12] uses experimental variation to causally identify position effects, finding that CTR declines with position are steeper for some advertisers (especially off-brand) than others and may differ by as much as 100% across advertisers. Similarly, [16] finds that the position discount is significantly steeper for less well-known brands than more well-known ones. Taken as a whole, the literature suggests that the separability assumption does not hold in practice.

As a result, there is a body of work that explores relaxing the separability assumption in position auctions. Our work is based on the Ad Types setting formalized by Colini-Baldeschi et al. [9]. When each ad is its own type, this model is identical to the one with arbitrary action rates that are still independent between ads, which has been studied before by Abrams et al. [1], Carvallo and Wilkens [5] and Wilkens et al. [4]. To our knowledge, no equilibrium

characterizations or Price of Anarchy bounds are known in these settings. The closest is a paper by Colini-Baldeschi et al. [8] that studies the relationship between envy, regret and social welfare loss in the Ad Types setting for an alternative version of GSP called “extended GSP” using the same semi-smoothness framework as proposed by Caragiannis et al. [3].

2 MODEL AND PRELIMINARIES

Advertisers There are n advertisers (each associated with a single ad) competing for m (ordered) slots. Each ad has a publicly known type τ_i , such as ‘video ad’, ‘link-click ad’ or ‘impression ad’. Ad i of type τ_i has value-per-conversion v_i . Ads of different types have different conversion events, e.g. for a link-click ad the conversion event is a link click and for a video ad the conversion event is the user watching a video ad.

Slots Slots are indexed by integers which increase moving down the feed. (So “lower” slots have higher indices.) Ads in lower slots see fewer conversions, and we consider a semi-separable model³ to capture this effect: $\Pr[\text{conversion on ad } i \text{ (of type } \tau_i) \text{ in slot } s]$ is given by $\delta_{\tau_i}^s \cdot \beta_i$ where $\delta_{\tau_i}^s$ is the slot effect for a particular ad type τ_i (e.g., the probability that a user will watch a video ad if it is shown in the s th slot) and β_i is the advertiser effect. We assume without loss of generality that the advertiser effect is included in the advertiser’s value; that is, if the value-per-conversion of the advertiser is v'_i , then $v_i = \beta_i \cdot v'_i$. Since advertisers effectively *discount* their value for the slot by $\delta_{\tau_i}^s$, we call $(\delta_{\tau_i}^s)$ for all slots s the *discount curve*.

Bidding and Payoffs Advertisers submit a single bid b_i for a conversion, which may or may not be their true valuation v_i . They are charged price p_i (calculated by the auction) if a conversion happens, so in expectation they are charged $\delta_{\tau(i)}^s p_i$. Thus, the expected payoff of an advertiser for a given slot at a given price is $u_i(s, p_i) = \delta_{\tau(i)}^s (v_i - p_i)$.

Discount Curves We assume that discount curves *monotonically decrease with the slot index*: that is, $1 \geq \delta_{\tau_i}^1 \geq \delta_{\tau_i}^2 \geq \dots \geq 0$. We will say that $s \geq_i s'$ (read as i prefers s to s') if $\delta_{\tau_i}^s \geq \delta_{\tau_i}^{s'}$. Since the conversion probability decreases moving down the feed for all types, advertisers agree on their preference ordering over any pair of slots, so we can drop the subscript and simply use \geq . Notice that since slots lower down the feed are indexed by higher numbers, $s \geq s' \iff s \leq s'$; we will often speak in terms of preference in order to avoid confusion. In some restricted settings, we consider *geometric* discount curves that can be written as $\delta_{\tau}^s = c \cdot \delta^s$ for some fixed c, δ , where s is an exponent on the right hand side.

Auction Algorithms. Any auction must answer two questions: who gets what (allocation), and how much do they pay (pricing). We use $\mathcal{A} : \mathbf{b} \rightarrow \mathbf{s}$ to designate allocation algorithms, and $\mathcal{P} : \mathcal{A}, \mathbf{b} \rightarrow \mathbf{p}$ to designate pricing algorithms. Here, \mathbf{b} is a vector of bids and \mathbf{s} is a vector of slot assignments. In other words, an allocation algorithm \mathcal{A} maps bid vectors to slot vectors. A pricing algorithm \mathcal{P} , however, takes both a vector of bids and an allocation *algorithm* \mathcal{A} . Thus the pricing algorithm is a meta-algorithm, rather than a particular algorithm. We refer to a pair $(\mathcal{A}, \mathcal{P}_{\mathcal{A}})$ as an *auction mechanism*. In

this paper we consider all combinations of two allocation algorithms and two pricing meta-algorithms:

- **Greedy (Allocation)** The *greedy* allocation begins with the highest slot, and among non-allocated bidders allocates the bidder whose *discounted* bid is highest (i.e., $\operatorname{argmax}_{i \in \mathcal{U}_s} \delta_{\tau(i)}^s b_i$, where \mathcal{U}_s is the set of unallocated bidders as of the time slot s is reached). For the Ad Types setting, the greedy algorithm generally does not yield the optimal allocation (see e.g. Example 1.1 in [9]).
- **Optimal (Allocation)** The *optimal* allocation computes the max-weight bipartite matching between ads and slot (where edge weights are discounted bids $\delta_{\tau(i)}^s b_i$), e.g. using the Kuhn-Munkres algorithm [19, 22].
- **GSP (Pricing)** The *Generalized Second Price* pricing rule executes the principle that a bidder pays the minimum bid under which they retain the slot they were assigned to, i.e. for allocation algorithm \mathcal{A} and bids \mathbf{b} : $[\mathcal{P}_{\mathcal{A}}(\mathbf{b})]_i := \operatorname{argmin}_{b: \mathcal{A}(b, \mathbf{b}_{-i}) = \mathcal{A}(\mathbf{b})} b$. Computing this bid is straightforward for the greedy allocation algorithm, while for the optimal algorithm we use the method of Carvallo et al [4].
- **VCG (Pricing)** The *Vickrey-Clarke-Groves* pricing rule [7, 14, 29] executes the principle that a bidder should pay their *externality*, i.e. for an allocation algorithm \mathcal{A} and bids \mathbf{b} : $[\mathcal{P}(\mathbf{b})]_i = \sum_{j \neq i} \delta_{\tau(j)}^{\mathcal{A}(\mathbf{b}_{-i})} b_j - \sum_{j \neq i} \delta_{\tau(j)}^{\mathcal{A}(\mathbf{b})} b_j$. When \mathcal{A} is the optimal allocation algorithm this yields the standard VCG algorithm. When \mathcal{A} is the greedy allocation algorithm, the resulting mechanism is not incentive compatible.

Given an auction $(\mathcal{A}, \mathcal{P}_{\mathcal{A}})$, bids \mathbf{b} , and valuations \mathbf{v} , the *social welfare* is $W(\mathcal{A}, \mathcal{P}_{\mathcal{A}}, \mathbf{b}, \mathbf{v}) = \sum_i \delta_{\tau(i)}^{\mathcal{A}(\mathbf{b})} \cdot v_i$ and the *revenue* is $\operatorname{Rev}(\mathcal{A}, \mathcal{P}_{\mathcal{A}}, \mathbf{b}, \mathbf{v}) = \sum_i \delta_{\tau(i)}^{\mathcal{A}(\mathbf{b})} \cdot \mathcal{P}_{\mathcal{A}}(\mathbf{b})_i$. At the risk of restating the obvious, notice that the auctioneer can only observe reported bids, not true valuations; hence, to the extent that each mechanism computes an “optimal” allocation, it is optimal with respect to the *bids*, not values. For non-incentive compatible mechanisms, these will not in general coincide, and we must take care in the analysis not to conflate the two; we will emphasize “apparent” with the “hat” symbol, e.g. we denote the apparent social welfare with respect to the bids as \widehat{W} .

Additional Notation To indicate vectors, we will use bold font: e.g. we denote the vector of bids as \mathbf{b} . We will use subscripts to denote a particular component, e.g. \mathbf{b}_i is the i th component of \mathbf{b} . At the risk of overloading notation, we will also use i as a subscript to track scalar functions for particular player. So, for instance, we can write b_i for player i ’s bid, or \mathbf{b}_i , depending on whether we are arguing about the auctioneer’s or player’s perspective. (The meaning of the subscript should be clear from context.) Also, we use the standard $_{-i}$ subscript to indicate “all but the i th component” of a vector. We will also use an analogous $^{-i}$ superscript for scalar functions, e.g. W^{-i} for the scalar welfare of all players but i .⁴

Since we consider multiple allocation and pricing formats, we write $\pi_{\mathcal{A}}(s, \mathbf{b})$ to indicate the *player* in *slot* s when \mathbf{b} is the bid profile and \mathcal{A} is the allocation algorithm. We will suppress the \mathcal{A}

³The model is semi-separable since ads share the same discount curve if and only if they are of the same type.

⁴To streamline notation, we will omit extra parentheses when writing the modified component and the “all-but-ith” component together where a vector would be required. For instance, $\mathcal{A}(b'_i, \mathbf{b}_{-i})$ instead of $\mathcal{A}((b'_i, \mathbf{b}_{-i}))$.

when it is clear from context. We use $\sigma_{\mathcal{A}}(i, \mathbf{b})$ to indicate the *slot* that player i receives when the bid profile is \mathbf{b} and the allocation algorithm is \mathcal{A} . We will sometimes overload notation to write $\tau(i)$ as a function returning player i 's ad type τ_i ; this will be useful when referring not to a specific player but rather to an arbitrary occupant of a given slot. We can also compose some or all of these together. For example, $\tau(\pi_{\mathcal{A}}(\sigma_{\mathcal{A}}(i, \mathbf{b}), \mathbf{b}'))$ is the type of the player assigned to the slot that i receives under the allocation algorithm \mathcal{A} and bid vector \mathbf{b} given that the bid vector is *instead* changed to \mathbf{b}' . Finally, we will denote the *optimal* allocation vector given a bid profile \mathbf{v} with the Greek letter ν . That is:

$$\nu := \operatorname{argmax}_{\sigma \in S_n} \sum_{i=1}^n \delta_{\tau(i)}^{\sigma_i} v_i,$$

where S_n is the set of all permutations of bidders. Maintaining our font conventions, we use $\nu(i)$ for the mapping of i to the slot index he is assigned under ν .

2.1 Solution Concepts and Learning

Each mechanism induces a game between agents that act strategically, so the *equilibrium concept* is an important modeling choice. In this paper we present equilibrium results for both full-information and Bayes-Nash equilibria:

Definition 2.1 (Nash Equilibrium). A bid profile \mathbf{b} is pure strategy Nash equilibrium if for each player i : $u_i(\mathbf{b}) \geq u_i(\mathbf{b}', \mathbf{b}_{-i})$ for all pure strategies \mathbf{b}' .

Definition 2.2 (Bayes-Nash Equilibrium). For a known value distribution \mathcal{V} , the vector of mappings $\mathbf{b}(\nu)$ is a Bayes-Nash equilibrium if for every player i :

$$\mathbb{E}_{\nu \sim \mathcal{V}}[u_i(\mathbf{b}(\nu))] \geq \mathbb{E}_{\nu \sim \mathcal{V}}[u_i(\mathbf{b}'_i(\nu_i), \mathbf{b}_{-i}(\nu_{-i}))]$$

for any other mapping $\mathbf{b}'_i(\nu_i)$.

For each of these equilibrium notions, an ϵ -approximate version is obtained by allowing the definitional inequality to be violated by no more than ϵ . A bid profile where no bidder can improve their payoff by more than ϵ is an ϵ -approximate Nash equilibrium. A Bayes-Nash equilibrium is *linear* if $b_i(v_i) = B_i v_i$ for some $B_i \in \mathbb{R}$. We say a bidder is *conservative* if he does not bid above his value; an equilibrium is conservative if it does not prescribe bidding above one's value. Some results in Section 3 assume conservative bidders.

For each equilibrium concept, there may be multiple equilibria with different welfare. The Price of Anarchy (PoA) captures the worst-case⁵ welfare compared to the optimal welfare knowing the valuations. Here, we write its definition adapted to our setting:

Definition 2.3 (Price of Anarchy). The Price of Anarchy is

$$\text{PoA(Nash)} := \max_{\mathbf{b} \in E} \frac{\sum_i \delta_{\tau(i)}^{\nu(i)} \cdot v_i}{\mathbb{E}[\sum_i \delta_{\tau(i)}^{\mathcal{A}(\mathbf{b})_i} \cdot v_i]}$$

where E is the set of Nash equilibria for $(\mathcal{A}, \mathcal{P}_{\mathcal{A}})$ and the randomness is over the strategy distributions. A similar definition can be made for a Bayesian PoA with randomness over the valuations.

⁵We also use the term ‘‘empirical PoA’’ to describe the ratio of average realized welfare to optimal welfare when speaking about specific or empirical cases. Strictly speaking the PoA is only the worst-case value, but the meaning should be clear.

In general, it may be difficult or impossible to analytically characterize equilibria in more complicated settings. Thus, in Section 5, we turn to *learning* equilibria using no-regret learning algorithms on data drawn from realistic valuation distributions. This approach, while powerful, is not guaranteed to recover either Nash or Bayes-Nash equilibria, but instead the more general notions of *Coarse Correlated Equilibrium* (CCE) and *Bayesian Coarse Correlated Equilibrium* (BCCE). As we do not rely on these notions for our analytical results, we refer the interested reader to e.g. [25] for details. Note that these notions can also be viewed as solution concepts in themselves for repeated play version of the games, and may arguably represent a more realistic model for long-run strategic play. While we will not directly analyze welfare bounds for these concepts, PoA *extension theorems* imply that the bounds we achieve carry over to the repeating setting [27].

3 PRICE OF ANARCHY

In this section, we provide characterizations of upper and lower bounds on the Price of Anarchy for (Greedy, GSP), (Greedy, VCG), and (Opt, GSP) with conservative bidders⁶. For upper bounds on the Price of Anarchy, we leverage the *semi-smoothness* framework of [3], itself a generalization of the *smoothness* framework of [24]. For lower bounds, we construct examples of equilibria⁷ that achieve less welfare than the optimal. For results that are primarily ancillary or require involved proofs, we provide proof sketches, and defer full proofs to an expanded online version of the paper.

For (Greedy, GSP) and (Greedy, VCG), we give a universal result – that is, under no requirements beyond the Ad Types setting – and this result matches known upper and lower bounds for the position auction (though our bounds are not yet as tight). For (Opt, GSP), we provide *instance-optimal* bounds; here, instance-optimal means allowing for dependence on the discount curves and number of slots but *not* over bidder valuations. It is very likely that our upper bounds on the Price of Anarchy in this setting are too pessimistic; we leave improvement of these bounds to future work.

	GSP	VCG
Greedy	2	3/2
Opt	4/3	NA

(a) Lower bounds on PoA.

	GSP	VCG
Greedy	4	4
Opt	$2 + 2(n-1) \frac{\delta_{\min}^{\max}}{\delta_{\min}}$ *	NA

(b) Upper bounds on PoA. * denotes instance-optimal bounds.

Table 1: Price of Anarchy Bounds

Our technique in each case will be to show that the game induced by the auction format and any valuation profile is *semismooth*, in the following sense:

⁶We omit (Opt, VCG); that bidding truthfully is a dominant strategy suggests alternative equilibria are unlikely in practice.

⁷Conservative bidders likely match reality when agents are not sophisticated or uncertain. This restriction is common in the literature, e.g. [3], and can be interpreted as strengthening our PoA lower bounds and weakening our upper bounds.

Definition 3.1 (Semismooth [3]). We say that a game is (λ, μ) -semismooth if there exists a (possibly randomized) strategy b' which depends only on a player's valuation such that:

$$u_i(b', \mathbf{b}_{-i}) \geq \lambda \sum_i \delta_{\tau_i}^{v(i)} v_i - \mu \sum_i \delta_{\tau_i}^{\sigma(i, \mathbf{b})} v_i.$$

for all bid profiles \mathbf{b} .

A game can be shown to be semismooth by showing that the the following inequality always holds:

$$u_i(b'_i, \mathbf{b}_{-i}) \geq \lambda \delta_{\tau(i)}^{v(i)} v_i - \mu \delta_{\tau(\pi(v(i), \mathbf{b}))}^{v(i)} v_{\pi(v(i), \mathbf{b})},$$

since if it holds, summing over players gives exactly the defining condition of semismoothness. And semismoothness directly yields Price of Anarchy bounds using the following theorem, from [3]:

THEOREM 3.2. *Suppose a game is (λ, μ) -semismooth, and social welfare is at least the sum of player utilities. Then its Price of Anarchy is upper bounded by $\frac{\mu+1}{\lambda}$.*

3.1 Greedy Allocation Proof Recipe

A common proof structure applies to both (Greedy, GSP) and (Greedy, VCG), because of their shared allocation algorithm and the fact that both pricing algorithms, when coupled with greedy allocation, guarantee that bidders are never overcharged. It is similar to the proof found in [3], but with additional subtlety due to the differing discount factors. We thus use the following Lemma:

LEMMA 3.3 (PARTIAL MONOTONICITY). *Suppose that \mathbf{b}, \mathbf{b}' are two bid profiles that only differ in element i , and $b'_i > b_i$. Let σ, σ' be the slots which i was assigned under \mathbf{b}, \mathbf{b}' respectively. Then under greedy allocation, we have that for each slot s strictly above σ :*

$$\delta_{\tau(\pi(s, \mathbf{b}'))}^s \mathbf{b}_{\pi(s, \mathbf{b}')}^s \geq \delta_{\tau(\pi(s, \mathbf{b}))}^s \mathbf{b}_{\pi(s, \mathbf{b})}^s$$

Informally, this lemma merely states that if bidder i deviates upwards from his bid under \mathbf{b} , the value obtained by players in the slots above his placement under \mathbf{b} can only increase. To see why this is true, recall that the greedy algorithm allocates from top to bottom. So for every slot s between σ' and σ (not including σ), the bidders considered when s was assigned under \mathbf{b} remain unallocated when considering s under \mathbf{b}' ; hence $\pi(s, \mathbf{b}')$ (i.e. whoever is assigned to s under \mathbf{b} must have at least as high of an effective value as $\pi(s, \mathbf{b}')$).

Now we can state the following theorem (proof in appendix):

THEOREM 3.4 (SEMI-SMOOTHNESS FOR GREEDY ALGORITHMS). *Let $(\mathcal{A}, \mathcal{P}_{\mathcal{A}})$ be an auction mechanism. Suppose that*

- (1) \mathcal{A} is the greedy algorithm, and
- (2) For any bid profile \mathbf{b} , for every bidder we have:

$$\mathcal{P}_{\mathcal{A}}(\mathbf{b})_i \leq b_i$$

Then $(\mathcal{A}, \mathcal{P}_{\mathcal{A}})$ is $(1/2, 1)$ -Semi Smooth.

3.2 Greedy Allocation and GSP Pricing

THEOREM 3.5. *Let $(\mathcal{A}, \mathcal{P}_{\mathcal{A}}) = (\text{Greedy}, \text{GSP})$. Then the Price of Anarchy is at most 4.*

PROOF. First, by assumption, \mathcal{A} is Greedy. Second, generalized second price will not charge a bidder more than their bid since under the greedy algorithm, the winner of a slot has a higher effective

bid than the second bidder's bid, which is what they are charged. Hence, the conditions of Theorem 3.4 are satisfied, so the induced game is $(\frac{1}{2}, 1)$ -semismooth and the bound follows. \square

On the other hand, the Price of Anarchy is at least 2.

THEOREM 3.6. *Let $(\mathcal{A}, \mathcal{P}_{\mathcal{A}}) = (\text{Greedy}, \text{GSP})$. Then the Price of Anarchy is at least 2.*

PROOF. Consider the following example: there are 2 slots and 2 bidders, one of type A and one of Type B. Let $\delta_A = (1, 0)$, $\delta_B = (1, 1)$, and let $v_A = (1 - \epsilon)v_B$, $\epsilon > 0$. Then the allocation (A, B) gets payoff $v_A + v_B = (2 - \epsilon)v_B$, while (B, A) gets welfare v_B .

We claim that the following is an equilibrium: A bids 0 and B bids v_B , giving the allocation (B, A) . To see that this is an equilibrium, notice that if these are the bids, $b_B > b_A$, so B will be given the first slot at a price of $b_A = 0$ for a total payoff of v_B . Since price is bounded below by 0, B could not gain by deviating any lower. On the other hand, in the second slot, A gets no value, but also is not charged, for a payoff of 0. To change anything, A would have to change the allocation, and so bid above $b_B = v_A$ - but then she would get a payoff of $v_A - v_B = (1 - \epsilon)v_B - v_B \leq 0$; hence she also would not like to switch. And note that since $0 \leq b_A$ and $v_B \leq v_A$, neither bidder is overbidding. But thus we see that

$$\frac{OPT}{EQ} = \frac{v_A + v_B}{v_B} = \frac{(2 - \epsilon)v_B}{v_B} = 2 - \epsilon$$

and so the Price of Anarchy can be made arbitrarily close to 2. \square

Note the equilibrium described is not unique - e.g., $b_A = (1 - \epsilon)v_B$, $b_B = v_B$ is also an equilibrium that achieves the same allocation.

3.3 Greedy Allocation and VCG Pricing

In this section, we consider the Price of Anarchy when $(\mathcal{A}, \mathcal{P}_{\mathcal{A}})$ is (Greedy, VCG). Again, using greedy allocation guarantees the first condition of Theorem 3.4. Though not obvious, bidders will not be overcharged, as we show in the following Lemma:

LEMMA 3.7. *Let $(\mathcal{A}, \mathcal{P}_{\mathcal{A}})$ be the greedy algorithm with VCG pricing. Then every bidder's charge will not exceed their effective bid.*

PROOF SKETCH. Let j^* be the bidder who would have taken i 's slot were i absent. It can be shown that i 's price p_i is exactly j^* 's value for i 's slot plus the difference between j^* 's price and j^* 's value for his slot. Combining with the fact that i must have greater value for his slot than j^* and applying strong induction to the assumption that bidders below i are not overcharged yields the claim. \square

Lemma 3.7 allows us to conclude that (Greedy, VCG) satisfies the conditions of Theorem 3.4, yielding the following Theorem:

THEOREM 3.8. *Let $(\mathcal{A}, \mathcal{P}_{\mathcal{A}}) = (\text{Greedy}, \text{VCG})$. Then the Price of Anarchy is at most 4.*

For lower bounds, we again find a suboptimal equilibrium:

THEOREM 3.9. *Let $(\mathcal{A}, \mathcal{P}_{\mathcal{A}}) = (\text{Greedy}, \text{VCG})$. The Price of Anarchy is at least $3/2$.*

PROOF SKETCH. Let $v_A = 1 + \epsilon$, $v_B = 1$, $v_C = 1 - \epsilon$. Let $\delta_A = (1, 1, 1 - 2\epsilon)$, $\delta_B = (1, 1, 0)$, $\delta_C = 1, \epsilon, \epsilon^2$. It can be shown that each player bidding their valuation is an equilibrium resulting in the allocation (A,B,C), for welfare $2 + \epsilon + \epsilon^2 - \epsilon^3$. But (C,B,A) gives welfare $3 - 2\epsilon - 2\epsilon^2$. With small enough ϵ , the ratio of optimal to equilibrium welfare can be driven arbitrarily close to $3/2$. \square

3.4 Optimal Allocation and GSP Pricing

In the case of optimal allocation and GSP pricing, we will obtain a smoothness result that depends on the largest and smallest discounts and the number of bidders, but not on the valuation profile. The result is as follows:

THEOREM 3.10. *Suppose $(\mathcal{A}, \mathcal{P}_{\mathcal{A}})$ is (Opt, GSP) . Then the game between bidders is $(\frac{1}{2}, \frac{\delta_{\max}}{\delta_{\min}}(n-1))$ -semismooth.*

To prove this result, we begin by observing that GSP pricing never charges a bidder more than his effective bid. Formally:

LEMMA 3.11. *In (Opt, GSP) , bid upper bounds price.*

COROLLARY 1. *The (Opt, GSP) mechanism has an instance-specific upper bound on Price of Anarchy of:*

$$PoA \leq 2 + 2(n-1) \frac{\delta_{\max}}{\delta_{\min}}.$$

If we assume that there are m slots and all discount curves are geometric and strictly ordered (e.g. $c_{\tau} = c_{\tau'}$ and $\delta_{\tau_1} \geq \delta_{\tau_2} \geq \dots \geq \delta_{\tau_k}$ for some k), then an upper bound is given by:

$$2 + 2 \cdot (n-1) \frac{\delta_{\tau_1}}{\delta_{\tau_k}^m}$$

We remark that this bound is potentially exponential in the number of bidders in the case of *geometric* discount curves, but linear in the case of *linear* discount curves (assuming a fixed set of discount curves). And while this bound is likely too pessimistic, we can give a lower bound as well:

THEOREM 3.12. *If $(\mathcal{A}, \mathcal{P}_{\mathcal{A}}) = (Opt, GSP)$, there is a conservative 3-bidder 3-slot example with competitive ratio arbitrarily close to $3/4$.*

PROOF SKETCH. Again, we construct a counterexample, prove it is an equilibrium, and optimize the welfare subject to equilibrium conditions. Here, a two-player two-slot example cannot suffer a high PoA, because the inefficient assignment of bidders would allow for a profitable deviation of the bidder in the worse slot (or the better slot if the price were too high). But with three bidders and three slots, we can find an example where two of the bidders effectively exert a “joint” externality, and no single bidder has any incentive to deviate despite the allocation being suboptimal overall. \square

Of all our results, this mechanism has the least-tight upper bound on the price of anarchy, and the weakest lower bound. But intuition suggests that the mechanism should perform relatively well: by construction, if the mechanism were given the true valuations, its allocation would be optimal. Without incentive-compatibility, it cannot assume this, but GSP, like VCG, does somewhat “protect” a bidder from the risk of overpaying. Thus bidders may have less incentive to greatly shade their bid. We leave formalizing this intuition and improving these PoA bounds to future work.

	GSP	VCG
Greedy	$(1 - \delta_A)v_A, (1 - \delta_B)v_B$	$\frac{1 - \delta_A}{1 - \delta_B}v_A, \frac{1 - \delta_B}{1 - \delta_A}v_B$
Opt	$(1 - \delta_A)v_A, (1 - \delta_B)v_B$	(v_A, v_B)

Table 2: 2 bidder, 2 type case, simple equilibrium strategies

4 EQUILIBRIUM CHARACTERIZATION

In this section, we report the first analytical characterization of Bayes-Nash equilibrium in the two-slot, two-bidder case with ad types and under the assumption that bidder values are drawn independently and from identical uniform distributions over the interval $[0, 1]$, with differing discount rates. For each auction type, we assume there are two slots, two discount types A and B with discount curves $(1, \delta_A)$ and $(1, \delta_B)$, i.e., geometric discount curves that both have a constant factor of 1. (This assumption can be easily relaxed at the cost of some extra notation.) Throughout, we will assume without loss of generality that $\delta_A < \delta_B$, and define $\Delta := \frac{1 - \delta_B}{1 - \delta_A} < 1$. For the sake of efficiency, we say a bidder ‘wins’ if they win the first slot. In this setting, we show the existence of simple equilibria that are symmetric in form and mostly natural; these equilibria are the *unique* linear equilibria. We provide the details for (Greedy, GSP) in the Appendix, and defer the (similar) proofs for the other mechanisms to an online appendix.

Comparison and Revenue Results Table 2 displays the simple linear equilibria we discover. Notice that in each setting, the *form* of the players’ strategies are symmetric up to relabeling, despite the fact that the particular strategy will differ due to different discount rates. Also, other than the VCG mechanism, each auction involves some shading. For GSP pricing, the downward shading coincides with each bidder’s marginal benefit of the first slot relative to the second. But under VCG pricing and greedy allocation, Bidder B shades down while Bidder A shades *up*.⁸

Table 3 gives the expected revenue for each of the equilibria described in Table 2. Several features are noteworthy. First, immediately we can see that both the two standard formats, as well as the two nonstandard formats, are (expected) revenue equivalent. This may be surprising, given the variation in payment rules and strategies; however, as can be shown via direct calculation, the strategies are such that the win condition and payment conditional on winning work out to be the same. Second, we note that if we allow $\delta_A = \delta_B = \delta$, we recover the equivalent revenue to the VCG mechanism for all four auction formats. This is because when discounts are the same and there are only two slots, the greedy allocation is equivalent to the optimal allocation, and GSP pricing coincides with externality pricing, so the matrix of auction formats collapses to a single row and column. Moreover, if $\delta = 0$, we recover the revenue of the standard second price auction with two uniform bidders (which is sensible, because if $\delta = 0$, the auction is effectively a second price auction for the only slot with any value). Finally, we note that it is not immediately obvious whether revenue increases or decreases with discount values (since Δ is a function of δ_A, δ_B); one can show via calculus that revenue decreases as either discount

⁸While this may be counterintuitive, note that with greedy allocation, bidding higher increases the win probability, and under GSP pricing, bidding higher does not (directly) increase the price paid. Of course, overbidding results in the possibility of winning at a price higher than one’s valuation, so there is a countervailing force stopping bidders from increasing their bids indefinitely.

	GSP	VCG
Greedy	$\frac{1-\delta_A}{6}\Delta^2 + \frac{1-\delta_B}{6}(3-2\Delta)$	$\frac{1-\delta_A}{6}\Delta^3 + \frac{1-\delta_B}{6}\Delta(3-2\Delta^2)$
Opt	$\frac{1-\delta_A}{6}\Delta^3 + \frac{1-\delta_B}{6}\Delta(3-2\Delta^2)$	$\frac{1-\delta_A}{6}\Delta^2 + \frac{1-\delta_B}{6}(3-2\Delta)$

Table 3: 2 bidder, 2 type case, equilibrium revenue

factor increases. This may be surprising, since bidders can derive more total welfare, but the principle is easy to see in the extreme: if $\delta = 1$, bidders may as well bid low⁹ and take the second slot.

Using Table 3 (and the fact that the *linear* equilibria are unique), straightforward, if involved, algebra allows us to make equilibrium¹⁰ revenue comparisons across auction formats:

THEOREM 4.1 (EQUILIBRIUM REVENUE). *Consider a two-bidder, two-type, two-slot setting with bidder valuations drawn from a standard uniform distribution. Then in linear equilibria:*

$$\mathcal{R}_{opt}^{vsg} = \mathcal{R}_{greedy}^{gsp} \geq \mathcal{R}_{greedy}^{vsg} = \mathcal{R}_{opt}^{gsp}$$

Importantly, these results only apply to our simple setting; it is unclear whether the revenue, welfare, or other predictions carry over into a general setting. While it is possible that more complicated analytic equilibrium may exist, it is difficult to foresee how such an equilibrium might be found analytically. Moreover, it is possible that equilibrium strategies, even if they do exist, are complicated to calculate and implement. Thus, in Section 5, we turn our attention to empirical study of revenue under realistic bid distributions, where (coarse correlated) equilibria are *learned* via no-regret learning techniques.

5 EMPIRICAL STUDY

In this section, we test our theoretical predictions of revenue and Price of Anarchy on both simulated and realistic data using *No-Regret Learning* (NRL) algorithms to model bidders. These algorithms have guarantees of convergence to (more general notions of) equilibrium, and have also been proposed as potential solution concepts in their own right [18]. We have two key results. First, the theoretical Bayes-Nash equilibria are (approximately) discovered by NRL bidders. Second, we find that, at least in the realistic data we study, the VCG mechanism tends to perform as well or better than the other mechanisms in terms of revenue and PoA. We also find that while there is significant variation in revenue, each mechanism performs better than the worst-case bounds predict.

Approach. We outline here the core commonality across the experiments. Each is based on NRL algorithms, which converge¹¹ to (Bayesian) *coarse correlated equilibrium* (CCE) (see, e.g. [25]). Though CCEs are more general than those we studied earlier in the paper, they may better reflect the real-world situation bidders face as they arise naturally from players learning to bid independently.

In each experiment and for each mechanism, we instantiate copies of the *exponential weights* (EW) algorithm to represent each

⁹We assume there is no reserve; we leave as an open problem questions around designing optimal auctions with ad types.

¹⁰By contrast, prior results in the position auction setting, such as that GSP prices are lower bounded by VCG prices for any *fixed* set of bids, cannot make predictions when bidders adjust their strategies to equilibrium.

¹¹We describe the formal meaning behind these statements, the technical details more broadly, and parameter specifications in a supplementary online appendix.

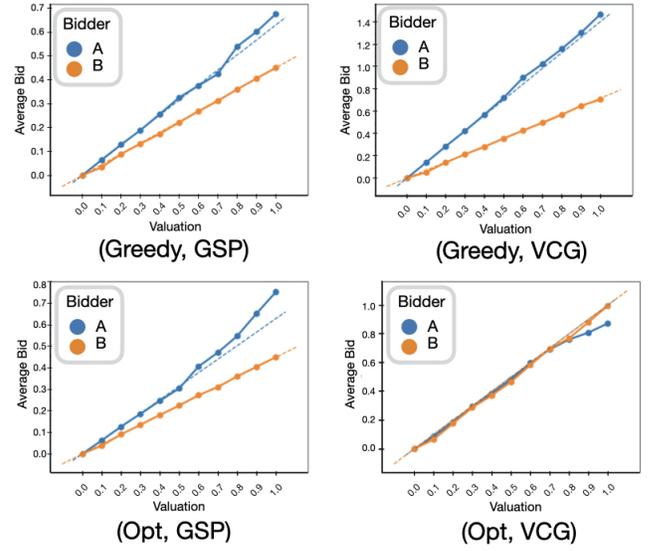


Figure 1: Mean bid by valuation in Bayesian setting (Experiment 1). Dashes indicate theoretical predictions.

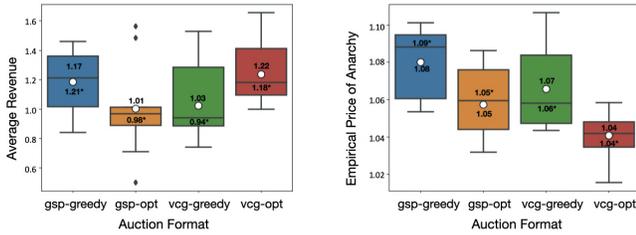
player. Play repeats over many rounds. At every round a distribution over bids from a discrete bid space. Each round, bidders draw their bids from this distribution, and the mechanism uses the bids to compute an allocation and price for each bidder. We record the total revenue and welfare under the realized outcome; for each player and each alternative potential bid, we also calculate a *counterfactual* outcome by re-running the mechanism with all *other* players' choices held fixed. The player then observes his counterfactual payoff and updates his distributions over actions accordingly.

Experiment 1: Bayesian no-regret learners converge to theoretical equilibria derived above.

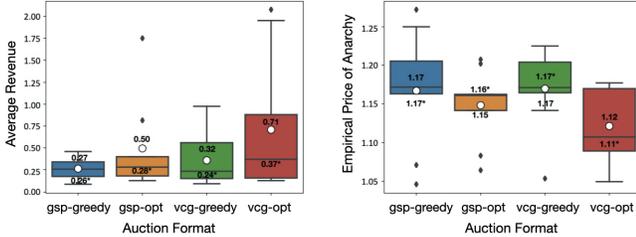
Experiment 1. Our first experiment approximates the two-bidder, two-slot, uniform distribution case we analyzed in Section 4. For this experiment, we adopt the *population interpretation* of Bayesian games, and so discretize the uniform valuation distributions into discrete uniform distributions over players with *fixed*, evenly-spaced valuations. Each round, nature selects a single player from each population, and each player maintains their own strategy. This approach is analyzed in [15]; we modify the approach by adding an extensive random exploration period. This modification is inspired by [11], which shows that including a sufficiently long exploration period in more standard auction settings allows bidders to provably converge to *specific* and natural Nash equilibria. In our case, we find that for each mechanism, we observe a very close correspondence between realized and theoretical bid distributions. Figure 1) displays the predicted bids (as a dashed line) and average observed bid for each valuation after a learning period.¹²

Experiments 2-3: With realistic data, there is wide variation in revenue, and VCG performs as well or better than other mechanisms. Each mechanism does well on PoA.

¹²We provide implementation details such as the additional protocols, the number of learning steps, and the discretization granularity of the bid and valuation spaces in the Appendix.



(a) Revenue (left) and Empirical Price of Anarchy (right) in Random Advertisers setting (Experiment 2).



(b) Revenue (left) and Empirical Price of Anarchy (right) in Random Auctions setting (Experiment 3).

Realistic Data for Experiments 2 and 3. We use real data from an online platform with a large advertising business and sample real bids to generate realistic valuation data. These are not *literal* valuation data for two reasons. First, bidders on the platform face a more complicated setting than modeled - for instance, bidders compete in multiple sequential and simultaneous auctions and rules may slightly vary. Second, we have modified the data to protect participant privacy. Hence, these bids are a reasonable proxy for real-world distributions, but may not be exactly such in practice.

The first dataset we collect is the *Random Advertisers* dataset, in which we sample 10 random advertisers who had between 100,000 and 200,000 impressions on a particular outlet and a day.¹³ For these 10 advertisers, we select 100,000 bids and normalize each advertiser’s bids to fall within the unit interval and clamp at the 5th and 95th percentiles. We can use this dataset to sample independently drawn valuations. The second is the *Random Auction* dataset: we again fix the outlet and day and randomly select 100,000 auctions, this time normalizing each auction individually. This dataset thus maintain correlation between bidders’ valuations in a given auction, which may be an important real-world feature of the domain.

Experiments 2 and 3. In Experiment 2, we use the Random Advertisers dataset. A protocol for a single round is as follows. We initialize an auction with 4 slots and 9 bidders of varying¹⁴ geometric discount factors (each with a fixed constant multiplier of 1). Each bidder has a valuation drawn independently from the Random Advertisers dataset, and is initialized with a fresh exponential weights algorithm over the (discretized) bidspace up to their valuation, as we would expect bids to be conservative in practice. After an exploration period, the players update via EW for 100 rounds, updating their bid distributions based on the realized and counterfactual bids each round. Then, we sample a bid profile from the time-averaged joint distribution by uniformly selecting a time period and drawing

a bid profile from the EW distributions of that given round; we estimate the revenue and welfare given that valuation profile by averaging across 200 such samples. For each auction format, we repeat this *entire* process for 200 valuation profile draws, and average these together as an estimate of the format’s revenue and welfare under the valuation distribution. We also compute the optimal allocation for each draw, and take the ratio of the average optimal value to the average welfare as an estimate of an empirical analogue (i.e. not worst-case) to the PoA. We plot revenue and empirical PoA in Figure 2a. Our third experiment is very similar to our second, except for the valuation sampling. Rather than sample valuations independently from the Random Advertiser dataset, we sample *auctions* randomly from the Random Auction dataset, and assign those bids to bidder valuations. Then we proceed as described in Experiment 1. We plot revenue and the Empirical PoA in Figure 2b.

Experiment 2 and 3 Results. Figures 2a and 2b show the average revenue and empirical price of anarchy for the Random Advertisers (Experiment 2) and Random Auctions (Experiment 3) settings.¹⁵ Notice first that there is a significant amount of variation (30% - 50%) in median revenue¹⁶ across the best- and worst-performing auctions. On the other hand, the gap between the best and worst auctions in terms of the PoA is on the order of 4%-5%, and all of them appear to be significantly better than the worst-case bounds proven in Section 3. Finally, while there appears to be a rough analogue to the revenue hierarchy of Table 3 in Experiment 2, the overall ranking of mechanisms seems to be dependent in general on the particular valuation distributions and data generation model. Taken together, these points highlight the importance in practice of testing the properties of these or other mechanisms “in the field”, i.e. against real bidders or at least realistic valuation distributions.

6 DISCUSSION AND OPEN QUESTIONS

We leave several open directions. In terms of Price of Anarchy: while we provide constant upper and lower bounds on the Price of Anarchy under greedy allocation, there remains a gap between these bounds. More substantially, while we provide a constant lower bound on the Price of Anarchy under optimal allocation with GSP pricing, our upper bound is instance-dependent and likely quite pessimistic; resolving this with either a constant upper bound, or identifying a family of arbitrarily bad examples, would be helpful. In terms of equilibrium characterization, it would be useful to identify (or rule out) analytical solutions in more complicated settings. In terms of empirics, understanding how our results would change as various features of the setting change would be valuable. For instance, even our large setting is still relatively small compared to modern instances encountered in online advertising today. Additionally, our understanding of how mechanism performance varies with discount curves in practice is not yet systematic; intuition suggests that bidders ought to bid less aggressively as their valuation of further slots increase, but how much less aggressively, and how this is affected by auction format, is unknown.

¹³Mobile Advertising, September 19, 2020

¹⁴Again, see the Appendix for details.

¹⁵We caveat these results by noting that because the bid space is high-dimensional, more samples of valuation profiles may be required for fidelity to the true distribution.

¹⁶We focus on medians due to the presence of outliers, but this is also true for means.

REFERENCES

- [1] ABRAMS, Z., GHOSH, A., AND VEE, E. Cost of conciseness in sponsored search auctions. In *Proc. of 3rd International Conference on Web and Internet Economics (2007)*, pp. 326–334.
- [2] AUSUBEL, L. M., AND MILGROM, P. The lovely but lonely vickrey auction. In *Combinatorial Auctions, chapter 1* (2006), MIT Press.
- [3] CARAGIANNIS, I., KAKLAMANIS, C., KANELLOPOULOS, P., KYROPOULOU, M., LUCIER, B., LEME, R. P., AND TARDOS, É. Bounding the inefficiency of outcomes in generalized second price auctions. *Journal of Economic Theory* 156 (2015), 343–388.
- [4] CAVALLO, R., SVIRIDENKO, M., AND WILKENS, C. A. Matching auctions for search and native ads. In *Proceedings of the 2018 ACM Conference on Economics and Computation* (2018), pp. 663–680.
- [5] CAVALLO, R., AND WILKENS, C. A. Gsp with general independent click-through-rates. In *Proc. of 10th International Conference on Web and Internet Economics (2014)*, T.-Y. Liu, Q. Qi, and Y. Ye, Eds., pp. 400–416.
- [6] CHAWLA, S., AND HARTLINE, J. D. Auctions with unique equilibria. In *Proceedings of the fourteenth ACM conference on Electronic commerce* (2013), pp. 181–196.
- [7] CLARKE, E. H. Multipart pricing of public goods. *Public choice* 11, 1 (1971), 17–33.
- [8] COLINI-BALDESCHI, R., LEONARDI, S., SCHRIJVERS, O., AND SODOMKA, E. Envy, regret, and social welfare loss. In *Proceedings of The Web Conference 2020* (2020), pp. 2913–2919.
- [9] COLINI-BALDESCHI, R., MESTRE, J., SCHRIJVERS, O., AND WILKENS, C. A. The ad types problem. In *Web and Internet Economics: 16th International Conference, WINE 2020, Ljubljana, Slovenia, December 17–20, 2017, Proceedings* (2020), Springer.
- [10] EDELMAN, B., OSTROVSKY, M., AND SCHWARZ, M. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American economic review* 97, 1 (2007), 242–259.
- [11] FENG, Z., GURUGANESH, G., LIU, C., MEHTA, A., AND SETHI, A. Convergence analysis of no-regret bidding algorithms in repeated auctions. In *The Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI-21)* (2021).
- [12] GOLDMAN, M., AND RAO, J. Experiments as instruments: Heterogeneous position effects in sponsored search auctions. *EAI Endorsed Trans. Serious Games* 3, 11 (2016), e2.
- [13] GOMES, R., AND SWEENEY, K. Bayes–nash equilibria of the generalized second-price auction. *Games and economic behavior* 86 (2014), 421–437.
- [14] GROVES, T. Incentives in teams. *Econometrica: Journal of the Econometric Society* (1973), 617–631.
- [15] HARTLINE, J., SYRGGANIS, V., AND TARDOS, E. No-regret learning in bayesian games. In *Advances in Neural Information Processing Systems* (2015), pp. 3061–3069.
- [16] JEZIORSKI, P., AND MOORTHY, S. Advertiser prominence effects in search advertising. *Management Science* 64, 3 (2018), 1365–1383.
- [17] KAPLAN, T. R., AND ZAMIR, S. Asymmetric first-price auctions with uniform distributions: analytic solutions to the general case. *Economic Theory* 50, 2 (2012), 269–302.
- [18] KLEINBERG, R. D., LIGETT, K., PILIOURAS, G., AND TARDOS, É. Beyond the nash equilibrium barrier. In *ICS* (2011), pp. 125–140.
- [19] KUHN, H. W. The hungarian method for the assignment problem. *Naval research logistics quarterly* 2, 1-2 (1955), 83–97.
- [20] LEME, R. P., AND TARDOS, E. Pure and bayes–nash price of anarchy for generalized second price auction. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science* (2010), IEEE, pp. 735–744.
- [21] LUCIER, B., AND PAES LEME, R. GSop auctions with correlated types. In *Proceedings of the 12th ACM conference on Electronic commerce* (2011), pp. 71–80.
- [22] MUNKRES, J. Algorithms for the assignment and transportation problems. *Journal of the society for industrial and applied mathematics* 5, 1 (1957), 32–38.
- [23] NARAYANAN, S., AND KALYANAM, K. Position effects in search advertising and their moderators: A regression discontinuity approach. *Marketing Science* 34, 3 (2015), 388–407.
- [24] ROUGHGARDEN, T. Intrinsic robustness of the price of anarchy. *Journal of the ACM (JACM)* 62, 5 (2015), 1–42.
- [25] ROUGHGARDEN, T. *Twenty lectures on algorithmic game theory*. Cambridge University Press, 2016.
- [26] ROUGHGARDEN, T., SYRGGANIS, V., AND TARDOS, E. The price of anarchy in auctions. *Journal of Artificial Intelligence Research* 59 (2017), 59–101.
- [27] SYRGGANIS, V., AND TARDOS, E. Composable and efficient mechanisms. In *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing* (New York, NY, USA, 2013), STOC '13, Association for Computing Machinery, p. 211–220.
- [28] VARIAN, H. R. Position auctions. *international Journal of industrial Organization* 25, 6 (2007), 1163–1178.
- [29] VICKREY, W. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance* 16, 1 (1961), 8–37.

A SUPPLEMENTARY MATERIAL

PROOF OF THEOREM 3.4. Recall that if for any bid profile:

$$u_i(b'_i, \mathbf{b}_{-i}) \geq \delta_{\tau(i)}^{v(i)} \frac{v_i}{2} - \delta_{\tau(\pi(v(i), \mathbf{b}))}^{v(i)} v_{\pi(v(i), \mathbf{b})}$$

then the statement holds. So suppose \mathbf{b} is a bid profile, and consider a deviation strategy of bidding half one's value. (Notice such a deviation guarantees a deviating bidder non-negative utility by Property 2.) Fix bidder i . Under this unilateral deviation, i receives $\sigma' := \mathcal{A}(b'_i, \mathbf{b}_{-i})_i$. There are two cases to consider: either $\sigma' \geq v(i)$ (i.e. σ' is $v(i)$ or better) or $\sigma' < v(i)$ (σ' is strictly worse than $v(i)$). If the first case holds, we achieve the desired inequality since:

$$\begin{aligned} u_i(b'_i, \mathbf{b}_{-i}) &= \delta_{\tau(i)}^{\sigma'} v_i - \delta_{\tau(i)}^{\sigma'} \mathcal{P}_i(b'_i, \mathbf{b}_{-i}, \mathcal{A}) \\ &\geq \delta_{\tau(i)}^{\sigma'} v_i - \delta_{\tau(i)}^{\sigma'} \frac{v_i}{2} = \delta_{\tau(i)}^{\sigma'} \frac{v_i}{2} \\ &\geq \delta_{\tau(i)}^{v(i)} \frac{v_i}{2} \geq \delta_{\tau(i)}^{v(i)} \frac{v_i}{2} - \delta_{\tau(\pi(v(i), \mathbf{b}))}^{v(i)} v_{\pi(v(i), \mathbf{b})}, \end{aligned}$$

where the first inequality follows by no-overcharging and the others by assumption or trivially.

Now suppose that instead, $\mathcal{A}(b'_i, \mathbf{b}_{-i})_i = \sigma' < v(i)$. We split this into two subcases. In the first subcase, $\frac{v_i}{2} \geq b_i$, i.e. b'_i is an upward deviation that results in i receiving σ' below $v(i)$. Combining (b'_i, \mathbf{b}_{-i}) into \mathbf{b}' , we can write:

$$\delta_{\tau(\pi(v(i), \mathbf{b}'))}^{v(i)} \mathbf{b}'_{\pi(v(i), \mathbf{b}')} \geq \delta_{\tau(i)}^{v(i)} \frac{v_i}{2}.$$

This follows because i was unallocated when $v(i)$ was considered, so if this did not hold, the greedy allocation would have allocated i to $v(i)$ instead of $\pi(v(i), \mathbf{b}')$.

Now, notice that we can view b_i as a downward deviation from b'_i , and a downward deviation cannot affect the allocation choices of any of the slots above its place before the deviation, including $v(i)$. But that means that the allocated bidder to $v(i)$ is the same under \mathbf{b} , so the inequality above also implies that:

$$\delta_{\tau(\pi(v(i), \mathbf{b}))}^{v(i)} \mathbf{b}_{\pi(v(i), \mathbf{b})} \geq \delta_{\tau(i)}^{v(i)} \frac{v_i}{2}.$$

Then using no-overcharging, we again have that

$$\delta_{\tau(i)}^{v(i)} \frac{v_i}{2} - \delta_{\tau(\pi(v(i), \mathbf{b}))}^{v(i)} \mathbf{b}_{\pi(v(i), \mathbf{b})} \leq 0 \leq u_i(b'_i, \mathbf{b}_{-i})$$

Finally, suppose $\frac{v_i}{2} < b_i$. As before, we must have that $\pi(v(i), \mathbf{b}')$ must have at least as high effective value as i . To see that $\pi(v(i), \mathbf{b})$ also has at least as high effective valuation as i , notice that we can view b_i as an upward deviation from b'_i . By assumption, $\sigma' < v(i)$, so Lemma 3.3 implies that in moving to \mathbf{b} , the values of bidders in slots above σ' , which include $v(i)$, must increase. But then:

$$\delta_{\tau(\pi(v(i), \mathbf{b}))}^{v(i)} \mathbf{b}_{\pi(v(i), \mathbf{b})} \geq \delta_{\tau(i)}^{v(i)} \frac{v_i}{2}.$$

and the desired inequality follows as before. \square

PROOF OF THEOREM 3.10. Consider deviation to $b'_i = v_i/2$, and let $v(i)$ be the slot of i under the optimal allocation. Then if i receives some slot $\sigma(i, b'_i, \mathbf{b}_{-i}) > v(i)$ by Lemma 3.11, we have:

$$\begin{aligned} u_i(b'_i, \mathbf{b}_{-i}) &= \delta_{\tau(i)}^{\sigma(i, b'_i, \mathbf{b}_{-i})} v_i - p_i \geq \delta_{\tau(i)}^{\sigma(i, b'_i, \mathbf{b}_{-i})} v_i - \delta_{\tau(i)}^{\sigma(i, b'_i, \mathbf{b}_{-i})} \frac{v_i}{2} \\ &= \delta_{\tau(i)}^{\sigma(i, b'_i, \mathbf{b}_{-i})} \frac{v_i}{2} \geq \delta_{\tau(i)}^{v(i)} \frac{v_i}{2} \end{aligned}$$

Otherwise, suppose deviating to b'_i gets i a slot $\sigma(i, b'_i, \mathbf{b}_{-i}) < v(i)$. Then since the allocation algorithm maximizes (apparent) welfare and allocating i to $v(i)$ was feasible, it must be that:

$$\delta_{\tau(i)}^{\sigma(i, b'_i, \mathbf{b}_{-i})} \frac{v_i}{2} + \sum_{j \neq i} \delta_{\tau(j)}^{\sigma(j, b'_i, \mathbf{b}_{-i})} v_j \geq \delta_{\tau(i)}^{v(i)} \frac{v_i}{2} + \sum_{j \neq i} \delta_{\tau(j)}^{v(j)} v_j$$

Here, the summation on the left-hand side is the apparent welfare (excluding i) given the allocation selected under deviation (b'_i, \mathbf{b}_{-i}) ; we will write this quantity as $\widehat{W}^{-i}(b'_i, \mathbf{b}_{-i})$. The summation on the right-hand side is what the apparent welfare (excluding i) would be if the (truly optimal) assignment v had been chosen instead; we will write this as \widehat{W}_v^{-i} . Then we write:

$$\delta_{\tau(i)}^{\sigma(i, b'_i, \mathbf{b}_{-i})} \frac{v_i}{2} \geq \delta_{\tau(i)}^{v(i)} \frac{v_i}{2} + \widehat{W}_v^{-i} - \widehat{W}^{-i}(b'_i, \mathbf{b}_{-i}).$$

Since the undiscounted price cannot exceed the bid (Lemma 3.11)

$$\begin{aligned} u_i(b'_i, \mathbf{b}_{-i}) &= \delta_{\tau(i)}^{\sigma(i, b'_i, \mathbf{b}_{-i})} v_i - p_i \geq \delta_{\tau(i)}^{\sigma(i, b'_i, \mathbf{b}_{-i})} \frac{v_i}{2} \\ &\geq \delta_{\tau(i)}^{v(i)} \frac{v_i}{2} + \widehat{W}_v^{-i} - \widehat{W}^{-i}(b'_i, \mathbf{b}_{-i}). \end{aligned}$$

We can drop \widehat{W}_v^{-i} and still have a true inequality, so we focus on how different $\widehat{W}^{-i}(b'_i, \mathbf{b}_{-i})$ can be from $W^{-i}(\mathbf{b})$. And since we assume conservative bids, we must have $\widehat{W}^{-i}(b'_i, \mathbf{b}_{-i}) \leq W^{-i}(b'_i, \mathbf{b}_{-i})$. Hence, we can rewrite the inequality we have as:

$$u_i(b'_i, \mathbf{b}_{-i}) \geq \delta_{\tau(i)}^{v(i)} \frac{v_i}{2} - W^{-i}(b'_i, \mathbf{b}_{-i})$$

Now, we need to bound $W^{-i}(b'_i, \mathbf{b}_{-i})$ in terms of $W(\mathbf{b})$. We will do this very coarsely. Notice that in any allocation, the algorithm will always fill all the slots. Let δ_{\max} be the maximum discount rate in the first slot - that is, $\max_j \delta_j^{\max}$ - and let δ_{\min} be the minimum discount rate for the *last* slot (i.e. $\min_j \delta_j^{\min}$). By monotonicity and full allocation, we know then that at the very most, we have:

$$W^{-i}(b'_i, \mathbf{b}_{-i}) \leq \sum_{j \neq i} \delta_j^{\max} v_j = \delta_{\max} \sum_{j \neq i} v_j$$

and at the very least, we have:

$$W(\mathbf{b}) \geq \sum_j \delta_j^{\min} v_j = \delta_{\min} \sum_j v_j.$$

But that means that whatever $W(\mathbf{b})$ is, we must have that:

$$W^{-i}(b'_i, \mathbf{b}_{-i}) \leq \frac{\delta_{\max}}{\delta_{\min}} \cdot \frac{\sum_{j \neq i} v_j}{\sum_j v_j} W(\mathbf{b}).$$

(To see this, note that $1 \leq W(\mathbf{b}) / (\delta_{\min} \sum_j v_j)$, multiply the inequality with W^{-i} by 1, and apply this inequality to $\delta_{\max} \sum_{j \neq i} v_j \cdot 1$.)

But now, using this upper bound for W^{-i} to upper bound the negative term in the inequality above, we can write that

$$u_i(b'_i, \mathbf{b}_{-i}) \geq \delta_{\tau(i)}^{v(i)} \frac{v_i}{2} - \frac{\delta_{\max}}{\delta_{\min}} \frac{\sum_{j \neq i} v_j}{\sum_j v_j} W(\mathbf{b}). \quad (1)$$

Inequality 1 thus holds in the case that i gets a worse slot than $v(i)$ under the deviation, but of course it also holds true in the case that

i gets a better slot. Thus, it always holds, so we can write:

$$\begin{aligned} \sum_i u_i(b'_i, \mathbf{b}_{-i}) &\geq \sum_i \delta^{v(i)} \frac{v_i}{2} - \frac{\delta^{\max}}{\delta^{\min}} \sum_i \frac{\sum_{j \neq i} v_j}{\sum_j v_j} W(\mathbf{b}) \\ &= \frac{\text{OPT}}{2} - \frac{\delta^{\max}}{\delta^{\min}} W(\mathbf{b}) \frac{1}{\sum_j v_j} \sum_i \sum_{j \neq i} v_j \\ &= \frac{\text{OPT}}{2} - \frac{\delta^{\max}}{\delta^{\min}} W(\mathbf{b}) \frac{1}{\sum_j v_j} \cdot (n-1) \sum_j v_j \end{aligned}$$

where the last inequality follows since each agent's valuation appears exactly $n-1$ times over the double sum. But then we have:

$$\sum_i u_i(b'_i, \mathbf{b}_{-i}) \geq \frac{\text{OPT}}{2} - \frac{\delta^{\max}}{\delta^{\min}} (n-1) W(\mathbf{b}).$$

Thus, this game is $(\frac{1}{2}, \frac{\delta^{\max}}{\delta^{\min}}(n-1))$ -semismooth. \square

THEOREM A.1. *In the two slot, two bidder, uniform case with $(\mathcal{A}, \mathcal{P}_{\mathcal{A}}) = (\text{Greedy}, \text{GSP})$, the strategy profile $((1-\delta_A)v_A, (1-\delta_B)v_B)$ is a Bayes-Nash equilibrium. Among linear equilibria, it is unique.*

PROOF OF THEOREM A.1. Consider Bidder A's perspective after she learns her valuation v_A . If A wins, she pays b_B and gets value v_A ; if she loses, she gets $\delta_A v_A$ and pays nothing. Then:

$$\begin{aligned} \mathbb{E}_{v_B \sim U[0,1]} [u_A | b_A] &= (v_A - \mathbb{E}[b_B | b_B < b_A]) \Pr[b_B < b_A] \\ &\quad + \delta_A v_A (1 - \Pr[b_B < b_A]) \end{aligned} \quad (2)$$

$$(3)$$

Since we wish to show that $(1-\delta_A)v_A$ is a best-response to $(1-\delta_B)v_B$, we can assume that $b_B = (1-\delta_B)v_B$. Hence, A wins if and only if $v_B < b_A/(1-\delta_B)$. Under the uniform distribution, $\Pr[x < c] = \min\{c, 1\}$ and $\mathbb{E}[x | x < c] = \frac{\min\{c, 1\}}{2}$ for $c > 0$. Thus we can apply these to Equation 2 to write:

$$\begin{aligned} \mathbb{E}_{v_B \sim U[0,1]} [u_A | b_A] &= \left(v_A - (1-\delta_B) \frac{b_A}{2(1-\delta_B)} \right) \frac{b_A}{1-\delta_B} + \delta_A v_A \left(1 - \frac{b_A}{1-\delta_B} \right) \\ &= \frac{v_A b_A}{1-\delta_B} - \frac{b_A^2}{2(1-\delta_B)} - \frac{\delta_A b_A v_A}{1-\delta_B} + \delta_A v_A, \end{aligned} \quad (4)$$

which we call $u_A(b_A)$ for brevity, whenever $b_A \leq (1-\delta_B)$, and or the "cap" $u_A = v_A - \frac{1-\delta_B}{2}$ otherwise. In principle, we must check whether $\text{argmax} u_A$ over $[0, 1-\delta_B]$ gives a better payoff than the cap. But as $u_A(1-\delta_B) = v_A - \frac{1-\delta_B}{2}$, increasing b_A beyond $1-\delta_B$ cannot improve payoff, so it suffices to simply find the maximum of u_A over $[0, 1-\delta_B]$. Notice that $u_A(b_A)$ is continuous and on $[0, 1-\delta_B]$, differentiable with:

$$u'_A(b_A) = \frac{v_A(1-\delta_A)}{1-\delta_B} - \frac{b_A}{1-\delta_A} \quad u''_A(b_A) = -\frac{1}{1-\delta_A}$$

Choosing $b_A^* := (1-\delta_A)v_A < 1-\delta_B$ satisfies the first order condition; since u_A is strictly concave on $[0, 1-\delta_B]$, b_A^* is there a global maximum. On the other hand, if $b_A^* \geq (1-\delta_B)$, then because u_A is increasing right up until $(1-\delta_B)$, u_A takes its maximum at $b_A = (1-\delta_B)$. But, bidding $(1-\delta_A)v_A$ results in the same payoff as bidding $1-\delta_B$ (because of the "cap"). Thus, regardless of what v_A is, the strategy $b_A^* = (1-\delta_A)v_A$ is a best-response if B is bidding $(1-\delta_B)v_B$. Reversing roles and considering B's perspective gives the same logic. Hence, the pair of strategies form an equilibrium. To see uniqueness among linear equilibria, notice that as long as b_B is linear, i.e. $b_B(v_B) = \beta v_B$ for some fixed $0 \leq \beta \leq 1$, Equation 4 holds, and the particular choice of β cancels out just as it did for $(1-\delta_B)$; hence, again, the optimal bid will be $(1-\delta_A)v_A$. A similar

argument holds for B, so any linear equilibrium must prescribe the same strategies as those identified here. \square

ALGORITHM 1: Abstract protocol for Experiment 2

```

for  $(\mathcal{A}, \mathcal{P}_{\mathcal{A}}) \in \{\text{Greedy}, \text{Opt}\} \times \{\text{GSP}, \text{VCG}\}$  do
  for  $t \in 1, 2, \dots, 200$  do
    Draw 9 valuations. Initialize bidders with values and fresh
    Exp. Weights.
    for  $t \in 1, 2, \dots, 100$  do
      Initialize a 4 slot auction using  $(\mathcal{A}, \mathcal{P}_{\mathcal{A}})$  and draw and fix
      bids. for  $i \in \mathcal{I}$  do
        for  $b' \in \{\frac{1}{d}, \frac{2}{d}, \dots, 1\}$  do
          Re-run auction with all other players' bids fixed,
          but player  $i$  using  $b'$ . Save payoff.
        end
      Update ExpWeights.
    end
  end
  for  $t' \in 1, 2, \dots, 200$  do
    Choose number uniformly at random from the total
    number of learning steps to sample from.
    Sample bid for each bidder from time average of
    ExpWeights at that round.
    Run auction with these bids and save outcomes.
  end
end

```

Variable	Description	Experiment 1
Eq?	Equilibrium Concept	BCCE
\mathcal{D}	Bid Datasets	Synthetic
d	Bid Discretization	20
$\mathcal{D}(v)?$	Discretization depends on value?	✓
M	Number of Bidders	2
$ \{v_b\} $	Number of valuations per bidder	11
S	Number of Slots	2
N_s	Number of Sampled valuations	NA
N_l	Number of Learning Steps	500000
N_t	Number of Testing Steps	NA
N_e	Number Exploration Steps	10367
OB?	Allow Overbidding?	✓
δ_0	Geometric Discount Constant	1
δ	Geometric Discount Factors	[.37, .85]

Table 4: Experimental Parameters - Experiment 1

Variable	Experiment 2	Experiment 3
Eq?	CCE	CCE
\mathcal{D}	Random Advertisers	Random Auction
d	20	20
$\mathcal{D}(v)?$	✓	✓
M	9	9
$ \{v_b\} $	1	1
S	4	4
N_s	200	200
N_l	100	100
N_t	200	200
N_e	0	0
OB?	✗	✗
δ_0	1	1
δ	[0.9, 0.9, 0.8, 0.8, 0.7, 0.7, 0.6, 0.6, 0.5]	←←←

Table 5: Experimental Parameters - Experiments 2 and 3