

Pacing Equilibrium in First-Price Auction Markets

Vincent Conitzer^{*1}, Christian Kroer^{†2}, Debmalya Panigrahi^{‡1}, Okke Schrijvers^{§2}, Eric Sodomka^{¶2}, Nicolas E. Stier-Moses^{||2}, and Chris Wilkens^{**3}

¹Computer Science Department, Duke University

²Core Data Science, Facebook Inc

³Tremor Technologies

November 20, 2018

Abstract

In ad auctions—the prevalent monetization mechanism of Internet companies—advertisers compete for online impressions in a sequential auction market. Since advertisers are typically budget-constrained, a common tool employed to improve their ROI is that of *pacing*, i.e., uniform scaling of their bids to preserve their budget for a longer duration. If the advertisers are excessively paced, they end up not spending their budget, while if they are not sufficiently paced, they use up their budget too soon. Therefore, it is important that they are paced at just the right amount, a solution concept that we call a *pacing equilibrium*. In this paper, we study pacing equilibria in the context of first-price auctions, which are popular in the theory of ad mechanisms. We show existence, uniqueness, and efficient computability of first-price pacing equilibria (FPPE), while also establishing several other salient features of this solution concept. In the process, we uncover a sharp contrast between these solutions and second price pacing equilibria (SPPE), the latter being known to produce non-unique, fragile solutions that are also computationally hard to obtain. Simulations show that FPPE have better revenue properties than SPPE, that bidders have lower ex-post regret, and that incentives to misreport budgets for thick markets are smaller.

Introduction

Advertising has emerged as the primary method of monetization for online services, search engines, social media platforms, etc. Large platforms implement their own auction markets to control ad placement and pricing. Abstractly, an auction market is a set of rules that governs the allocation of *impressions*—opportunities to show ads to online users—to *bidders*—advertisers who pay the platform to show impressions. Typically, bidders specify a *target* set of users, as well as the *price* they are willing to pay for each impression. They also set a *budget* that limits their total expenditure within a given time interval, say during a single day. Given advertisers’ targeting criteria, bids, and budgets, the role of the platform is to allocate impressions and set prices that respect bidders’ bids and budgets.

*conitzer@cs.duke.edu

†chrkroer@fb.com

‡debmalya@cs.duke.edu

§okke@fb.com

¶sodomka@fb.com

||stier@fb.com

**c.a.wilkens@gmail.com

A simple implementation consists in running an independent auction for each impression, removing a bidder from the set of participants once her budget runs out. However, this as-soon-as-possible approach has several downsides. A bidder who makes large bids but has a small budget will quickly expend it on the first impressions that become available. From the perspective of maximizing overall utility in the system—i.e., social welfare—this is suboptimal because the bidder will likely win *generic* impressions of interest to many other bidders instead of winning *specific* impressions that the bidder uniquely values. Similarly, this is suboptimal for the bidder because the increased competition on generic impressions will induce higher prices, leading to lower utility than if the bidder had bought specific impressions.

There are two broad approaches to taking budgets into consideration: bidder selection and bid modification. The goal in bidder selection is to choose a subset of participants for each auction, from the set of bidders whose budgets have not been exhausted, in a way that optimizes the use of bidders’ overall budgets. This has been considered in both first and second price settings. In the context of first price auctions, the celebrated work of Mehta et al. (2007) gives an algorithm for online allocation of impressions to bidders to maximize overall ad revenue. Their allocation algorithm, and those in the large body of follow-up work (see the survey by Mehta (2013)), can be interpreted as running a first price auction for each impression after removing a subset of bidders from the auction based on their remaining budgets. Bidder selection has also been explored in the context of Generalized Second Price (GSP) auctions, particularly for multi-objective optimization in search engines (Abrams et al., 2008; Azar et al., 2009; Goel et al., 2010; Karande, Mehta, and Srikant, 2013). An important feature of bidder selection is that a bidder who is chosen to participate in the auction does so with her original bid, i.e., the platform does not modify bid values of participating bidders.

This paper focuses on bid modification, wherein the platform shades an advertiser’s bids in order to preserve her budget for the future. This is commonly implemented by scaling an advertiser’s bid by a *pacing multiplier* (of value < 1). Many ad platforms provide a free option to advertisers to automatically have their bids scaled, and an impressive body of work has focused on advertiser strategies for bid modification in order to maximize their ROI (Rusmevichientong and Williamson, 2006; Borgs et al., 2007; Cary et al., 2007; Feldman et al., 2007; Hosanagar and Cherepanov, 2008). Recent work has studied pacing equilibria in the context of second-price auctions: Balseiro et al. (2017) study several budget smoothing methods including multiplicative pacing in a stochastic context where ties do not occur; Balseiro and Gur (2017) study how an individual bidder might adapt their pacing multiplier over time in a stochastic continuous setting; Conitzer et al. (2018) study second price pacing when bidders, goods, budgets, and valuations are known, and prove that equilibria exist under fractional allocations. In this paper, we study pacing in the context of first price auctions. While second price auctions displaced first price auctions in Internet advertising because of their many desirable robustness guarantees, particularly related to stability (Edelman and Ostrovsky, 2007) and strategyproofness, first price auctions are regaining popularity because they are simple to operate and offer a degree of transparency and ease of use that second price auctions do not (Chen, 2017; Sluis, 2017). In the context of position auctions, Dütting, Fischer, and Parkes (2018) further showed that first price does not suffer from the same equilibrium selection problems that GSP and VCG does. Moreover, as we will see later, first price auctions provide a clean characterization of equilibrium solutions in the context of pacing, unlike in the case of second price auctions (Conitzer et al., 2018). Indeed, as stated earlier, perhaps the most robust literature in ad auctions relates to bidder selection in first price auctions (Mehta, 2013). Our work complements this line of work by focusing on bid modification instead of bidder selection as the preferred method for budget management.¹

¹In Mehta et al. (2007) and most other related papers, bidder selection is performed by using a “scaling parameter” for bids based on remaining budgets, but once the winner of an auction is selected, she pays her entire (unscaled) bid

Contributions

We are given a set of bidders, goods, valuations, and budgets. A first price auction is conducted for each individual good. All bidders participate in every auction, but we are allowed to choose per-bidder *pricing multipliers* α_i which scale their bids. A set of pricing multipliers is said to be *budget feasible* (abbreviated BFPM) if no bidder’s total payments exceeds her budget. We say that a set of BFPM is *maximal* if they dominate any other BFPM. Furthermore, we say that a set of BFPM yields a *pricing equilibrium* if the pricing multipliers are simultaneously optimal for all bidders in the following sense: *every bidder either spends her entire budget or her pricing multiplier is set to 1*. We will allow goods to be allocated fractionally if there are ties; this is well-motivated in e.g. the ad-auction setting, where goods may be thousands of impressions. Note that pricing equilibria can be defined in both first and second price settings – we abbreviate these respectively by FPPE and SPPE. The goal of this paper is to characterize FPPE, study their properties, and design algorithms to find them.

Existence and Uniqueness. A priori, it is not clear if an FPPE always exists. Our first result shows that not only does an FPPE always exist, but also that it is *essentially unique*.² In fact, we show that the FPPE exactly coincides with the (unique) *maximal* set of pricing multipliers. This also leads to the observation that the FPPE is revenue-maximizing among all BFPM. Furthermore, we also show that the FPPE yields a *competitive equilibrium*, i.e., every bidder is exactly allocated her *demand set*. It is worth noting the contrast with SPPE, which are not unique in general (Conitzer et al., 2018).

Computability. We show an interesting connection between FPPE and a generalization of the classical Eisenberg-Gale (EG) convex program for quasi-linear utilities (Cole et al., 2017). Using Fenchel duality, we use the EG program to infer that the unique maximal BFPM, which is also revenue maximizing, yields the unique FPPE. Moreover, this connection with the EG program immediately yields a (weakly) polynomial algorithm to compute the FPPE. This also contrasts with SPPE, for which maximizing revenue (or other objectives like social welfare) is known to be NP-hard (Conitzer et al., 2018). Indeed, no polynomial-time algorithm is known for even finding an arbitrary pricing equilibrium in the second price setting.

Monotonicity and Sensitivity: We show that FPPE satisfies many notions of monotonicity, both in terms of revenue and social welfare. Adding an additional good weakly increases both revenue and social welfare, and adding a bidder or increasing a bidder’s budget weakly increases revenue. Note that this also distinguishes FPPE from SPPE, which generally do not satisfy such monotonicity conditions.

In fact, not only is the revenue obtained in an FPPE monotonically non-decreasing with increasing budgets, but it also changes smoothly in the sense that increasing the budget by some Δ can only increase the revenue by Δ . Again, this does not hold in SPPE where the revenue can increase by a substantial amount even for a small change in the budgets.

Shill-proofness: Using monotonicity, we also establish that there is no incentive for the provider to enter fake bids (shill-proof), unlike in SPPE.

Simulations: Using our convex program we perform simulations on synthetic data. We find that the ex-post regret associated with FPPE gets low quickly as markets get thick, especially when bidders

for the impression. In contrast, in our work, the individual bids are scaled using the respective pricing multipliers, and the winner pays the scaled bid instead of her original bid for the impression. Another point of difference is that the scaling multiplier changes from one impression to the next in the budgeted allocation literature since it is only a tool for winner selection, whereas we choose a single pricing multiplier for a bidder that is used to scale her bids for all the impressions in the problem instance.

²More precisely, pricing multipliers are unique but there may be different equivalent allocations due to tie breaking. This does not impact revenue, social welfare or individual utilities.

can only change the scale of their utilities, as is usually the case in ad markets when bidding on conversions (e.g. clicks). We furthermore show that even with just a bit of market thickness bidders have no incentive to misreport budgets in order to shift the FPPE outcome. Finally, we compare FPPE and SPPE revenue and social welfare across instances. We find that FPPE always has better revenue, while welfare splits evenly on which solution concept performs better.

First-price Pacing Equilibria

We consider a single-slot auction market in which a set of bidders $N = \{1, \dots, n\}$ target a set of (divisible) goods $M = \{1, \dots, m\}$. Each bidder i has a valuation $v_{ij} \geq 0$ for each good j , and a budget $B_i > 0$ to be spent across all goods. We assume that the goods are sold through independent (single slot) first-price auctions, and the valuations and budgets are assumed to be known to the auctioneer. When multiple bids are tied for an item we assume that the item can be fractionally allocated, and we allow the auctioneer to choose the fractional allocation (though our later results on equivalence to competitive equilibrium show that the fractional choices are optimal for the bidders as well).

The goal is to compute a vector of *pacing multipliers* α that smooths out the spending of each bidder so that they stay within budget. A pacing multiplier for a bidder i is a real number $\alpha_i \in [0, 1]$ that is used to scale down the bids across all auctions: for any i, j , bidder i participates in the auction for good j with a bid equal to $\alpha_i v_{ij}$; we refer to these bids as *multiplicatively paced*. We define feasibility as follows:

Definition 1. A set of budget-feasible first-price pacing multipliers (BFPM) is a tuple (α, x) , of pacing multipliers $\alpha_i \in [0, 1]$ for each bidder $i \in N$, and fractional allocation $x_{ij} \in [0, 1]$ for bidder $i \in N$ and good $j \in M$ with the following properties:

- (Prices) Unit price $p_j = \max_{i \in N} \alpha_i v_{ij}$ for goods $j \in M$.
- (Goods go to highest bidders) If $x_{ij} > 0$, then $\alpha_i v_{ij} = \max_{i' \in N} \alpha_{i'} v_{i'j}$ for each bidder $i \in N$ and good $j \in M$.
- (Budget-feasible) $\sum_{j \in M} x_{ij} \cdot p_j \leq B_i$ for bidders $i \in N$.
- (Demanded goods sold completely) If $p_j > 0$, then $\sum_{i \in N} x_{ij} = 1$ for each good $j \in M$.
- (No overselling) $\sum_{i \in N} x_{ij} \leq 1$ for each good $j \in M$.

Within the feasible space, we're particularly interested in outcomes that are stable in some sense. Specifically, solutions where no bidder is unnecessarily paced, which we call first-price pacing equilibria (FPPE).

Definition 2. A first-price pacing equilibrium (FPPE) is a BFPM tuple (α, x) , of pacing multipliers α_i for each bidder i , and fractional allocation x_{ij} for bidder i and good j with these additional properties:

- (No unnecessary pacing) If $\sum_{j \in M} x_{ij} p_j < B_i$, then $\alpha_i = 1$ for each bidder $i \in N$.

Existence, uniqueness, and structure of FPPE

A priori it's not obvious whether there are BFPM that satisfy the additional constraint of no unnecessary pacing. We show that not only are FPPE guaranteed to exist, they are also unique and maximize the seller's revenue over all BFPM. In the following, the inequality symbols are *component-wise*.

Lemma 1. *There exists a Pareto-dominant BFPM (α, x) (i.e., $\alpha \geq \alpha'$ component-wise for any other BFPM (α', x')).*

Proof. First, we will show that given any two BFBM $(\alpha^{(1)}, x^{(1)})$ and $(\alpha^{(2)}, x^{(2)})$, there exists a BFBM with pacing multipliers $\alpha^* = \max(\alpha^{(1)}, \alpha^{(2)})$ that are the component-wise maximum of $\alpha^{(1)}$ and $\alpha^{(2)}$. Note that, on each item, the resulting paced bid for a bidder is the higher of her two paced bids on this item in the original two BFBMs. Let the corresponding allocation x be: for each good j , identify which of the two BFBMs had the highest paced bid for j , breaking ties towards the first BFBM; then, allocate the good to the same bidder as in that BFBM (if the good was split between multiple bidders, allocate it in the same proportions), at the same price. Note that these prices coincide with the winning bidders' paced bids in the new solution. Thus, we charge the correct prices, goods go to the highest bidders, demanded goods are sold completely, and there is no overselling.

All that remains to be verified is that the new solution is budget-feasible. Consider bidder i in the bidder-wise-max α^* . In either $\alpha^{(1)}$ or $\alpha^{(2)}$, bidder i had the exact same multiplier—say it was in $\alpha^{(b)}$ for $b \in \{1, 2\}$ (breaking ties towards 1); the others' multipliers in $\alpha^{(b)}$ were the same or lower compared to α^* . Thus, if i did not have the highest bid on an item j in $\alpha^{(b)}$, then i also does not win this item under α^* . If i did have the highest bid on an item j in $\alpha^{(b)}$, then in x^* , at most i can win the same fraction of the item as in $x^{(b)}$, at the same price. This is because of the following reasons. If i is winning a nonzero fraction of j in x^* , then either the highest bid on j under $\alpha^{(b)}$ was strictly higher than that under the other BFBM, in which case the same allocation as in $x^{(b)}$ is used; or, the highest bid on j was the same in both of the original BFBMs and the tie was broken towards the allocation of j in $x^{(1)}$, which implies that $\alpha_i^{(1)} v_{ij} \geq \alpha_i^{(2)} v_{ij}$ and therefore $b = 1$. (It cannot be the case that the highest bid on j under $\alpha^{(b)}$ was strictly lower than that under the other BFBM, because in that case bidder i 's bid would not be the highest under α^* .) It follows that i spends no more than under BFBM b , which is budget-feasible. Hence, the new BFBM is budget-feasible.

We now complete the proof. Let $\alpha_i^* = \sup\{\alpha_i \mid \alpha \text{ is part of a BFBM}\}$. We will show that α^* is part of a BFBM, proving the result. For any $\epsilon > 0$ and any i , there exists a BFBM where $\alpha_i > \alpha_i^* - \epsilon$. By repeatedly taking the component-wise maximum for different pairs of i , we conclude there is a single BFBM $(\alpha^\epsilon, x^\epsilon)$ such that for every i , $\alpha_i^\epsilon > \alpha_i^* - \epsilon$. Because the space of combinations of multipliers and allocations is compact, the sequence $(\alpha^\epsilon, x^\epsilon)$ (as $\epsilon \rightarrow 0$) has a limit point (α^*, x^*) . This limit point satisfies all the properties of a BFBM by continuity. \square

In addition to there being a maximal set of pacing multipliers over all BFBM, this maximal BFBM is actually an FPPE.

Lemma 2 (Guaranteed existence of FPPE). *The Pareto-dominant BFBM (α, x) with maximal pacing multipliers α has no unnecessarily paced bidders, so it forms a FPPE.*

Proof. We prove the statement by contradiction. If bidder i is paced unnecessarily under maximal BFBM (α, x) , either it tied with another bidder on some good j or not. If it's not tied for any good, then we can increase bidder i 's pacing multiplier by a sufficiently small $\epsilon > 0$ such that i is still not tied for any item and still within budget, contradicting the fact that α was maximal. So bidder i must be tied for at least some good. Define $N(i)$ as all the bidders that are tied for any good with bidder i , i.e. $N(i) = \{\text{bidder } k : \exists \text{ good } j \text{ with } \alpha_i \cdot v_{ij} = \alpha_k \cdot v_{kj}\}$. Now take the transitive closure T , i.e., include all bidders who are tied for an item with a bidder in $N(i)$, etc. Once one has the transitive closure, redistribute the items that are tied such that none of the bidders in T are budget constrained (while still allocating items completely). Simultaneously increase the pacing multipliers of all bidders in T by $\epsilon > 0$ so that all bidders in T are still not budget-constrained, and no new ties are created; call this set of pacing multipliers α' and the redistributed goods x' . This contradicts that α was the maximal BFBM to begin with, as (α', x') is a BFBM yet it has pacing multipliers that are higher than in (α, x) . \square

The converse of Lemma 2 is also true: any BFBM for which at least one bidder has a pacing multiplier α_i lower than the maximal BFBM must have an unnecessarily paced bidder.

Lemma 3. Consider two BFPM $(\alpha^{(1)}, x^{(1)})$ and $(\alpha^{(2)}, x^{(2)})$, where $\alpha^{(1)} \geq \alpha^{(2)}$ and $\alpha_i^{(1)} > \alpha_i^{(2)}$ for some bidder i . Then, $(\alpha^{(2)}, x^{(2)})$ must have an unnecessarily paced bidder.

Proof. Consider the set I of all bidders whose pacing multipliers are strictly lower in $\alpha^{(2)}$ than in $\alpha^{(1)}$ (by definition there must be at least one bidder in this set). Collectively, I wins fewer (or the same) items under $\alpha^{(2)}$ than under $\alpha^{(1)}$ (the bids from outside I have stayed the same, those from I have gone down), and at lower prices. Since $(\alpha^{(1)}, x^{(1)})$ was budget feasible, I was not breaking its collective budget before. Since I is spending strictly less, at least some bidder must not spend their entire budget and thus is unnecessarily paced. \square

This implies that the pacing multipliers of FPPE are uniquely determined.

Corollary 1 (Essential Uniqueness). *The pacing multipliers of any FPPE are uniquely determined and correspond to the pacing multipliers of the maximal BFPM.*

While the pacing multipliers are uniquely determined, the allocation is not: Tie-breaking may give different goods to different bidders. However, the tie-breaking is inconsequential in the sense that the bidder utilities (and thus social welfare), item prices (and thus revenue), and which set of bidders are budget constrained, are all uniquely determined.

Given two BFPM, if the pacing multipliers of one dominate the other, then the revenue of that BFPM must also be at least as high. In the following, let $\text{Rev}(\alpha, x)$ refer to the revenue of a BFPM (α, x) .

Lemma 4. Given two BFPM $(\alpha^{(1)}, x^{(1)})$ and $(\alpha^{(2)}, x^{(2)})$, where $\alpha^{(1)} \geq \alpha^{(2)}$, $\text{Rev}(\alpha^{(1)}, x^{(1)}) \geq \text{Rev}(\alpha^{(2)}, x^{(2)})$.

Proof. Since $\alpha^{(1)} \geq \alpha^{(2)}$, prices under $(\alpha^{(1)}, x^{(1)})$ must be at least as large as under $(\alpha^{(2)}, x^{(2)})$. By the definition of BFPM, all demanded items must be sold completely. Therefore, under $(\alpha^{(1)}, x^{(1)})$ we sell at least all the items as we did under $(\alpha^{(2)}, x^{(2)})$ at prices that are at least those under $(\alpha^{(2)}, x^{(2)})$, hence the $\text{Rev}(\alpha^{(1)}, x^{(1)}) \geq \text{Rev}(\alpha^{(2)}, x^{(2)})$. \square

Corollary 2 (Revenue-maximizing). *The FPPE is revenue-maximizing among all BFPM.*

Theorem 1. *Given input (N, M, V, B) , an FPPE is guaranteed to exist. In addition, the uniquely-determined maximal pacing multipliers α maximize the revenue over all BFPM.*

Proof. Follows from Lemma 2 and Corollaries 1 and 2. \square

Properties of First-Price Pacing Equilibria

We first show that an FPPE is also a competitive equilibrium. In fact, we show that the concept of FPPE is equivalent to a natural refinement of competitive equilibrium.

Definition 3. A competitive equilibrium consists of prices p_j of goods and feasible allocations x_{ij} of goods to bidders such that the following properties hold:

1. Each bidder maximizes her utility under the given prices, that is, for every i it holds that $\mathbf{x}_i \in \arg \max_{\mathbf{x}_i \in [0,1]^m: \sum_j p_j x_{ij} \leq B_i} \{\sum_j (v_{ij} - p_j) x_{ij}\}$.
2. Every item with a positive price is sold completely, that is, for all j , $p_j > 0 \Rightarrow \sum_i x_{ij} = 1$.

We now introduce a refinement of competitive equilibrium that requires that each individual dollar that a bidder has is spent (or not spent) in a way that maximizes the utility obtained by the bidder for that dollar. Thus, there exists a rate β_i for each bidder that indicates her return on a dollar.

Definition 4. An equal-rates competitive equilibrium (ERCE) is a competitive equilibrium such that for every bidder i , there is a number β_i such that:

1. If $x_{ij} > 0$, then $v_{ij}/p_j = \beta_i$.
2. If i does not spend her entire budget, then $\beta_i = 1$.

We obtain the following characterization of FPPE.

Theorem 2. A combination of prices p_j and allocations x_{ij} is an ERCE if and only if it is an FPPE.

Proof. We first note that budget feasibility, the no-overselling condition, and the condition that items with a positive price must be sold completely, appear in the definitions of both concepts, so we only need to check the other conditions.

We first prove that an FPPE is also an ERCE. Let $\beta_i = 1/\alpha_i$. First, consider a bidder with $\alpha_i = 1$. If $x_{ij} > 0$ for some j , then $v_{ij}/p_j = v_{ij}/v_{ij} = 1 = \beta_i$, proving both conditions in the definition of an ERCE. Moreover, for any item j , we have $v_{ij}/p_j \leq v_{ij}/v_{ij} = 1$. Therefore, the bidder is spending optimally given the prices.

Next, consider a bidder with $\alpha_i < 1$. If $x_{ij} > 0$ for some j , then $v_{ij}/p_j = v_{ij}/(\alpha_i v_{ij}) = \beta_i$, proving the first condition in the definition of an ERCE. Moreover, by the definition of FPPE, such a bidder must spend her entire budget, proving the second condition. Moreover, for any item j , we have $v_{ij}/p_j \leq v_{ij}/(\alpha_i v_{ij}) = \beta_i$. Hence, the bidder is spending all her budget on the optimal items for her, and leaving money unspent would be suboptimal because $\beta_i > 1$. Therefore, the bidder is spending optimally given the prices. We conclude that an FPPE is also an ERCE.

We next prove that an ERCE is also an FPPE. For a bidder i with $\sum_j x_{ij} > 0$, consider the set of items $S_i = \{j : x_{ij} > 0\}$. By the ERCE property, we have that for $j \in S_i$, $v_{ij}/p_j = \beta_i$. Let $\alpha_i = 1/\beta_i$. We must have $\alpha_i \leq 1$, because otherwise i 's dollars would be better left unspent, contradicting the first property of competitive equilibrium. Also, if $\alpha_i < 1$ then all of i 's budget must be spent, establishing that no bidder is unnecessarily paced. For a bidder with $\sum_j x_{ij} = 0$, define $\alpha_i = \beta_i = 1$.

Now, we show that no part of an item j can be won by a bidder i for whom $\alpha_i v_{ij}$ is not maximal; if it were, by the first ERCE condition we would have $v_{ij}/p_j = \beta_i \Leftrightarrow \alpha_i v_{ij} = p_j$ and another bidder i' for whom $\alpha_{i'} v_{i'j} > \alpha_i v_{ij} = p_j$. Hence, $v_{i'j}/p_j > 1/\alpha_{i'} = \beta_{i'}$, but this would contradict that i' is receiving an optimal allocation under the prices.

Next, we prove that the prices are set correctly for an FPPE. For any item j , if it is sold completely, consider a bidder with $x_{ij} > 0$. Again, by the first ERCE condition we have $v_{ij}/p_j = \beta_i \Leftrightarrow \alpha_i v_{ij} = p_j$, and we have already established that this bidder must maximize $\alpha_i v_{ij}$. If the item is not entirely sold, then by the second condition of competitive equilibrium it must have price 0. This in turn implies that all bidders have value 0 for it, for otherwise there would be a bidder i with $\beta_i = v_{ij}/p_j = \infty$, who hence should be able to obtain a utility of ∞ since every one of the bidder's dollars must result in that amount of utility for her—but this is clearly impossible with finitely many resources. Thus, we have established all the conditions of an FPPE. \square

Next, we show that, unlike with second-price payments, under FPPE the seller does not benefit from adding fake bids.

Definition 5. A solution concept is shill-proof if the seller does not benefit from adding fake bids.

Proposition 1. FPPE are shill-proof.

Proof. Note that if we start from an FPPE and remove both a bidder and the items it wins, we still have an FPPE, since the remaining bidders are spending the same as before, and the remaining

items are allocated as before and thus fully allocated. Consider an instance of a market with three FPPE: (a) an FPPE with shill bids, (b) an FPPE of the market without shill bids, and (c) the FPPE generated by removing both the shill bids and the items they won from (a). Notice that the seller makes the same revenue in (a) and (c). Moreover, by Proposition 3 we know that the revenue of (b) is weakly greater than the revenue of (c), and therefore it is weakly greater than the revenue of (a). Thus, the seller cannot benefit from shill bids. \square

Akbarpour and Li (2018) observe that a first-price single-item auction is also *credible*, in the sense that the seller/auctioneer cannot benefit from lying about what other agents have done in the mechanism. However, FPPE do not satisfy this property.

Example 1. *Suppose $B_1 = 2, v_{11} = 2, v_{12} = 2$ and $v_{22} = 1$. The FPPE sets $p_1 = p_2 = 1$ and allocates both items to bidder 1. But the auctioneer could lie to bidder 1 claiming that someone else had bid 3 for item 2, and charge bidder 1 a price of 2 for item 1; and meanwhile charge bidder 2 a price of 1 for item 2, for a higher revenue overall.*

An FPPE does have a *price predictability* guarantee: given any allocation, a bidder either pays its full value or pays its budget. Even though individual item prices may not be known, this guarantees a degree of transparency to bidders about the price they will pay.

An FPPE is also robust to deviations by groups of bidders:

Definition 6. *An allocation with a set of payments is in the core if no group of bidders has an incentive to form a coalition with the seller to attain an outcome that is strictly better for all agents in the coalition.*

Proposition 2. *An FPPE is in the core.*

Proof. Since an FPPE is a competitive equilibrium, if we treat money as a good then we have a traditional locally non-satiated Walrasian equilibrium in an exchange economy. Since a locally non-satiated Walrasian Equilibrium is in the core, an FPPE is also in the core. \square

Monotonicity and Sensitivity Analysis

We showed that FPPEs are guaranteed to exist, that they are unique (up to inconsequential tiebreaking), and that they satisfy a number of attractive properties. We now look at how well-behaved FPPE is under changing conditions. We would ideally like the solution concept to be stable, so that changes in input cannot lead to disproportionate changes in output. We will show that this is largely the case. This is in contrast to SPPE, where Conitzer et al. (2018) showed that the equilibrium can be very sensitive: firstly, SPPE is not unique, and the revenue and welfare can change drastically across equilibria. Secondly, even when there is a unique equilibrium, small budget changes may lead to large revenue changes.

Monotonicity

We investigate whether FPPE is monotonic when adding bidders or goods, or when increasing budgets or valuations. Table 1 summarizes our results.

Revenue monotonicity is maintained for adding bidders, goods, and budget, but not incremental additions to valuations. Our proofs of revenue monotonicity all rely on Corollary 2: the fact that multipliers in an FPPE are maximal among BFPMS. Bidder and budget monotonicity both follow from a particularly simple argument: the original solution remains a BFPM, and thus by maximality

| | | | | |
|------|-------------|----------|--------------|----------------------|
| | Add Bidder | Add Good | Incr. Budget | Incr. Value v_{ij} |
| Rev. | ≥ 0 | ≥ 0 | ≥ 0 | Can go down |
| SW | Can go down | ≥ 0 | Can go down | Can go down |

Table 1: Overview of monotonicity results.

| | | |
|-----------------|--------------------------------------|--------------|
| | Decrease UB | Increase UB |
| Rev. (additive) | 0 | Δ |
| SW (relative) | $\frac{1-\Delta-\Delta^2}{1+\Delta}$ | $1 + \Delta$ |

Table 2: Overview of sensitivity results. For revenue, the number is the upper bound on change in revenue as a result of increasing a bidder’s budget by Δ , i.e. $B'_i = B_i + \Delta$. For social welfare, the number is the upper bound on relative change in social welfare as a result of a relative increase in budget of $1 + \Delta$, i.e. $B'_i = (1 + \Delta) \cdot B_i$.

of FPPE over BFPMS monotonicity is maintained. Explicit proofs and statements are given in the appendix. We show revenue monotonicity for goods here as an example.

Proposition 3. *In FPPE, adding a good weakly increases revenue.*

Proof. Let (α, x) be the FPPE for N, M , and let (α', x') be the FPPE for $N, M \cup \{j\}$ which includes the new good $j \notin M$. We first prove that $\alpha'_i \leq \alpha_i$ for all bidders $i \in N$: Suppose there are bidders whose multipliers go up (strictly); consider the set of all such bidders S . Collectively, these bidders are now winning weakly more goods (because there are more goods and nobody else’s (paced) bids went up). That means they are, collectively, paying strictly more (it’s first price and they’re bidding higher). But this is impossible, because all of them were running out of budget before (since they were paced).

Using the fact that $\alpha \geq \alpha'$, any bidder who was paced in α is still paced in α' and spending her whole budget. Let T be the set of buyers whose pacing multiplier hasn’t changed, i.e. $T = \{i \in N \mid \alpha_i = \alpha'_i\}$. They must win weakly more items: Any item they were tied originally with bidders outside T must now go completely to bidders in T . Additionally, bidders in T may win (part of) the new item. Since the pacing multipliers of bidders in T did not change, their prices did not change, hence winning weakly more items means they’re spending weakly more.

So bidders whose pacing multiplier changed are spending the same, and the remaining bidder spend weakly more. Hence revenue is weakly higher. \square

For social welfare, monotonicity is only maintained for goods. Adding bidders, or increasing budgets or valuations, can lead to drops in social welfare. The cause of nonmonotonicity is that there can be a mismatch between valuation and budget: a high-value but low-budget bidder can be lose out to a low-value high-budget bidder. Due to space constraints the proofs for social welfare appear in the appendix.

Sensitivity Analysis

We now investigate the sensitivity of FPPE to budget changes. An overview of our results is shown in Table 2. When adding Δ to the budget of a bidder, revenue can only increase, and at most by Δ . This shows that FPPE is, in a sense, revenue (and thus paced-welfare) stable with respect to budget changes: the change in revenue is at most the same as the change in budget. In contrast to this, Conitzer et al. (2018) show that in SPPE revenue can change drastically, at least by a factor of 100. The proofs are in the appendix.

For social welfare, due to the nature of multiplicative pacing additive bounds (such as the ones given for revenue) don't immediately make sense.³ Therefore we focus on sensitivity results for a *relative* change in budget, leading to a *relative* change in social welfare.

Our social welfare proofs rely on the fact that when a budget changes by factor $1 + \Delta$, pacing multipliers can only change by at most a factor $1 + \Delta$.

Lemma 5. *In FPPE, changing one bidder's budget $B'_i = (1 + \Delta)B_i$ for $\Delta \geq 0$, yields FPPE (α', x') with $\alpha \leq \alpha' \leq (1 + \Delta)\alpha$.*

Proof. Fix instance (N, M, V, B) with FPPE (α, x) , let $B'_i = (1 + \Delta)B_i$ and let (α', x') be the FPPE for (N, M, V, B') . Note that (α, x) is a BFPM for (N, M, V, B') , so by Corollary 1 $\alpha \leq \alpha'$. For the other inequality, note that $(\frac{\alpha'}{1+\Delta}, x')$, forms a BFPM for the original instance (N, M, V, B) : all prices drop by exactly a factor $\frac{1}{1+\Delta}$, which means that with the same allocation x' , spend for all bidders goes down by a factor $\frac{1}{1+\Delta}$ so no bidder exceeds their budget. By Corollary 1 the pacing multipliers α of the FPPE on (N, M, V, B') can only be higher, yielding $\alpha \geq \frac{\alpha'}{1+\Delta}$. Rearranging yields the claim. \square

To complete the proofs for social welfare, note that in an FPPE, pacing multipliers correspond to the bang-for-buck of buyers (i.e., the ratio between value and spend), so the bound in revenue change implies a bound in social welfare change. The proofs are given in the appendix.

Algorithms via Convex Programming

We now turn to computing the FPPE corresponding to an instance and adapt a well-known method for competitive equilibria. Solutions to the Eisenberg-Gale convex program for Fisher markets with quasi-linear utilities correspond exactly to FPPE in our setting. Cole et al. (2017) give the following primal convex program (as well as a dual not included here) for computing a solution to a Fisher market with quasi-linear utilities.

$$\begin{aligned}
 (\mathcal{CP}) \quad & \max_{x \geq 0, \delta \geq 0, u} \sum_i B_i \log(u_i) - \delta_i \\
 & u_i \leq \sum_j x_{ij} v_{ij} + \delta_i, \forall i \tag{1}
 \end{aligned}$$

$$\sum_i x_{ij} \leq 1, \forall j \tag{2}$$

The variables x_{ij} denote the amount of item j that bidder i wins, δ_i being nonzero denotes bidder i saving some of their budget, and u_i denotes a measure of utility for bidder i (exactly their utility when $\delta_i = 0$, otherwise it is not exact).

The dual variables β_i, p_j correspond to constraints (1) and (2), respectively. They can be interpreted as follows: β_i is the inverse bang-per-buck: $\min_{j: x_{ij} > 0} \frac{p_j}{v_{ij}}$ for buyer i , and p_j is the price of good j .

The leftover budget is denoted by δ_i , it arises from the dual program: it is the dual variable for the dual constraint $\beta_i \leq 1$, which constrains bidder i to paying at most a cost-per-utility rate of 1. See Cole et al. (2017) for the dual.

³To see why, take any instance (N, M, V, B) with budget-constrained bidders and compare it with an instance $(N, M, 2V, B)$ where the valuations are multiplied by 2. Changing a budget will yield the same allocation for both instances (and pacing multipliers are precisely a factor 2 off), but the change in social welfare will be twice as large in the second instance.

We now show via Fenchel duality that \mathcal{CP} computes an FPPE. Informally, the result follows because β_i specifies a single utility rate per bidder, duality guarantees that any item allocated to i has exactly rate β_i , and thus since \mathcal{CP} is known to compute a competitive equilibrium we have shown that it computes an ERCE. Theorem 2 then gives the result.

Theorem 3. *An optimal solution to \mathcal{CP} corresponds to a FPPE with pacing multiplier $\alpha_i = \beta_i$ and allocation x_{ij} , and vice versa.*

Proof. We start by listing the primal KKT conditions:

- | | |
|--|--|
| 1. $\frac{B_i}{u_i} = \beta_i \Leftrightarrow u_i = \frac{B_i}{\beta_i}$ | 4. $x_{ij}, \delta_i, \beta_i, p_j \geq 0$ |
| 2. $\beta_i \leq 1$ | 5. $p_j > 0 \Rightarrow \sum_i x_{ij} = 1$ |
| 3. $\beta_i \leq \frac{p_j}{v_{ij}}$ | 6. $\delta_i > 0 \Rightarrow \beta_i = 1$ |
| | 7. $x_{ij} > 0 \Rightarrow \beta_i = \frac{p_j}{v_{ij}}$ |

It is easy to see that x_{ij} is a valid allocation: \mathcal{CP} has the exact packing constraints. Budgets are also satisfied (here we may assume $u_i > 0$ since otherwise budgets are satisfied since the bidder wins no items): by KKT condition 1 and KKT condition 7 we have that for any item j that bidder i is allocated part of:

$$\frac{B_i}{u_i} = \frac{p_j}{v_{ij}} \Rightarrow \frac{B_i v_{ij} x_{ij}}{u_i} = p_j x_{ij}$$

If $\delta_i = 0$ then summing over all j gives

$$\sum_j p_j x_{ij} = B_i \frac{\sum_j v_{ij} x_{ij}}{u_i} = B_i$$

This part of the budget argument is exactly the same as for the standard Eisenberg-Gale proof (Nisan et al., 2007). Note that (1) always holds exactly since the objective is strictly increasing in u_i . Thus $\delta_i = 0$ denotes full budget expenditure. If $\delta_i > 0$ then KKT condition 6 implies that $u_i = B_i$ which implies $\delta_i = \frac{B_i}{u_i} \delta_i$. This gives:

$$\sum_j p_j x_{ij} + \delta_i = B_i \frac{\sum_j v_{ij} x_{ij}}{u_i} + \frac{B_i}{u_i} \delta_i = B_i$$

This shows that $\delta_i > 0$ denotes the leftover budget.

If bidder i is winning some of item j ($x_{ij} > 0$) then KKT condition 7 implies that the price on item j is $\alpha_i v_{ij}$, so bidder i is paying their bid as is necessary in a first-price auction. Bidder i is also guaranteed to be among the highest bids for item j : KKT conditions 7 and 3 guarantee $\alpha_i v_{ij} = p_j \geq \alpha_{i'} v_{i'j}$ for all i' .

Finally each bidder either spends their entire budget or is unpaced: KKT condition 6 says that if $\delta_i > 0$ (that is, some budget is leftover) then $\beta_i = \alpha_i = 1$, so the bidder is unpaced.

Now we show that any FPPE satisfies the KKT conditions for \mathcal{CP} . We set $\beta_i = \alpha_i$ and use the allocation x from the FPPE. We set $\delta_i = 0$ if $\alpha < 1$, otherwise we set it to $B_i - \sum_j x_{ij} v_{ij}$. We set u_i equal to the utility of each bidder. KKT condition 1 is satisfied since each bidder either gets a utility rate of 1 if they are unpaced and so $u_i = B_i$ or their utility rate is α_i so they spend their entire budget for utility B_i/α_i . KKT condition 2 is satisfied since $\alpha_i \in [0, 1]$. KKT condition 3 is satisfied since each item bidder i wins has price-per-utility $\alpha_i = \frac{p_j}{v_{ij}} = \beta_i$, and every other item has a higher price-per-utility. KKT conditions (4) and (5) are trivially satisfied by the definition of FPPE. KKT condition 6 is satisfied by our solution construction. KKT condition 7 is satisfied because a bidder i being allocated any amount of item j means that they have a winning bid, and their bid is equal to $v_{ij} \alpha_i$. \square

This shows that we can use \mathcal{CP} to compute an FPPE. Cole et al. (2017) show that \mathcal{CP} admits rational equilibria, and thus an FPPE can be computed in polynomial time with the ellipsoid method as long as all inputs are rational. Furthermore, the relatively simple structure of \mathcal{CP} means that it can easily be solved via standard convex-programming packages such as CVX (Grant and Boyd, 2008, 2014) or scalable first-order methods.

The equivalence between solutions to \mathcal{CP} and FPPE provides an alternative view on many of our results. Since Theorem 3 can be shown directly via Fenchel duality (as we do in the appendix), it allows us to prove via duality theory that FPPE corresponds to ERCE, and that FPPE always exists (\mathcal{CP} is easily seen to always be feasible and the feasible set is compact. Thus \mathcal{CP} always has a solution. By Theorem 3 that solution will be an FPPE.).

Experiments

We investigated the properties of FPPE via numerical simulations. We aimed to answer two questions: (1) Under FPPE, how high is bidder regret for reporting truthfully? (2) How does FPPE compare to SPPE in terms of revenue and social welfare?

To investigate these questions, we generated instances according to the complete-graph model by Conitzer et al. (2018). In this model, every bidder is interested in every item, and each valuation is drawn i.i.d. from $\text{unif}(0, 1)$. We generated 5 instances for each point in the Cartesian product of $2, 4, \dots, 8$ bidders and $4, 6, \dots, 14$ goods. For each instance, we computed an FPPE as the solution to \mathcal{CP} , computed with CVX (Grant and Boyd, 2014). We computed an SPPE for every objective function using the MIP given by Conitzer et al. (2018). We considered at most 8 bidders and 14 goods because of the limited scalability of the SPPE MIP; we were able to solve all \mathcal{CP} instances in less than 2ms. Figure 4 in the appendix shows the distribution of pacing multipliers; almost all bidders are budget constrained.

First we look at the ex-post regret that each bidder has in FPPE as compared to being able to unilaterally deviate, while keeping the FPPE multipliers fixed for all other bidders. We consider deviation either by individually setting bids in each auction, or choosing their own pacing multiplier after the fact. The former corresponds to a setting where bidders can interact with individual auctions. The latter corresponds to online ad-auction markets, where bidders often only specify a value for a conversion, targeting criteria, and a budget, whereas the individual auction valuations are based on conversion value scaled by a conversion rate (that the bidder cannot set). Changing the value for a conversion or the budget will affect all paced bids proportionately. The regrets are computed under the assumption that the valuations are truthful; alternatively this can be thought as the regret that the proxy agent has for not adjusting bids or multipliers after the fact. The results are shown in Figure 1. The figure shows summary statistics over max relative ex-post regret, which is the fraction of the truthful-response value that the bidder improves by if they deviate. For each instance, for each bidder, we compute the optimal best response, subject to budget constraints (in practice, we compute this by computing the optimal best response for second-price auctions; the bidder could achieve this result by sufficiently shading each bid to equal the bid of the next-highest bidder for every auction it wants to win). The max is over bidders in the auction, and the statistic is across instances. The middle line on each box is the median; the lower and upper hinges show the first and third quartiles. The line extending from the box shows outliers within 1.5 times the inter-quartile range. When the bidders are able to individually set bids in auctions the max regret is sometimes quite high, although it goes down as market thickness increases (note that market thickness is likely much higher in real-world auctions as compared to the at-most-10-bidder instances used here). On the other hand, when bidders are only able to set multipliers, as is often the case in ad-auction markets, we see that regret goes down rapidly, and already with 6-bidder

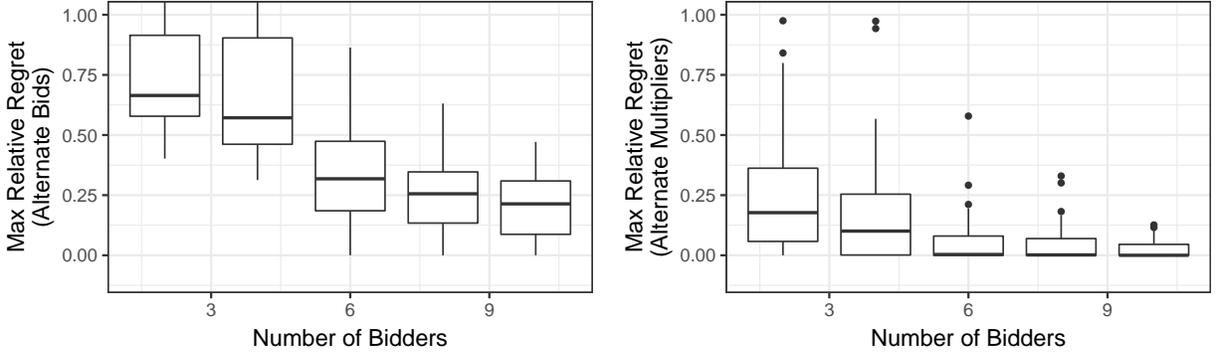


Figure 1: Summary statistics over max relative ex-post regret (max taken over bidders in a given auction, the statistic taken over the max across instances), when the bidder can choose an alternate bid for every auction (left) or a single alternative multiplier (right).

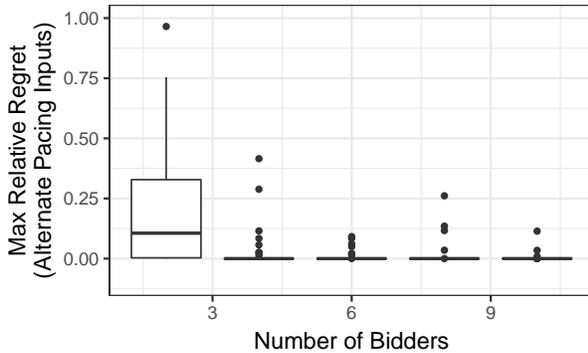


Figure 2: Summary statistics for when bidders can misreport their values or budget in order to change the FPPE.

instances the regret is usually near zero. Thus FPPE might indeed be an attractive solution concept for thick markets where bidders end up with (near) best responses.

We found that bidders have low ex-post regret as soon as markets have more than a few bidders. However another issue we may worry about is whether bidders can influence the FPPE outcome itself by misreporting their input to the FPPE computation. Figure 2 shows results for when bidders can misreport by scaling their valuation and budgets by multipliers $(\lambda_v, \lambda_b) \in [0.1, 0.2, \dots, 1.1]^2$ times their actual amounts. Already with 4 bidders there is practically never an incentive to do this: Only a few outlier instances have nonzero gain.

Secondly we compare revenue and social welfare under FPPE and SPPE. The results are shown in Figure 3. The left figure shows the cdf over the ratio of FPPE revenue and SPPE revenue, while the right figure shows the cdf over the ratio of FPPE welfare and SPPE welfare. We see that FPPE revenue is always higher than SPPE revenue, though it is the same for about 75% of instances, and almost never more than 1.5 times as high. For welfare we find that, perhaps surprisingly, neither solution concept is dominating, with most instances having relatively similar welfare in either solution concept, though FPPE does slightly better. There are two caveats to keep in mind for these results: firstly we did not compute the social-welfare-maximizing SPPE so it is possible that there is a better one (although this is highly unlikely given that Conitzer et al. (2018) find that there’s usually a single unique equilibrium in this class of instances); secondly almost every bidder is budget constrained in the FPPE of our setting, and so this might not generalize to settings where many bidders are not budget constrained (see the appendix for statistics on multipliers in the two solution concepts). These experiments show that FPPE are not necessarily worse than SPPE for welfare (at least with nonstrategic bidders), while potentially having significantly higher revenue.

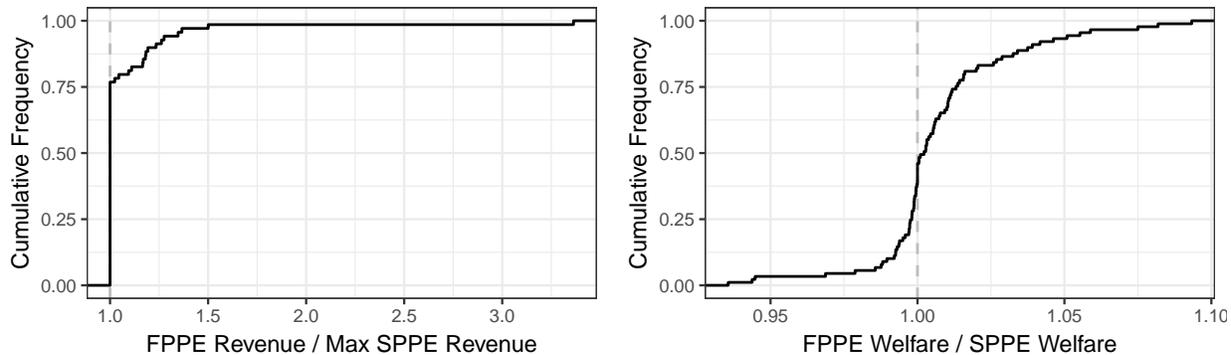


Figure 3: Summary statistics over the FPPE/ SPPE ratio of revenue (left) and social welfare (right).

Conclusion

In an ad platform, we must continually remember that the auction is only a small piece of a larger, complex system; when setting auction rules it is not just the properties of the auction that matter, but the platform’s consequent aggregate behavior. In this paper, we have seen that the equilibrium properties of an ad platform in which a first-price auction is used to sell each impression are in fact quite good, in many ways better than those when a second-price auction is used [Conitzer et al. \(2018\)](#).

In retrospect, the benefits of using a first-price auction are not surprising. In simple settings, second-price auctions win most theoretical head-to-head contests over first-price auctions; however, it is well-known that the luster of the second price auction fades as it is generalized to a Vickrey-Clarke-Groves (VCG) auction, so much that VCG earned the title “lovely but lonely” for its lack of use in practice [Ausubel and Milgrom \(2006\)](#). Indeed, some of the strengths of an ad platform based on first-price auctions are analogous to those seen in other complex auction settings [Milgrom \(2004\)](#) — uniqueness of equilibria, relation to competitive equilibria and core membership, shill-proofness and false-name-proofness, etc. — suggesting that first-price auctions may, in fact, have a serious role to play in today’s ad marketplaces.

References

- Abrams, Z.; Sathiya Keerthi, S.; Mendeleevitch, O.; and Tomlin, J. 2008. Ad delivery with budgeted advertisers: A comprehensive lp approach. 9.
- Akbarpour, M., and Li, S. 2018. Credible mechanisms. In *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, 371.
- Ausubel, L. M., and Milgrom, P. 2006. The lovely but lonely vickrey auction. In *Combinatorial Auctions, chapter 1*. MIT Press.
- Azar, Y.; Birnbaum, B.; Karlin, A. R.; and Nguyen, C. T. 2009. On revenue maximization in second-price ad auctions. In Fiat, A., and Sanders, P., eds., *Algorithms - ESA 2009*, 155–166. Berlin, Heidelberg: Springer Berlin Heidelberg.
- Balseiro, S., and Gur, Y. 2017. Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Available at SSRN 2921446*.
- Balseiro, S.; Kim, A.; Mahdian, M.; and Mirrokni, V. 2017. Budget management strategies in

- repeated auctions. In *Proceedings of the 26th International World Wide Web Conference, Perth, Australia*.
- Borgs, C.; Chayes, J.; Immorlica, N.; Jain, K.; Etesami, O.; and Mahdian, M. 2007. Dynamics of bid optimization in online advertisement auctions. In *Proceedings of the 16th international conference on World Wide Web*.
- Cary, M.; Das, A.; Edelman, B.; Giotis, I.; Heimerl, K.; Karlin, A. R.; Mathieu, C.; and Schwarz, M. 2007. Greedy bidding strategies for keyword auctions. In MacKie-Mason, J. K.; Parkes, D. C.; and Resnick, P., eds., *EC*, 262–271. ACM.
- Chen, Y. 2017. Programmatic advertising is preparing for the first-price auction era. <https://digiday.com/marketing/programmatic-advertising-readying-first-price-auction-era/>. Accessed: 2018-08-30.
- Cole, R.; Devanur, N.; Gkatzelis, V.; Jain, K.; Mai, T.; Vazirani, V. V.; and Yazdanbod, S. 2017. Convex program duality, fisher markets, and nash social welfare. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, 459–460. ACM.
- Conitzer, V.; Kroer, C.; Sodomka, E.; and Stier-Moses, N. E. 2018. Multiplicative pacing equilibria in auction markets. In *Conference on Web and Internet Economics (WINE'18)*.
- Dütting, P.; Fischer, F.; and Parkes, D. C. 2018. Expressiveness and robustness of first-price position auctions. *Mathematics of Operations Research*.
- Edelman, B., and Ostrovsky, M. 2007. Strategic bidder behavior in sponsored search auctions. *Decision Support Systems* 43(1):192–198.
- Feldman, J.; Muthukrishnan, S.; Pal, M.; and Stein, C. 2007. Budget optimization in search-based advertising auctions. In *Proceedings of the 8th ACM conference on Electronic commerce*.
- Goel, A.; Mahdian, M.; Nazerzadeh, H.; and Saberi, A. 2010. Advertisement allocation for generalized second-pricing schemes. *Oper. Res. Lett.* 38(6):571–576.
- Grant, M., and Boyd, S. 2008. Graph implementations for nonsmooth convex programs. In Blondel, V.; Boyd, S.; and Kimura, H., eds., *Recent Advances in Learning and Control*, Lecture Notes in Control and Information Sciences. Springer-Verlag Limited. 95–110.
- Grant, M., and Boyd, S. 2014. CVX: Matlab software for disciplined convex programming, version 2.1. <http://cvxr.com/cvx>.
- Hosanagar, K., and Cherepanov, V. 2008. Optimal bidding in stochastic budget constrained slot auctions. In *Proceedings 9th ACM Conference on Electronic Commerce (EC-2008), Chicago, IL, USA, June 8-12, 2008*, 20.
- Karande, C.; Mehta, A.; and Srikant, R. 2013. Optimizing budget constrained spend in search advertising. In *Proceedings of the Sixth ACM International Conference on Web Search and Data Mining, WSDM '13*, 697–706. New York, NY, USA: ACM.
- Mehta, A.; Saberi, A.; Vazirani, U.; and Vazirani, V. 2007. Adwords and generalized online matching. *Journal of the ACM (JACM)* 54(5).
- Mehta, A. 2013. Online matching and ad allocation. *Foundations and Trends in Theoretical Computer Science* 8(4):265–368.

- Milgrom, P. 2004. *Putting Auction Theory to Work*. Cambridge University Press.
- Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V. V. 2007. *Algorithmic game theory*. Cambridge university press.
- Rusmevichientong, P., and Williamson, D. P. 2006. An adaptive algorithm for selecting profitable keywords for search-based advertising services. In *Proceedings 7th ACM Conference on Electronic Commerce (EC-2006), Ann Arbor, Michigan, USA, June 11-15, 2006*, 260–269.
- Sluis, S. 2017. Big changes coming to auctions, as exchanges roll the dice on first-price. <https://adexchanger.com/platforms/big-changes-coming-auctions-exchanges-roll-dice-first-price/>. Accessed: 2018-08-30.

Omitted Monotonicity and Sensitivity Proofs

Monotonicity

Proposition 4. *In FPPE, adding a bidder weakly increases revenue.*

Proof. Let N be the original set of bidders, $i \notin N$ a new bidder, and M the set of goods. Let (α, x) be the FPPE on N and M . After adding bidder i , for each bidder $k \in N \setminus i$ and good $j \in M$, let $\alpha'_k = \alpha_k$ and $x'_{kj} = x_{kj}$. Set $\alpha'_i = x_{ij} = 0$ for bidder i and good $j \in M$, to obtain (α', x') . By construction (α', x') is a BFPM, so by Lemma 2, the revenue of the FPPE for $N \cup \{i\}$ and M must be at least as high. \square

Proposition 5. *In FPPE, increasing a bidder's budget from B_i to $B'_i > B_i$ weakly increases revenue.*

Proof. Let (α, x) be the FPPE where the budget of bidder i is still B_i . After increasing the budget to B'_i , (α, x) is still a BFPM. Therefore, by Lemma 2, the revenue of the new FPPE weakly increases. \square

Proposition 6. *In FPPE, increasing a bidder i 's value for some good j from v_{ij} to $v'_{ij} > v_{ij}$ can decrease revenue.*

Proof. Consider the following instance: 2 bidders, 2 goods, $v_{11} = 10, v_{12} = 5, v_{21} = 0, v_{22} = 5$, with $B_1 = 10, B_2 = 5$. The FPPE will be $\alpha_1 = \alpha_2 = 1, x_{11} = x_{22} = 1, x_{12} = x_{21} = 0$. Both bidders are spending their whole budget for total revenue 15.

However, consider $v'_{12} = 10 > v_{12}$. The FPPE is now $\alpha_1 = \frac{1}{2}, \alpha_2 = 1, x_{11} = x_{22} = 1, x_{12} = x_{21} = 0$. The bidders still receive the same goods, but the price for the first good dropped to 5 for total revenue of 10 instead of 15. \square

Social Welfare

Proposition 7. *In FPPE, adding a bidder can decrease social welfare by a factor $\frac{1}{2}$.*

Proof. Consider the following instance: 1 bidder, 1 good. We have $v_{11} = K$ for some parameter $K > 2$ and $B_1 = 1$. The FPPE is $\alpha_1 = 1/K, x_{11} = 1$ for social welfare K .

Now add bidder 2 with $v_{21} = 2, B_2 = 1$. The new FPPE is $a_1 = \frac{2}{K}, a_2 = 1$ and $x_{11} = x_{21} = \frac{1}{2}$. Social welfare now is $\frac{K}{2} + \frac{1}{2}$. As $K \rightarrow \infty$, the new social welfare is half of what it was before. \square

Proposition 8. *In FPPE, adding a good weakly increases social welfare.*

Proof. Fix N, M , and additional good $j \notin M$. Let (α, x) be the FPPE for N and M , (α', x') the FPPE for $N, M \cup \{j\}$ and let S be the set of bidders who are paced in (α, x) , i.e. $S = \{i \mid \alpha_i < 1\}$. We'll compare the contribution to social welfare of S and $N \setminus S$ separately.

First the set S . From the proof of Proposition 3, adding a good weakly decreases pacing multipliers. Since the bidders in S spent their entire budget in (α, x) , they must also spend their entire budget in (α', x') . The bang-per-buck of bidder i is $\frac{1}{\alpha_i}$, since by Definition 1 they pay $\alpha_i \cdot v_{ij}$ per unit of good j , and they receive v_{ij} of value per unit of good j . Since pacing multipliers weakly decreased, the bang-per-buck of bidders in S weakly increased, and as they spend their entire budget, their contribution to social welfare weakly increased.

Now the set $N \setminus S$. By Proposition 3, the total revenue weakly increased. Since the bidders in S spend exactly the same amount as before, the increase in revenue must have come from bidders in $N \setminus S$. Moreover, they were unpaced in (α, x) and so had a bang-for-buck of 1. In (α', x') , they have back-for-buck at least 1, hence their contribution to social welfare weakly increased.

Since the contribution to social welfare weakly increased for both sets S and $N \setminus S$, the total social welfare weakly increased. \square

Proposition 9. *In FPPE, increasing a bidder's budget from B_i to $B'_i > B_i$ can decrease social welfare.*

Proof. Consider the following instance (which is similar to the one in Proposition 7): 2 bidders, 1 good. We have values $v_{11} = K, v_{21} = 2$ and budgets $B_1 = B_2 = 1$. The FPPE is $a_1 = \frac{2}{K}, s_2 = 1, x_{11} = x_{21} = \frac{1}{2}$ for total social welfare of $\frac{K}{2} + \frac{1}{2}$.

Now increase bidder 2's budget to $B'_2 = 2$. The new FPPE is $a_1 = \frac{3}{K}, s_2 = 1, x_{11} = \frac{1}{3}, x_{21} = \frac{2}{3}$ for total social welfare of $\frac{K}{3} + \frac{2}{3}$. As $K \rightarrow \infty$, we lose $\frac{1}{6}$ of the social welfare. \square

Proposition 10. *In FPPE, increasing a bidder i 's value for some good j from v_{ij} to $v'_{ij} > v_{ij}$ can decrease social welfare.*

Proof. Consider the following instance: 2 bidders, 1 good. We have values $v_{11} = K, v_{21} = 1$, and budgets $B_1 = \frac{1}{2}, B_2 = 2$. The FPPE is $a_1 = \frac{1}{K}, a_2 = 1$ and $x_{11} = x_{21} = \frac{1}{2}$, for total social welfare $\frac{k+1}{2}$. Now increase $v'_{21} = 2$. The new FPPE will have $a_1 = \frac{2}{K}, a_2 = 1$, and $x_{11} = \frac{1}{4}, x_{21} = \frac{3}{4}$ for social welfare of $\frac{k+3}{4}$, as $K \rightarrow \infty$ we lose $\frac{1}{2}$ of the social welfare. \square

Sensitivity

Proposition 11. *In FPPE, increasing a bidder i 's budget by Δ , i.e. $B'_i = B_i + \Delta$, yields a revenue increase of at most Δ .*

Proof. Fix instance (N, M, V, B) , and let B' be the budget profile where $B'_i = B_i + \Delta$ for some bidder i . Let (α, x) be the FPPE on (N, M, V, B) , and let (α', x') be the FPPE on (N, M, V, B') . Since (α, x) is a BFPM for the new instance, we have $\alpha' \geq \alpha$, the new pacing multipliers are weakly higher than the old pacing multipliers. Let S^+ be the bidders for whom $\alpha'_k > \alpha_k$, and let $\text{Rev}_{S^+}^{\text{old}}$ be the revenue from them in (α, x) . Since the pacing multipliers for all bidders in S^+ strictly increased, they must have had $\alpha_k < 1$, so by the definition of a FPPE they must have spent their entire budget and $\text{Rev}_{S^+}^{\text{old}} = \sum_{k \in S^+} B_k$. In the new FPPE, they can't spend more than their budget, so $\text{Rev}_{S^+}^{\text{new}} \leq \sum_{k \in S^+} B'_k \leq (\sum_{k \in S^+} B_k) + \Delta = \text{Rev}_{S^+}^{\text{old}} + \Delta$.

What's left to show is that the revenue from the bidders S^- with $\alpha'_k = \alpha_k$ cannot have gone up. If there were any goods that S^+ and S^- were tied for, then after increasing the pacing multipliers of S^+ , the prices of those goods increased and S^+ won all of them. Moreover, the prices of goods that S^- as a set still win have not changed. Thus S^- is winning a subset of the goods they won previously at the same per-unit cost, hence their spend cannot have gone up. \square

Along with Proposition 5, this shows that when a bidder's budget increases by Δ , $\text{Rev}^{\text{new}} - \text{Rev}^{\text{old}} \in [0, \Delta]$. It's not difficult to see that these extremes can also both be attained: for the lower bound, increasing the budget of a non-budget-constrained bidder will not change the FPPE, hence revenue is unchanged. On the upper bound, take 1 bidder, 1 good, $v_{11} = 2\Delta, B_1 = \Delta$. Setting $B'_1 = B_1 + \Delta$ will increase revenue by Δ .

From Proposition 9 below, we know that social welfare can decrease when we increase a bidder's budget. The following lemma bounds that loss. In the following, let SW^{old} be the social welfare prior to changing the budget, and SW^{new} be the social welfare after changing the budget.

Proposition 12. *In FPPE, changing one bidder's budget $B'_i = (1 + \Delta)B_i$ for $\Delta \geq 0$, yields $SW^{\text{new}} \geq \left(\frac{1 - \Delta - \Delta^2}{1 + \Delta}\right) SW^{\text{old}}$.*

Proof. Let i be the bidder with $B'_i = (1 + \Delta)B_i$. Let (α, x) be the FPPE before the change, and let (α', x') be the FPPE after the budget change. Let S^p be the set of bidders who are paced in α' , and let $S_1 = N \setminus S_p$ the unpaced bidders. We'll lower-bound the new revenue from S_p and S_1 separately.

For the bidders in S_p , they spent their entire budget in both (α, x) and (α', x') : by Lemma 5 $\alpha \leq \alpha'$ hence bidders that are paced in α' are also paced in α and by the definition of FPPE that means they spend their entire budget. Moreover, by Lemma 5, their pacing multipliers cannot have gone up by more than $1 + \Delta$, hence their bang-per-buck is at least $\frac{1}{1+\Delta}$ that in (α, x) . Combining these statements, in (α', x') bidders in S_p spend at least as much as in (α, x) , and their bang-per-buck is at least $\frac{1}{1+\Delta}$ times that in (α, x) , hence their contribution to social welfare $SW_k^{\text{new}} \geq \frac{SW_k^{\text{old}}}{1+\Delta}$ for each $k \in S_p$, and therefore $SW_{S_p}^{\text{new}} \geq \frac{SW_{S_p}^{\text{old}}}{1+\Delta}$.

For the set S_1 of bidders who are unpaced in α' , their combined decrease in spend can be at most $\Delta \cdot B_i$: the total spend cannot have decreased by Prop 5, bidder i 's spend increased by at most $\Delta \cdot B_i$, the paced bidders (excluding bidder i) in α' were also all paced in α so their spend stayed constant, hence the largest possible reduction in spend by unpaced bidders in α' is $\Delta \cdot B_i \leq \Delta \cdot SW^{\text{old}}$. For unpaced bidders, spend equals contribution to social welfare, so we have $SW_{S_1}^{\text{new}} \geq SW_{S_1}^{\text{old}} - \Delta \cdot SW^{\text{old}}$.

Combining everything, we have $SW^{\text{new}} = SW_{S_p}^{\text{new}} + SW_{S_1}^{\text{new}} \geq \frac{SW_{S_p}^{\text{old}}}{1+\Delta} + SW_{S_1}^{\text{old}} - \Delta \cdot SW^{\text{old}} \geq \frac{SW^{\text{old}}}{1+\Delta} - \Delta \cdot SW^{\text{old}} = \left(\frac{1-\Delta-\Delta^2}{1+\Delta}\right) SW^{\text{old}}$. \square

Proposition 13. *In FPPE, changing one bidder's budget B_i to $(1+\Delta)B_i$ for $\Delta \geq 0$, yields $SW^{\text{new}} \leq (1+\Delta)SW^{\text{old}}$.*

Proof. Let i be the bidder whose budget increases from B_i to $(1+\Delta)B_i$. Let (α, x) be the FPPE before the change, and (α', x') be the FPPE after the budget change. By Lemma 5, increasing a budget can only increase pacing multipliers. Let S^+ be the set of bidders whose pacing multiplier increased (for convenience excluding bidder i), let S^- be the set who had pacing multiplier strictly lower than 1 whose multiplier did not change, and let S^1 be the set of bidders who were unpaced in (α, x) . Let SW^{old} be the old social welfare, and SW^{new} the new one. Define SW_i , SW_+ , SW_- , and SW_1 as the contribution to social welfare of bidder i , bidders S_+ , S_- , and S_1 respectively.

We use extensively that at pacing multiplier α_k , the spend $s_k = \alpha_k \cdot SW_k$.

For bidder i , we have $SW_i^{\text{new}} \leq (1+\Delta)SW_i^{\text{old}}$. The pacing multiplier of i can only have increased, so the bang-for-buck can only have decreased. Spend increased at most by $(1+\Delta)$, back-for-buck was at most the same, hence SW cannot be more than $1+\Delta$ more.

For bidders S_+ , we have $SW_+^{\text{new}} < SW_+^{\text{old}}$. Their spend cannot have increased as they spent their budget completely in (α, x) . Meanwhile, their bang-for-buck strictly decreased due to increasing pacing multipliers.

For bidders S_- , we have $SW_-^{\text{new}} = SW_-^{\text{old}}$. Since they were and are paced, they must spend their entire budget. Since their pacing multiplier hasn't changed, their bang-for-buck stayed the same. Thus their contribution to SW stayed the same.

Finally, for bidders S_1 , we have $SW_1^{\text{new}} \leq SW_1^{\text{old}}$. From the proof of prop. S4, the total spend of bidders $S_- + S_1$ cannot have increased. Since the spend of bidders in S_- must have stayed the same, the spend of bidders in S_1 cannot have increased. Since their bang-for-buck is 1, their SW cannot have increased.

Summing over all groups: $SW^{\text{new}} = SW_i^{\text{new}} + SW_+^{\text{new}} + SW_-^{\text{new}} + SW_1^{\text{new}} \leq (1+\Delta)SW_i^{\text{old}} + SW_+^{\text{old}} + SW_-^{\text{old}} + SW_1^{\text{old}} \leq (1+\Delta)SW^{\text{old}}$. \square

Additional experimental results

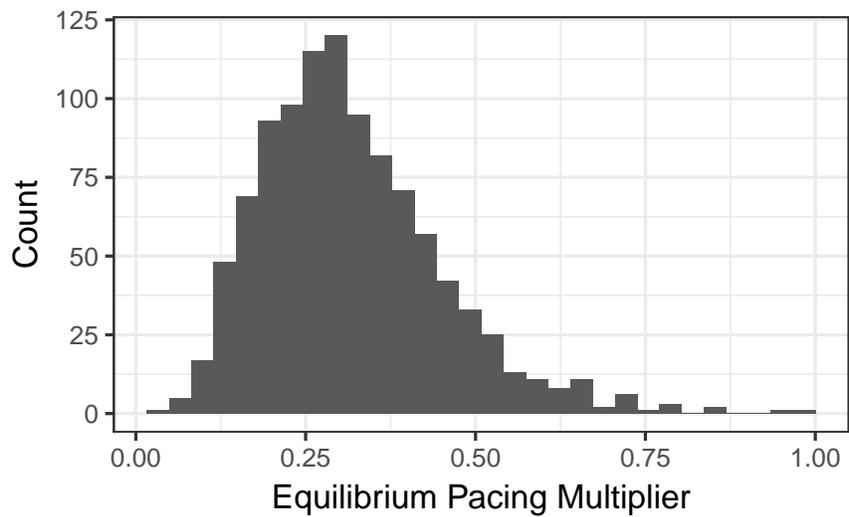


Figure 4: A histogram of pacing multipliers across all bidders and all instances. As can be seen almost all bidders are budget constrained.

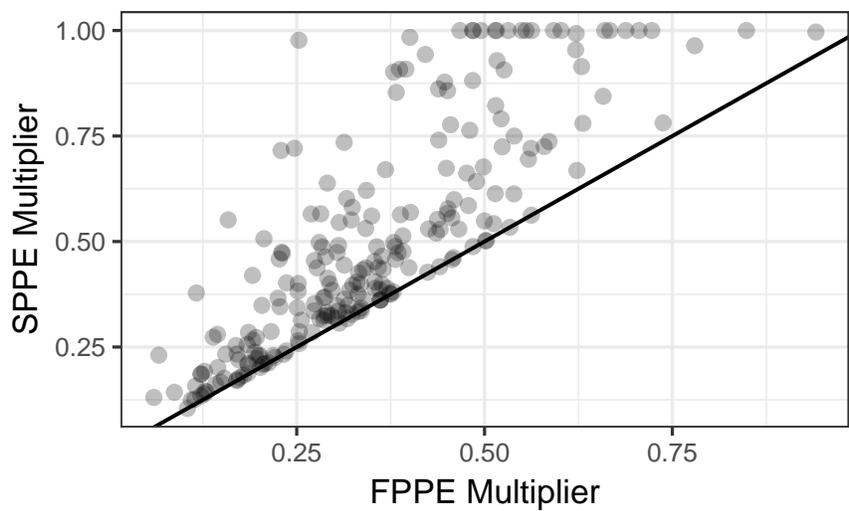


Figure 5: A plot of pacing multipliers across FPPE and SPPE solutions. Each dot is a bidder in an instance. The x-axis shows the FPPE multiplier. The y-axis shows the SPPE multiplier. The multipliers are higher for SPPE in every instance.

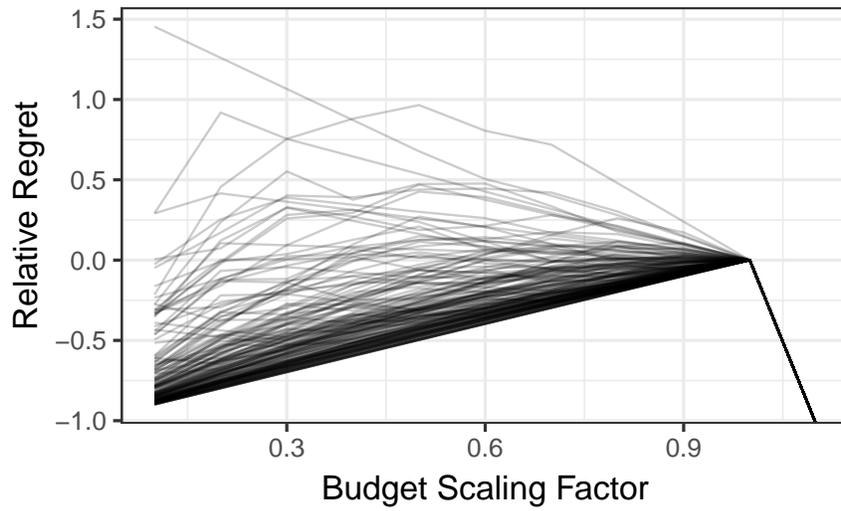


Figure 6: Regret faced by a bidder that truthfully reports bids and budgets into a first-price pacing system, as opposed to truthfully reporting bids but *misreporting budget* by a multiplicative factor. Each line corresponds to a bidder in a particular pacing instance.

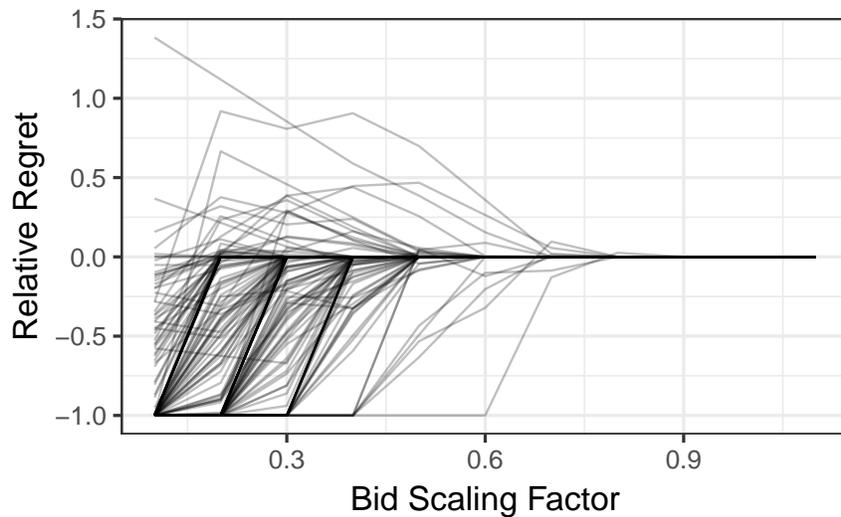


Figure 7: Regret faced by a bidder that truthfully reports bids and budgets into a first-price pacing system, as opposed to truthfully reporting budgets but *misreporting bids* by a multiplicative factor. Each line corresponds to a bidder in a particular pacing instance.

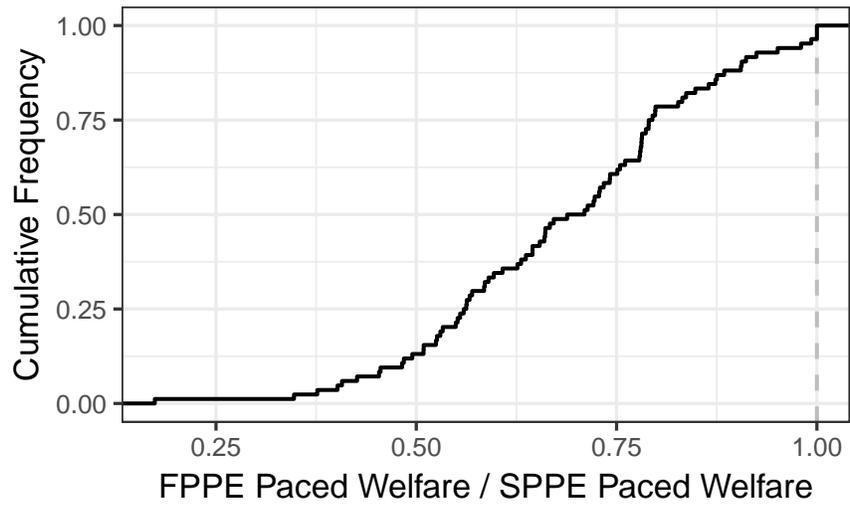


Figure 8: Summary statistics over the FPPE / SPPE ratio of paced welfare, where paced welfare is defined as welfare scaled down by the bidders equilibrium pacing multiplier.