



Management Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Pacing Equilibrium in First Price Auction Markets

Vincent Conitzer, Christian Kroer, Debmalya Panigrahi, Okke Schrijvers, Nicolas E. Stier-Moses, Eric Sodomka, Christopher A. Wilkens

To cite this article:

Vincent Conitzer, Christian Kroer, Debmalya Panigrahi, Okke Schrijvers, Nicolas E. Stier-Moses, Eric Sodomka, Christopher A. Wilkens (2022) Pacing Equilibrium in First Price Auction Markets. *Management Science* 68(12):8515-8535. <https://doi.org/10.1287/mnsc.2022.4310>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2022, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Pacing Equilibrium in First Price Auction Markets

Vincent Conitzer,^{a,b} Christian Kroer,^{c,d,*} Debmalya Panigrahi,^b Okke Schrijvers,^c Nicolas E. Stier-Moses,^c Eric Sodomka,^c Christopher A. Wilkens^e

^a Econorithms LLC, Chapel Hill, North Carolina 27517; ^b Computer Science Department, Duke University, Durham, North Carolina 27708;

^c Core Data Science, Meta, Menlo Park, California 94025; ^d Industrial Engineering and Operations Research Department, Columbia University, New York, New York 10027; ^e Tremor Technologies, Boston, Massachusetts 02110

*Corresponding author

Contact: conitzer@cs.duke.edu, <https://orcid.org/0000-0003-1899-7884> (VC); christian.kroer@columbia.edu, <https://orcid.org/0000-0002-9009-8683> (CK); debmalya@cs.duke.edu (DP); okke@fb.com (OS); nicostier@yahoo.com, <https://orcid.org/0000-0001-6380-1516> (NES-M); sodomka@fb.com (ES); c.a.wilkens@gmail.com (CAW)

Received: July 31, 2019

Revised: November 6, 2020; July 19, 2021

Accepted: August 19, 2021

Published Online in Articles in Advance: April 6, 2022

<https://doi.org/10.1287/mnsc.2022.4310>

Copyright: © 2022 INFORMS

Abstract. Mature internet advertising platforms offer high-level campaign management tools to help advertisers run their campaigns, often abstracting away the intricacies of how each ad is placed and focusing on aggregate metrics of interest to advertisers. On such platforms, advertisers often participate in auctions through a proxy bidder, so the standard incentive analyses that are common in the literature do not apply directly. In this paper, we take the perspective of a budget management system that surfaces aggregated incentives—instead of individual auctions—and compare first and second price auctions. We show that theory offers surprising endorsement for using a first price auction to sell individual impressions. In particular, first price auctions guarantee uniqueness of the steady-state equilibrium of the budget management system, monotonicity, and other desirable properties, as well as efficient computation through the solution to the well-studied Eisenberg–Gale convex program. Contrary to what one can expect from first price auctions, we show that incentives issues are not a barrier that undermines the system. Using realistic instances generated from data collected at real-world auction platforms, we show that bidders have small regret with respect to their optimal *ex post* strategy, and they do not have a big incentive to misreport when they can influence equilibria directly by giving inputs strategically. Finally, budget-constrained bidders, who have significant prevalence in real-world platforms, tend to have smaller regrets. Our computations indicate that bidder budgets, pacing multipliers, and regrets all have a positive association in statistical terms.

History: Accepted by Gabriel Weintraub, revenue management and market analytics.

Funding: D. Panigrahi was supported in part by the National Science Foundation [Awards CCF 1535972, CCF 1750140, and CCF 1955703].

Supplemental Material: The data files are available at <https://doi.org/10.1287/mnsc.2022.4310>.

Keywords: mechanism design • auction systems • pacing equilibria

1. Introduction

The early days of internet advertising were the Wild West. Savvy search advertisers at that time would carefully tune their keywords and bids, building complex tools that precipitated cyclic bidding patterns (Edelman and Ostrovsky 2007). For example, it would pay off to differentiate between users who search for “flower” and “flowers” in search advertising. Deft display advertisers would exploit real-time bidding (RTB) capabilities—bidding programmatically on a per-impression basis—to identify high-value users where competitors lacked data, reaching them at bargain basement prices. Crafty clickers would sabotage competitors’ campaigns to gain advantage. Such complex strategic behavior was enabled by low-level bidding tools and driven by advertisers’ focus on the value of each customer.

As platforms matured and developed higher-level functionality, programmatic advertising attracted advertisers

who cared more about the population of users they reached than about the value of an individual user. Early programmatic systems like search advertising and RTB bidding, where advertisers more or less bid on each opportunity, catered to advertisers targeting events deep in the sales funnel where the value of each event was substantial and measurable. Even as these systems came to dominate search and remnant display advertising¹ in the first decade of the twenty-first century, advertisers simply looking to reach an audience would often go through salespeople to buy so-called guaranteed placements. With only a rough measure of the value in expectation of reaching each individual, competitive bidding on a per-user basis made little sense for these advertisers. The gap between these worlds was bridged by platforms that abstracted away the auctions to focus on reach. Advertisers would input basic data to set up a

campaign, and platforms helped them run the campaign relying on a proxy bidder. For example, an advertiser at a social media platform today might simply say, "I want to show my ad to as many junior tech professionals in Colorado as possible for \$1,000" but not indicate the value per individual.

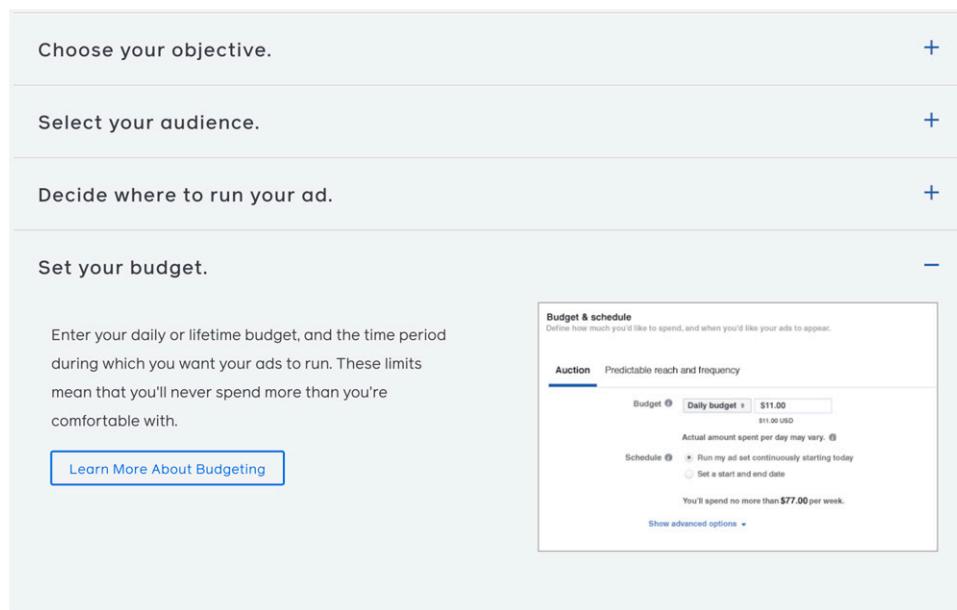
Facebook's current advertiser workflow exemplifies this mindset as illustrated in Figure 1. The campaign creation flow asks an advertiser four questions. (1) What do you want people to do (see, watch, click, buy, etc.)? (2) Who do you want to reach (targeting criteria)? (3) Where do you want to show your ad (e.g., Facebook, Instagram, Messenger)? (4) How much money do you want to spend (budget)? Notably missing in the basic flow is the willingness to pay to reach an individual person (bid).²

This shift from the primacy of the individual user to the primacy of the population being reached substantially changes the mechanism designer's problem. In the standard theory, the bid is assumed to be the primary strategic implement, and budgets, when present, are often assumed to be public knowledge to make analysis tractable. However, when focusing on the population reached, advertisers' primary strategic levers are budgets and audience targeting, with control of the bid ceded to a proxy bidder. This turns much of the standard auction literature on its head. The advertiser's bid, which was already the product of a constant per-event bid and a platform-generated impression-specific event probability, is now entirely managed by a proxy bidder tasked with managing the advertiser's budget across impressions. This leads us to the motivation of the present paper; in a system where

advertisers' primary or only strategic lever is the budget, what is the best way to allocate and price ads?

The goal of our study is to understand the strengths and weaknesses of different auction rules in the context of proxy bidders that help run campaigns. We study a specific question on this topic. When the proxy bidder is designed to shade bids so the advertiser's budget is exhausted at the end of the budget horizon, would the system as a whole perform better if individual impressions are sold through a first price or a second price mechanism? Because the proxy bidders shade bids on behalf of advertisers to maximize their utilities over the course of the campaign, the standard analyses of first and second price auctions no longer apply. Instead, the focus is the study of equilibria resulting from the independent and strategic behavior of all proxy bidders. First, we show that when a first price auction is used to sell each impression, the resulting *first price pacing equilibrium* (FPPE) has many theoretical guarantees that would not hold in an analogous *second price pacing equilibrium* (SPPE), including uniqueness and monotonicity. Second, we show that FPPEs are not an arbitrary construction; they correspond to market equilibria, which can be efficiently computed using the well-studied Eisenberg–Gale (EG) convex program (unlike SPPEs, which are PPAD complete to compute (Chen et al. 2021)). Finally, we run simulations on real data collected from Facebook and Instagram advertising auctions in order to study the impact on incentives for advertisers. In these simulations, we find that the ex post incentive for an individual advertiser to shade their bids via reporting a lower value per conversion or budget is very small in almost all

Figure 1. (Color online) Flow to Create a Facebook Ad as Described at <https://www.facebook.com/business/ads>



cases. In an individual first price auction, the winner has a significant incentive to bid strategically when the gap between the first and second bid is large, and closing the gap requires a “thick” market where many bidders are present. However, in a pacing equilibrium, the proxy bidder is picking a single bid to be used in all auctions, so it is enough that the gap between the first and second bids is small in a fraction of the auctions in which the proxy bidder will participate. We also find that the incentive for misreporting the value per conversion or budget as inputs to the mechanism is vanishingly small. We hypothesize that this conclusion has to do with the coarseness at which manipulations can be performed when buyers do not have the ability to shade bids in individual auctions.

1.1. Pacing Equilibria in Internet Advertising

The primary function of proxy bidders in internet advertising is to optimize the way advertisers’ budgets are spent. Most modern implementations are based on control systems that tune one or more control parameters and apply them to the incoming opportunities. A common budget management method known as bidder selection or *throttling* tries to enforce budget constraints by adaptively selecting which advertisers participate in each auction. This is commonly done by tuning a probability for each ad to randomize the participation in the auctions it targets. The platform adjusts the ad’s participation probability continuously; when an ad keeps spending its budget too fast, its participation probability is continuously reduced and vice versa. Other platforms instead modify the bid of each ad by applying a shading factor, referred to as a *pacing multiplier*, to manage the budget. Tuning the pacing multiplier changes the spending rate, thereby attempting to exhaust the ad’s budget precisely at the end of the budget horizon. (For further details, we expand the description of pacing in the next section.) Between these two options, shading-based systems tend to increase advertisers’ utilities by allocating the cheapest impressions to them, whereas throttling offers advertisers a more representative sample of the available opportunities. There is no one-size-fits-all solution, and the design space of budget management systems is large and rife with trade-offs.

Both methods have precedent. Throttling has been studied 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 in 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 multiobjective optimization in search engines (Abrams et al. 2008, Azar et al. 2009, Goel et al. 2010, Karande et al. 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).

Our work 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 between zero and one. Many advertising 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 return on investment (Rusmevichientong and Williamson 2006, Feldman et al. 2007, Hosanagar and Cherepanov 2008). In addition to advertiser strategies, Cary et al. (2007) also studies the limit point if all advertisers participate in a repeated position auction. In that setting, prices converge to Vickrey–Clarke–Groves (VCG) prices, but there is no heterogeneity in impression opportunities. Closest to our work is that of Borgs et al. (2007), who study first price auctions with budget constraints in a perturbed model. They show that in the limit of their perturbations, prices converge to those of an (equal rates) competitive equilibrium. The limit point they describe is an **FPPE**; hence, this shows guaranteed existence and the relation to competitive equilibria. It leaves open the question of uniqueness, revenue, and other properties of **FPPE**. They also point out a convex set of constraints that can be used to compute this limit point and thus, an **FPPE**.

Recent work has studied pacing equilibria in the context of second price auctions; Balseiro et al. (2015) investigate budget management in auctions through a fluid mean-field approximation, which leads to elegant existence results and closed form descriptions of equilibria in certain settings. Balseiro et al. (2021) studies several budget smoothing methods, including multiplicative pacing in a stochastic context where ties do not occur. Balseiro and Gur (2019) studies how an individual bidder might adapt their pacing multiplier over time in a stochastic continuous setting. Conitzer et al. (2021) studies second price pacing when bidders, goods, budgets, and valuations are known and proves that equilibria exist under fractional allocations. In this paper, we study pacing in the context of first price auctions.

Although 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. In 2017, exchanges such as AppNexus, Index Exchange, and OpenX began moving their inventory to from second to first price auctions (Chen 2017, Sluis 2017). Google transitioned all display advertising publisher inventory to first price auctions in 2019, citing a desire for greater transparency and simplicity in the ad market; they reported a neutral to positive impact on revenue for publishers and a more competitive market (Bigler 2019). Mopub (Twitter's ad platform) transitioned all publisher inventory to a first price auction from June to October of 2020, also citing a desire for increased transparency, reduced complexity, and increased fairness (Jordan 2020). To explain this recent shift to first price auctions, Paes Leme et al. (2020) analyzed the equilibria of a game played between auction markets, where each market can choose its auction rules. We focus on a game between proxy bidders as opposed to a game between auction markets. We remark that this shift to first price auctions has happened in display advertising, which typically uses real-time bidding as opposed to a centralized pacing system (although intermediaries likely have their own budget pacing systems for their advertisers). By contrast, in this work, we explore the properties that arise when first price auctions are combined with centralized pacing systems.

In the context of position auctions, Dütting et al. (2019) further showed that first price does not suffer from the same equilibrium selection problems as GSP and VCG. 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. 2021). Indeed, as stated earlier, one of the most developed lines of research 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.³

1.2. Contributions

We introduce the **FPPE** problem, where each buyer participates in many first price auctions with budget constraints that span the auctions. The platform seeks to find a vector of *pacing multipliers*, one for each buyer, and buyers bid their value times their pacing multiplier. At a high level, the platform seeks to find multipliers such that buyers are only just paced enough that they do not overspend their budget. We allow goods to be allocated fractionally if there are ties; this is well motivated in the ad auction setting, where a good may be interpreted as representing thousands of impressions. We show that **FPPE** can be understood as a maximal point within the larger class of *budget-feasible pacing multipliers* (**BFPMs**), which is the set of vectors of pacing multipliers that satisfy all

budgets but potentially pace beyond what is necessary in order to satisfy budgets.

1.2.1. Existence and Uniqueness. From Borgs et al. (2007), it follows that **FPPEs** are guaranteed to exist. Our first result shows that not only does an **FPPE** always exist but also, that it is *essentially unique*.⁴ In fact, we show that the **FPPE** exactly coincides with the (unique) *maximal* set of pacing multipliers. This also leads to the observation that the **FPPE** is revenue maximizing among all **BFPMs**. Our structural characterization of **FPPE** in terms of **BFPMs** is a powerful tool for reasoning about properties of **FPPE** and is the basis of several of our results and additional ones by Peysakhovich and Kroer (2019) in follow-up work on competitive equilibria from equal incomes. In contrast, **SPPEs** are guaranteed to exist but are not unique in general (Conitzer et al. 2021). Furthermore, we also show that an **FPPE** yields a *competitive equilibrium* (i.e., every bidder is exactly allocated her *demand set*). **SPPE** also corresponds to a competitive equilibrium but only for buyers that are *supply aware* when computing their demands, whereas **FPPE** is a competitive equilibrium both for supply-aware and supply-unaware buyers.

1.2.2. Computability. We show an interesting connection between **FPPE** and a generalization of the classical **EG** convex program for quasilinear utilities (Cole et al. 2017). Using Fenchel duality, we use the **EG** program to infer that the unique maximal **BFPMs**, 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. 2021), and very recent results show that the general problem is PPAD complete (Chen et al. 2021). Although Borgs et al. (2007) already showed a convex program, this connection to **EG** is important because it lends itself to scalable first-order methods (Nesterov and Shikhman 2018, Kroer et al. 2021).

1.2.3. Monotonicity and Sensitivity. We show that **FPPE** satisfies many notions of monotonicity in terms of both 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. Again, this stands in sharp contrast with **SPPEs**, which generally do not satisfy such monotonicity conditions.

In fact, in an **FPPE**, not only is the revenue monotonically nondecreasing with budgets, but it also changes smoothly in the sense that increasing the budget by a constant Δ can only increase the revenue by Δ . This does not hold for **SPPE**, where the revenue

can increase by a substantial amount even for a small change in the budgets.

1.2.4. Shill Proofness. Using monotonicity, we also establish that there is no incentive for the platform to enter fake bids to an FPPE mechanism.

1.2.5. Simulations. To test the properties mentioned earlier and compare equilibria computationally, we rely on the EG convex program formulation to perform simulations on realistic instances constructed from elements of real-world data. We find that the ex ante and ex post regret associated with FPPE is very small when bidders can only change the scale of their utilities, as is usually the case in ad markets when bidding on clicks or conversions. Finally, we compare the revenue and social welfare of FPPE and SPPE across instances. We find that FPPE always provide higher revenue, whereas welfare splits about evenly on which solution concept performs better. The contributions and the comparison between FPPE and SPPE are summarized in Table 1.

1.2.6. Conclusion. In Section 7, we discuss the assumptions we have made and put them in perspective versus more general or alternative assumptions.

2. First Price Pacing Equilibria

We consider a single-slot auction market in which a set of bidders $N = \{1, \dots, n\}$ targets 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 (although our results on equivalence to competitive equilibrium show that the fractional choices made by the auctioneer 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 BFPMS is a tuple (α, x) of pacing multipliers $\alpha_i \in [0, 1]$ for each bidder $i \in N$ and fractional allocations $x_{ij} \in [0, 1]$ for each bidder $i \in N$ and good $j \in M$ that satisfies the following properties.

- (Prices) Unit price $p_j = \max_{i \in N} \alpha_i v_{ij}$ for each good $j \in M$.
- (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) $\sum_{j \in M} x_{ij} p_j \leq B_i$ for each bidder $i \in N$.
- (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 are particularly interested in outcomes that are stable in some sense. Specifically, we are interested in solutions where no bidder is unnecessarily paced, which we call FPPE. This captures that the platform does not want to pace an advertiser that is not spending the full budget.

Definition 2. An FPPE is a BFPMS tuple (α, x) of pacing multipliers α_i for each bidder i and fractional allocation x_{ij} for bidder i and good j with this additional property.

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

Table 1. A Comparison of FPPE and SPPE

	SPPE	FPPE
Exists?	Yes	Yes
Buyers best responding?	Yes	No
Is market equilibrium?	Yes	Yes (even for supply-unaware buyers)
Is unique?	No	Yes, in utilities, multipliers, and prices
Is efficiently computable?	PPAD complete	Convex program
Is welfare monotone?	No	Yes, in goods
Is revenue monotone?	No	Yes, in goods/bidders/budgets
Is shill proof?	No	Yes
Simulated regret/IC	No regret, very small IC	Small regret, small IC
Simulated revenue	$SPPE \leq FPPE$	
Simulated welfare	Ambiguous	

Notes. In the second row, we note that under SPPE, each bidder's multiplier is a mutual best response given that other agent bids are fixed—and would be even if the set of deviating strategies was expanded to allow the bidder to submit arbitrary per-auction bids (for a proof, see proposition 1 in Conitzer et al. 2021). Under FPPE, a buyer may increase its utility by shading its multiplier (or shading individual bids if it could submit arbitrary bids for each auction). IC stands for incentive compatibility violation.

2.1. Existence, Uniqueness, and Structure of FPPE

Borgs et al. (2007) show that FPPEs are guaranteed to exist. We show that not only are FPPEs guaranteed to exist but that they are also unique and maximize the seller's revenue over all BFPMS. In the following, inequalities should be interpreted *component wise* when applied to vectors or sets.

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

Proof. First, we will show that given any two BFPMS $(\alpha^{(1)}, x^{(1)})$ and $(\alpha^{(2)}, x^{(2)})$, there exists a BFPMS 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 BFPMSs. Let the corresponding allocation x^* be that for each good j , identify which of the two BFPMSs had the highest paced bid for j , breaking ties toward the first BFPMS; then, allocate the good to the same bidder as in that BFPMS (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 toward 1). To show budget feasibility, we now prove the following: if a bidder i wins fraction x of an item j in (α^*, x^*) , then she must also win at least fraction x of item j in $(\alpha^{(b)}, x^{(b)})$. The key observation is that for any bidder and item, her paced bid in (α^*, x^*) is at least as much as her paced bid in either $(\alpha^{(1)}, x^{(1)})$ or $(\alpha^{(2)}, x^{(2)})$ —we call this the *monotone* property. Now, consider two situations. First, if i did not have the highest bid on an item j in $(\alpha^{(b)}, x^{(b)})$, then the bidder who had a higher bid in $(\alpha^{(b)}, x^{(b)})$ continues to have a higher bid in (α^*, x^*) by the monotone property. Hence, i does not win this item under (α^*, x^*) . Second, suppose that i did have the highest bid on an item j in $(\alpha^{(b)}, x^{(b)})$ and won a fraction x of the item. In this case, every other bidder who matched i 's bid in $(\alpha^{(b)}, x^{(b)})$ continues to have at least as high a bid in (α^*, x^*) by the monotone property. Therefore, i shares item j in (α^*, x^*) with at least the set of bidders she shared it with in $(\alpha^{(b)}, x^{(b)})$. There are three further subcases. First, if the highest-paced bid for j is unique, then the highest must be achieved by $(\alpha^{(b)}, x^{(b)})$; in this case, i gets exactly fraction x of item j . Second, if the highest-paced bid for j is tied and $b = 1$, then item j is divided exactly as in $(\alpha^{(1)}, x^{(1)})$. Again, i gets exactly

fraction x of item j . Third, if the highest-paced bid for j is tied but $b = 2$, then we claim that i does not get item j in (α^*, x^*) . This is because the allocation of j under (α^*, x^*) is identical to that under $(\alpha^{(1)}, x^{(1)})$, but because $b = 2$, we have $\alpha_i^{(1)} < \alpha_i^{(2)}$, which implies that i does not have the highest bid for j in $(\alpha^{(1)}, x^{(1)})$. In summary, in any of these three cases, the fraction of item j that bidder i wins under (α^*, x^*) is at most the fraction she wins under $(\alpha^{(b)}, x^{(b)})$. As a consequence, i spends no more under the new BFPMS than under BFPMS b , which is budget feasible. Hence, the new BFPMS is budget feasible.

We now complete the proof. Let $\alpha_i^* = \sup \{\alpha_i \mid \alpha$ is part of a BFPMS}. We will show that α^* is part of a BFPMS, proving the result. For any $\epsilon > 0$ and any i , there exists a BFPMS where $\alpha_i > \alpha_i^* - \epsilon$. By repeatedly taking the component-wise maximum for different pairs of i , we conclude there is a single BFPMS $(\alpha^\epsilon, x^\epsilon)$ 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 BFPMS by continuity. \square

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

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

Proof. Suppose bidder i is unnecessarily paced under the maximal BFPMS (α, x) . If i is not tied for any good, then we can increase her pacing multiplier by a sufficiently small $\epsilon > 0$ such that i is still not tied for any item and is within budget, contradicting the fact that α was maximal. So, bidder i must be tied for at least one good. Define $N(i)$ as all the bidders that are tied for any good with bidder i (i.e., $N(i) = \{\text{bidder } k : \exists \text{ good } j \text{ with } \alpha_i v_{ij} = \alpha_k v_{kj}\}$). Now, take the transitive closure T of this set $N(i)$ (i.e., include all bidders who are tied for an item with a bidder in $N(i)$, etc.). Next, redistribute the items that are tied such that none of the bidders in T are budget constrained while still allocating items completely. This is always possible because we can slightly increase the share of i for all items she is tied on while simultaneously reducing the share of all other bidders tied with her. In the next step, we can slightly increase the share of all these bidders for other items they are tied on while reducing the share of the new bidders they are tied with for those items and so on. Next, increase the pacing multipliers of all bidders in T by a small-enough $\delta > 0$ 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 redistribution of goods x' . This

contradicts that α was the maximal BFP to begin with as (α', x') is a BFP, yet it has pacing multipliers that are higher than in (α, x) . \square

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

Lemma 3. Consider two BFPMs $(\alpha^{(1)}, x^{(1)})$ 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. Because $(\alpha^{(1)}, x^{(1)})$ was budget feasible, I was not breaking its collective budget before. Because I is spending strictly less, at least some bidders must not spend their entire budget and thus, are 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 BFP.

Although the pacing multipliers are uniquely determined, the allocation is not. Tiebreaking may give different goods to different bidders. However, tiebreaking is inconsequential in the sense that the bidder utilities (and thus, social welfare), item prices (and thus, revenue), and the set of budget constrained bidders are all uniquely determined.

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

Lemma 4. Given two BFP $(\alpha^{(1)}, x^{(1)})$ and $(\alpha^{(2)}, x^{(2)})$, where $\alpha^{(1)} \geq \alpha^{(2)}$, we must have that $\text{Rev}(\alpha^{(1)}, x^{(1)}) \geq \text{Rev}(\alpha^{(2)}, x^{(2)})$.

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

Corollary 2 (Revenue Maximizing). The FPPE is revenue maximizing among all BFP.

The following theorem summarizes the main properties of FPPE that follow from Lemma 2 and Corollaries 1 and 2.

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 BFP.

3. 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 prevailing prices; that is, for all $i \in N$, it holds that

$$\mathbf{x}_i \in \arg \max_{\mathbf{x}_i \in [0,1]^m} \left\{ \sum_{j \in M} (v_{ij} - p_j) x_{ij} : \sum_{j \in M} p_j x_{ij} \leq B_i \right\}.$$

2. Every item with a positive price is sold completely; that is, $p_j > 0 \Rightarrow \sum_i x_{ij} = 1$ for all $j \in M$.

We now introduce a refinement of competitive equilibrium that requires that *each individual dollar* of a bidder 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 $i \in N$, 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 $\beta_i = 1$.

Through the previous definition, 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 was, 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 that 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 j is not entirely sold, then by the second condition of competitive equilibrium, it must have price zero. This, in turn, implies that all bidders have value zero 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 ∞ because 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 the platform does not benefit from adding fake bids under an FPPE mechanism, unlike in the case of second price payments.

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

Proposition 1. FPPEs are *shill proof*.

Proof. Note that if we start from an FPPE and remove both a bidder and the items she wins, we still have an FPPE because 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 FPPEs: (a) an FPPE with shill bids, (b) an FPPE 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 at

least as much as the revenue of (c), and therefore, also at least as much as the revenue of (a). Thus, the seller cannot benefit from shill bids. \square

Akbarpour and Li (2020) observe that a first price single-item auction satisfies a property referred to as *credible*. This means that the platform cannot benefit from misrepresenting what other agents have done in the mechanism. The example illustrates that the FPPE mechanism does not necessarily 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. However, the auctioneer could lie to bidder 1, claiming that someone else had bid three for item 2, and charge bidder 1 a price of two for item 1. Meanwhile, she could charge bidder 2 a price of one for item 2 for a higher revenue overall.

However, an FPPE does have a *price predictability* guarantee; given any allocation, a bidder either pays its full value or pays her 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 who might form a coalition to benefit themselves.

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. Because an FPPE is a competitive equilibrium, if we treat money as a good, then we have a traditional locally nonsatiated Walrasian equilibrium in an exchange economy. Because a locally nonsatiated Walrasian equilibrium is in the core, an FPPE is also in the core (see, e.g., Mas-Colell et al. 1995 for a general exposition of Walrasian equilibria and cooperative game theory and Powell 2018 for an exposition under assumptions that match those of this paper). \square

4. Monotonicity and Sensitivity Analysis

In the previous section, we showed that FPPEs are guaranteed to exist, that they are essentially unique (up to ties that are largely inconsequential), and that they satisfy a number of attractive properties. We now look at how well-behaved FPPEs are under changing conditions. We would ideally like the solution concept to be stable, so that changes in the input do not produce disproportionate changes in the output. We will show that this is largely the case. This is in stark contrast to SPPE, where Conitzer et al. (2021) showed that the equilibrium can be very sensitive. First, SPPE is not unique, and the revenue and welfare can vary drastically across equilibria. Second, even when there is a unique SPPE, small changes in the budget can cause disproportionately large changes in revenue.

4.1. Monotonicity

We investigate whether FPPEs are monotonic when adding bidders or goods or when increasing budgets or valuations. Table 2 summarizes our results.

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

Proposition 3. *In an 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 that 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 (they are bidding higher, and it is first price). However, this is impossible because all of them were running out of budget before because 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 has not changed (i.e., $T = \{i \in N \mid \alpha_i = \alpha'_i\}$). They must win weakly more items. Any item they were tied to originally with bidders outside T must now go completely to bidders in T . Additionally, bidders in T may win (part of) the new item. Because the pacing multipliers of bidders in T did not change, their prices did not change; hence, winning weakly more items means they are spending weakly more.

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

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

Proof. Let N be the original set of bidders, $i \notin N$ be a new bidder, and M be 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 an 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 B_i . After increasing the budget to B'_i , the solution (α, x) is still a BFPM. Therefore, by Lemma 2, the revenue of the new FPPE weakly increases. \square

Proposition 6. *In an 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: two bidders, two goods, $v_{11} = 10, v_{12} = 5, v_{21} = 0, v_{22} = 5$, with $B_1 = 10, B_2 = 5$. The FPPE consists of $\alpha_1 = \alpha_2 = 1, x_{11} = x_{22} = 1, x_{12} = x_{21} = 0$. Both bidders are spending their whole budget, so the total revenue is 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 a total revenue of 10 instead of 15. \square

4.1.2. Social Welfare. 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.

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

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

Now, add bidder 2 with $v_{21} = 2, B_2 = 1$. The new FPPE is $\alpha_1 = \frac{2}{K}, \alpha_2 = 1$, 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 an FPPE, adding a good weakly increases social welfare.*

Proof. Fix N, M , and an additional good $j \notin M$. Let (α, x) be the FPPE for N and M and (α', x') the FPPE

Table 2. Overview of Monotonicity Results

	Add bidder	Add good	Increase budget	Increase value v_{ij}
Revenue difference	≥ 0	≥ 0	≥ 0	Can be negative
Social welfare difference	Can be negative	≥ 0	Can be negative	Can be negative

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 will compare the contribution to social welfare of S and $N \setminus S$ separately.

First, we look at the set S . From the proof of Proposition 3, adding a good weakly decreases pacing multipliers. Because 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 $\frac{1}{\alpha_i}$ because 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 . Because 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, we do it for the set $N \setminus S$. By Proposition 3, the total revenue weakly increased. Because 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 per buck of one. In (α', x') , they have bang per buck of at least one; hence, their contribution to social welfare weakly increased.

Because 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 an 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): two bidders, one good. We have values $v_{11} = K, v_{21} = 2$ and budgets $B_1 = B_2 = 1$. The FPPE is $a_1 = \frac{2}{K}, s_2 = 1, x_{11} = x_{21} = \frac{1}{2}$ with a total social welfare of $\frac{K}{2} + \frac{1}{2}$.

Now, increase bidder 2's budget to $B'_2 = 2$. The new FPPE is $a_1 = \frac{3}{K}, s_2 = 1, x_{11} = \frac{1}{3}, x_{21} = \frac{2}{3}$ with a 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 an 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: two bidders, one 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 $x_{11} = x_{21} = \frac{1}{2}$, with a total social welfare of $\frac{k+1}{2}$. Now, increase $v'_{21} = 2$. The new FPPE consists of $a_1 = \frac{2}{K}, a_2 = 1$ and $x_{11} = \frac{1}{4}, x_{21} = \frac{3}{4}$ with a social welfare of $\frac{k+3}{4}$. As $K \rightarrow \infty$, we lose $\frac{1}{2}$ of the social welfare. \square

4.2. Sensitivity Analysis

We now investigate the sensitivity of FPPE to budget changes. An overview of our results is shown in Table 3. When adding Δ to the budget of a bidder, revenue can only increase and at most, by Δ . This shows that, in a sense, an FPPE is 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. (2021) show that in an SPPE, the revenue can change drastically, at least by a factor of 100.

Because of the nature of multiplicative pacing, additive bounds for social welfare (such as the ones given for revenue) do not immediately make sense.⁵ 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 an FPPE, changing one bidder's budget from B_i to $B'_i = (1 + \Delta)B_i$ for $\Delta \geq 0$ yields a modified FPPE (α', x') with $\alpha \leq \alpha' \leq (1 + \Delta)\alpha$.*

Proof. Fix the instance (N, M, V, B) , and let (α, x) be its FPPE. Let $B'_i = (1 + \Delta)B_i$ and (α', x') be the FPPE for (N, M, V, B') . Note that (α, x) is a BFPM for (N, M, V, B') , so $\alpha \leq \alpha'$ by Corollary 1. For the other inequality, note that $(\frac{\alpha'}{1+\Delta}, x')$ forms a BFPM for the original instance (N, M, V, B) . Indeed, all prices drop by exactly a factor $\frac{1}{1+\Delta}$, which means that with the same allocation x' , the spend for all bidders goes down by a factor $\frac{1}{1+\Delta}$ so that 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 per buck of buyers (i.e., the ratio between value and spend), so the bound in revenue change implies a bound in social welfare change.

Proposition 11. *In an 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 the 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') . Because (α, x) is a BFPM for the new instance, we have $\alpha' \geq \alpha$; that is, 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 $\text{Rev}_+^{\text{old}}$ be the revenue from them in (α, x) .

Table 3. Overview of Sensitivity Results

	Maximal decrease	Maximal increase
Revenue (additive)	0	Δ
Social welfare (relative)	$\frac{1-\Delta-\Delta^2}{1+\Delta}$	$1+\Delta$

Notes. 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., $B'_i = (1 + \Delta) \cdot B_i$).

Because the pacing multipliers for all bidders in S_+ strictly increased, they must have had $\alpha_k < 1$, so by the definition of an FPPE, they must have spent their entire budget and $\text{Rev}_+^{\text{old}} = \sum_{k \in S_+} B_k$. In the new FPPE, they cannot spend more than their budget, so $\text{Rev}_+^{\text{new}} \leq \sum_{k \in S_+} B'_k \leq (\sum_{k \in S_+} B_k) + \Delta = \text{Rev}_+^{\text{old}} + \Delta$.

What remains to be shown 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 wins 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 is not difficult to see that these extremes can both be attained. For the lower bound, increasing the budget of a nonbudget-constrained bidder will not change the FPPE; hence, revenue is unchanged. On the upper bound, take one bidder, one good, $v_{11} = 2\Delta, B_1 = \Delta$. Setting $B'_1 = B_1 + \Delta$ will increase revenue by Δ .

From Proposition 9, 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 an FPPE, changing one bidder's budget from B_i to $B'_i = (1 + \Delta)B_i$ 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 $S_1 = N \setminus S_p$ be the unpaced bidders. We will lower bound the new revenue from S_p and S_1 separately.

For the bidders in S_p , they spend their entire budget in both (α, x) and (α', x') ; we have that $\alpha \leq \alpha'$ by Lemma 5, and thus, bidders that are paced in α' are also paced in α . 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 is $SW_p^{\text{new}} \geq \frac{SW_p^{\text{old}}}{1+\Delta}$ for each $k \in S_p$, and therefore, $SW_p^{\text{new}} \geq \frac{SW_p^{\text{old}}}{1+\Delta}$.

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 Proposition 5, bidder i 's spend increased by at most $\Delta \cdot B_i$, and 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 $\Delta \cdot B_i \leq \Delta \cdot SW_1^{\text{old}}$. For unpaced bidders, spend equals contribution to social welfare, so we have $SW_1^{\text{new}} \geq SW_1^{\text{old}} - \Delta \cdot SW_1^{\text{old}}$.

Combining everything, we have $SW^{\text{new}} = SW_p^{\text{new}} + SW_1^{\text{new}} \geq \frac{SW_p^{\text{old}}}{1+\Delta} + SW_1^{\text{old}} - \Delta \cdot SW_1^{\text{old}} \geq \frac{SW_p^{\text{old}}}{1+\Delta} - \Delta \cdot SW_1^{\text{old}} = \left(\frac{1-\Delta-\Delta^2}{1+\Delta}\right)SW^{\text{old}}$. \square

Proposition 13. *In an FPPE, changing one bidder's budget from B_i to $B'_i = (1 + \Delta)B_i$ 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 multipliers strictly lower than one and that 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} be the new one. Define SW_i , SW_+ , SW_- , and SW_1 as the contribution to social welfare of bidders i , 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 per buck can only have decreased. Spend increased at most by $(1 + \Delta)$; bang per buck was at most the same, and hence, SW cannot exceed $1 + \Delta$ 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 per buck strictly decreased because of increasing pacing multipliers.

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

Finally, for bidders S_1 , we have $SW_1^{\text{new}} \leq SW_1^{\text{old}}$. First, note that the total spend of bidders in $S_- \cup S_1$ cannot have increased; their pacing multipliers stayed the same, so this would require them to win strictly more of some good they were previously sharing with a bidder not in $S_- \cup S_1$, but they are no longer winning any such goods. Because we have seen that the spend of bidders in S_- stayed the same, the spend of bidders in S_1 cannot have increased. Because their bang per buck is one, 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

5. 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 quasilinear utilities correspond exactly to FPPE in our setting. Cole et al. (2017) give the following primal and dual convex programs for computing a solution to a Fisher market with quasilinear utilities:

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

$$\sum_i x_{ij} \leq 1, \quad \forall j \quad (2)$$

$$\min_{p \geq 0, \beta \geq 0} \quad \sum_j p_j - \sum_i B_i \log(\beta_i) \quad \forall i, p_j \geq v_{ij} \beta_i \quad \beta_i \leq 1.$$

We show the primal convex program on top, which we denote by CP, and the corresponding dual convex program on the bottom. The variables x_{ij} denote the amount of item j that bidder i wins. The leftover budget is denoted by δ_i , and 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.

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 .

We now show via Fenchel duality that CP computes an FPPE. Informally, the result follows because β_i specifies a single utility rate per bidder and duality guarantees that any item allocated to i has exactly rate β_i ; thus, because 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 CP corresponds to an FPPE with pacing multiplier $\alpha_i = \beta_i$ and allocation x_{ij} and vice versa.*

Proof. We start by listing the primal Karush-Kuhn-Tucker (KKT) conditions.

1. $\frac{B_i}{u_i} = \beta_i \Leftrightarrow u_i = \frac{B_i}{\beta_i}$
2. $\beta_i \leq 1$
3. $\beta_i \leq \frac{p_j}{v_{ij}}$
4. $x_{ij}, \delta_i, \beta_i, p_j \geq 0$
5. $p_j > 0 \Rightarrow \sum_i x_{ij} = 1$
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; CP has the exact packing constraints. Budgets are also satisfied (here, we may assume $u_i > 0$; otherwise, budgets are satisfied because the buyer wins no goods). By KKT conditions (1) and (7), we have that for any good 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 Constraint (1) of CP always holds exactly because 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 $u_i = B_i$. 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 good j ($x_{ij} > 0$), then KKT condition (7) implies that the price on good 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 good j ; KKT conditions (3) and (7) guarantee that $\alpha_i v_{ij} = p_j \geq \alpha_{i'} v_{i'j}$ 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 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 because each bidder either gets a utility rate of one 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 because $\alpha_i \in [0, 1]$. KKT condition (3) is satisfied because each good bidder i wins has price per utility $\alpha_i = \frac{p_j}{v_{ij}} = \beta_i$, and every other good 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 good 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 CP to compute an FPPE. Cole et al. (2017) show that 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 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.

Borgs et al. (2007) also gave a convex program. However, that convex program is a set of feasibility constraints (in fact, those constraints correspond to the KKT conditions of CP). Because of this, their convex program cannot easily be solved via first-order methods, which are suitable for very large-scale problems. Similarly, online convex optimization can potentially be applied to CP (or its dual) for solving online variants of FPPE, whereas one cannot do this for the convex program of Borgs et al. (2007). See, for example, Gao and Kroer (2020, 2021) for recent work that builds on our results in this direction.

The equivalence between solutions to CP and FPPE provides an alternative view on many of our results. Because Theorem 3 can be shown directly via Fenchel duality, it allows us to prove via duality theory that FPPE corresponds to ERCE and that FPPE always exists.⁶

6. Experiments

In previous sections, we have shown that FPPEs have many satisfying theoretical properties, such as guaranteed existence, uniqueness, and polynomial time computability. However, the empirical question remains. If participating in a single first price auction has such bad properties that the industry shifted away from them in the past, does that assessment change when the auctions are within a budget management system? Additionally, although FPPE have nice properties compared with SPPE, how do the equilibria for first price auctions compare with those for second price auctions for typical objectives like social welfare and for typical instances? We investigated the properties of FPPE via numerical simulations⁷ on realistic instances generated from data collected at Facebook and Instagram. We aimed to answer two concrete questions. (1) Under FPPE, how high is bidder regret for reporting truthfully? (2) How does FPPE compare with SPPE in terms of revenue and social welfare?

6.1. Data

Similarly to instances generated by Conitzer et al. (2021), we construct realistic instances from real-world auction markets in two steps. We first take bidding data for a region during a period and use it to create n bidders and m goods. To get the n bidders, we identify the top n advertisers that participate in

the most auctions in that period in that region. Each of those advertisers will map to a buyer in the final instance. As an intermediate step, we define the goods in the instance as the real-life auctions that include at least one of the n bidders. In this intermediate step, we set the bid in each bidders-auction pair to be the value of the bidder for the corresponding good.⁸ In a second step and to complete the construction, we reduce the size of the market so we can compute results more efficiently and thoroughly answer the questions of interest. To achieve that, we cluster the real-world auctions (the goods of the intermediate step) into a smaller number of goods. We apply the k -means algorithm using the n -dimensional vector of values for that good as features. The goods in the final instances are the resulting auction clusters. Each bidder valuation for a (cluster-level) good is set to the average of their valuations for goods in the cluster. To generate the budgets, we set it equal to the expected value that the bidder would receive in a uniform random allocation of goods to bidders (i.e., $B_i = \frac{1}{n} \sum_j v_{ij}$). The motivation for this is that it leads to a good mixture of budget-constrained and unconstrained bidders because in aggregate, this constrains the sum of prices to be the sum of average valuations, whereas it would be the sum of maximum valuations if every buyer was unconstrained. The distribution of budget-constrained and unconstrained bidders is also similar to that of Conitzer et al. (2021), despite a slightly different process for setting budgets.⁹

We collected bidding data from Facebook and Instagram for 7 days for auctions in a chosen country. For each day and each of those two platforms, we collect the bids of the top n advertisers (where $n \in \{6, 8, 10, 12, 14\}$) with the most impressions that day. We consider $m \in \{10, 20, 30\}$ goods where each good is a representative cluster of auctions. In the end, each combination of platform, day, number of bidders, and number of goods becomes an instance of our study, resulting in a total of $2 \times 7 \times 5 \times 3 = 210$ instances. To compare with SPPE, we take the same data and consider a set of instances with $\{3, \dots, 8\}$ bidders and $\{4, \dots, 8\}$ goods for a total of $2 \times 7 \times 6 \times 5 = 420$ instances. The SPPE instances are smaller because equilibria are harder to compute.

Although we present our results on moderate-sized instances, we also performed some tests on larger synthetic data to verify that CP can indeed be solved efficiently. For these tests, we used the Mosek interior-point method (Dahl and Andersen 2021). We found that for uniform random valuations, instances with 4,000 buyers and 4,000 items can be solved in about 12 minutes, and run time scaled approximately linearly with the size of the valuation matrix. In addition to these encouraging experiments via conic solvers, it was recently shown that CP can be solved using first-order methods (Gao and Kroer 2020), which should

enable further scalability, at the cost of slightly more inaccurate solutions.

6.2. Incentive Properties

We start by numerically analyzing the incentive properties of FPPE. Our analysis will focus on two types of incentives: first, *ex post* incentives to shade bids because of the first price auction rule, when other buyer's bids are held constant, and second, *ex ante* incentives to misreport in order to shift the FPPE outcome itself when viewed as a mechanism. For both types of incentives, we will focus on the proxy-bidder setting, where advertisers typically report a value per conversion and budget, whereas the valuation for individual auctions is determined by the platform as the value per conversion times the conversion rate. Thus, the space of possible misreports is substantially reduced; buyers can misreport their overall value per conversion and/or their budget (as opposed to individually shading or misreporting in each auction).

6.2.1. Analysis of Regret. First, we look at the *ex post* regret that each bidder has in FPPE as compared with being able to unilaterally deviate to a different pacing multiplier while keeping the FPPE multipliers fixed for all other bidders. The regrets resulting from our computational study are shown in Figure 2. The figure shows summary statistics over relative *ex post* regret, which is the fraction of utility for the best-response pacing multiplier that is lost if the bidder reports truthfully. Each data point represents the relative regret that one of the bidders experienced in the FPPE of each instance; we generate a data point for each bidder in the instance. The middle line in each box (which is at 0.00 for all breakdowns) is the median; the lower and upper edges of the boxes represent the first and third quartiles, respectively. The lines extending from the box show the rest of the distribution

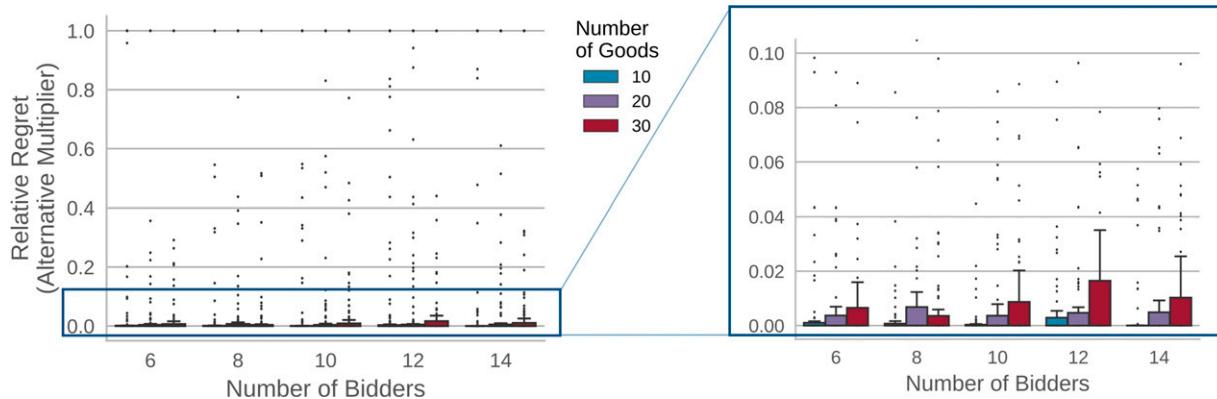
except for outliers that are determined by a function of the interquartile range, and the dots represent individual outliers outside that range.

The results show that across all breakdowns of number of bidders and number of goods, the relative regret is small. In all cases, the median is at 0% relative regret, and the highest upper quartile represents less than 2% relative regret; therefore, bidders across the board do almost as well in FPPE as they would if they could optimally pick their pacing multiplier (given fixed competition). We do see that the benefit to deviating increases somewhat with the number of goods (instances with 10 goods show consistently lower regret than instances with 20 or 30 goods), whereas there does not seem to be a strong dependence on the number of bidders.

To put the previous analysis in perspective, Appendix A presents results for an unrealistic model where bidders can optimally set their paced bids for each individual auction independently.¹⁰ Even under these conditions, across all breakdowns, 75% of bidders experience less than 30% relative regret, and as the number of bidders grows (compared with the number of goods), the majority of bidders in instances have 0% relative regret.

6.2.2. Misreporting to Influence Equilibria. After providing evidence that bidders have low *ex post* regret if they had chosen the best pacing multiplier while keeping the competition fixed, we look at whether bidders can influence the FPPE outcome itself by submitting conveniently chosen (but potentially wrong) parameters to the FPPE mechanism. We assume that a focal bidder can misreport by scaling the valuations and budget by scalars chosen from the set $\{0.1, 0.2, \dots, 1.1\}^2$. In other words, the bidder would select a tuple (λ_v, λ_b) and submit the true valuations multiplied by λ_v and the true budget multiplied by λ_b . Note that the bidder does

Figure 2. (Color online) Summary Statistics of Relative Ex Post Regret in FPPE (Ratio of Best-Response Utility Keeping Competitors' Bids Fixed to Utility Under the FPPE)



Note. The plot on the right is a zoomed-in version of the plot on the left.

not have the ability to scale values differently for different auctions (but see Appendix A for an analysis of that case). All other bidders' bids and budgets remain unchanged. Figure 3 shows the relative regret with respect to the highest utility at an FPPE amongst the possible ways to misreport. In almost all cases, there is no or negligible gain. We note also that the incentive to misreport is even lower than the incentive for ex post shading discussed in the previous section (of the order of 0.001 versus of the order of 0.01 relative regret). This is likely because, in most cases, the ideal ex post shading for a single buyer leads to budget infeasibility for other buyers, which in turn, makes that outcome infeasible.

6.2.3. Regret as a Function of the Tightness of Budget Constraints.

Although the majority of bidders in these experiments have surprisingly low regret under FPPE, there are some outliers that have high relative regret. To investigate where this variability comes from, we look at the regret that a bidder experiences as a function of their pacing multiplier at equilibrium. Intuitively, a bidder who has a low pacing multiplier gets a high bang per buck for each impression they buy. This implies that underbidding loses value faster than it saves cost, whereas a bidder with a pacing multiplier equal to one does not get any utility in each auction won. Our computations indicate that bidder budgets, pacing multipliers, and regrets all have a positive association.

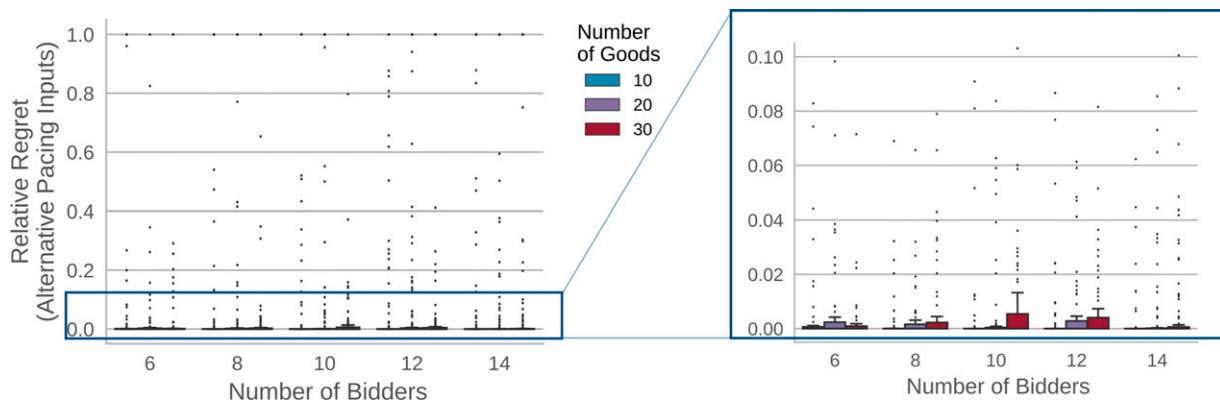
Figure 4 plots regrets for these instances as a function of the pacing multiplier, where each point represents a bidder of a pacing instance. The left panel is a scatterplot of the relative regret defined as the fraction of utility that a bidder loses in the FPPE compared with the best-response pacing multiplier. Because the regret for the majority of data points is approximately zero, we also plot a curve corresponding to a locally weighted scatterplot smoothing (LOWESS) regression

to capture the "typical" behavior.¹¹ The typical relative regret increases when the pacing multiplier increases because of the smaller relative frequency of points with no regret and a noticeable mass of points located at (1, 1). At first glance, this implies that when the pacing multiplier exceeds 0.4, the typical regret grows linearly from 0 to about 0.1. The points at (1, 1) represent unpaced bidders that receive no utility in the FPPE; although there is no numerical error in the calculation, relative regrets do not fully depict the situation. For example, if a bidder gets impressions they value at \$100 and pays \$100 in the FPPE, although a best-response pacing multiplier may make them pay \$98 to get \$99 worth of impressions, their relative regret is one, but in absolute terms, the outcomes are very similar. Consequently, we complement the analysis with an absolute notion of regret that captures the utility gained under a best response. To more easily compare bidders against each other, we normalize the utility gained by the total value of goods received for the best-response pacing multiplier. The results are displayed in the right panel of Figure 4. The point at (1, 1) has disappeared, and the best-fit line for bidders with a pacing multiplier equal to one shows a regret down to roughly 1.5%. From the distribution of instances (the black dots), it can also be seen that small pacing multipliers are more common. So, although positive deviations do exist, on average the benefit is small.

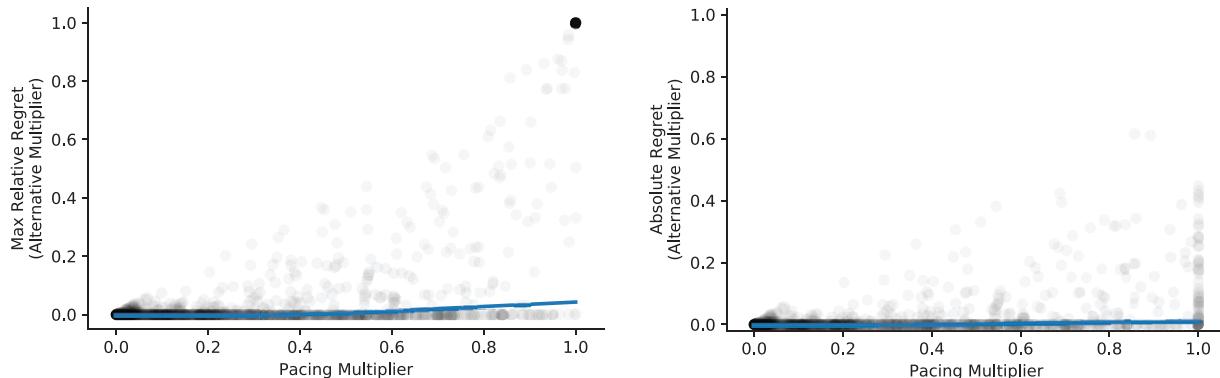
6.3. Revenue and Social Welfare

We now compare revenue and social welfare under FPPE and SPPE, as shown in Figure 5. The left panel shows the cumulative distribution function (cdf) of the ratio of FPPE revenue to that of SPPE. The right panel is similar but with social welfare. We see that FPPE revenue is always higher than SPPE revenue, although both coincide for about 30% of instances, and almost never more than 4.5 times as high. For

Figure 3. (Color online) Summary Statistics When Bidders Can Misreport Their Value per Conversion or Budget to Manipulate the FPPE



Note. The plot on the right is a zoomed-in version of the plot on the left.

Figure 4. (Color online) Regret as a Function of the FPPE Pacing Multiplier

Notes. Each point represents the regret of an individual bidder; the curve summarizes the points using an LOWESS regression. The left and right panels display relative and absolute regret, respectively.

social welfare, we find that, perhaps surprisingly, neither solution concept is dominating, with most instances having relatively similar social welfare under either solution concept, although FPPE does slightly better. There are two caveats to keep in mind for these results. (a) 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. (2021) find that most instances admit a single equilibrium). (b) Many bidders are budget constrained in the FPPE of our setting, and so, these insights might not translate to cases where many bidders are not budget constrained (see Appendix B.2 for statistics on multipliers in the two solution concepts). These experiments show that an FPPE is not necessarily worse than an SPPE with respect to social welfare (at least with nonstrategic bidders) while potentially having a significantly higher revenue.

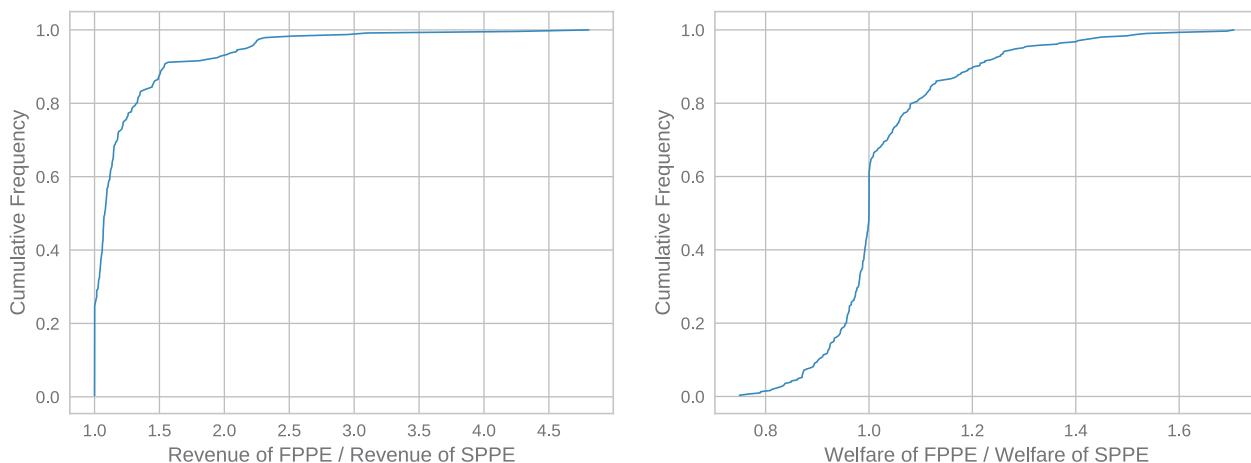
6.4. Robustness

In order to test the robustness of our results, we ran the same set of experiments on synthetic data (see

Appendix B for details). Qualitatively, the results on instances with values drawn from a $\text{unif}(0, 1)$ distribution are similar, although regret tends to be higher in the synthetic data sets. One hypothesis is that values in the realistic data are correlated (for example, because certain users are overall more likely to interact with ads), and this results in stronger competition for high-value auctions.

7. Discussion of Modeling Assumptions

Our model captures several salient features present in contemporary real-world internet advertising systems; budget management is done via a multiplicative scalar after buyers report a value per conversion and a budget. Because of the deterministic valuations, we can allow general valuation structures as opposed to independent and/or identical stochastic valuations. That said, there are also a number of practical issues that we do not capture. We discuss these limitations as well as potential ideas on how to relate them to our model.

Figure 5. (Color online) Summary Statistics over the FPPE/SPPE Ratio of Revenue (Left Panel) and Social Welfare (Right Panel)

7.1. Stochastic Valuations

In line with the information available by the setting explored here, our model assumes that all valuations are known ahead of time, and consequently, we have a completely determined market. The literature generally considers models in which user types are realized according to some sort of stochastic process, and each buyer's valuation is sampled according to a distribution that depends on the user type. Other authors have modeled the buyer's valuations as being drawn independently each time a user arrives (Balseiro et al. 2015, 2021; Balseiro and Gur 2019). Although such models address the stochastic aspects of the process, they fail to capture correlations. Conversely, our model does capture the correlation but fails to capture the stochastic aspect. In practice, we would ideally want a model that captures both, but such models are more cumbersome to analyze. Nonetheless, our model can be viewed as a discrete approximation to this kind of process.

7.2. Dynamics

In practice, platforms would typically have a model of what will occur; they know which advertisers are active in the system, their stated budgets, and values per conversion as well as the distribution of users. However, the platform does not know *a priori* which subset of users will actually arrive on a given day and thus, will not know which exact auctions will materialize. Instead, these auctions materialize sequentially as the day unfolds and users visit the platform. In practice, this uncertainty is handled by a control algorithm that tunes the pacing multiplier to the right spending rate with the goal of exhausting the budget at the end of the planning horizon. Although we do not model these dynamics, our model can be related to dynamic markets in several ways. First, our model captures the *ideal* outcome; when each proxy bidder operates an independent control process on a pacing multiplier, the hope is that these pacing multipliers stabilize in a steady state. Our model describes what that steady state should be. Second, in practice, the platform would be able to create a good model of the overall composition of the market. Our computational results can be used to compute static pacing multipliers on such a market model. These static pacing multipliers can be used directly in the dynamic market. They can also be used to seed the control algorithm, which further refines the pacing multiplier as the market progresses. Third, showing that pacing equilibria are directly connected to the EG convex program opens the door to applying *online convex optimization* techniques for converging to FPPE in an online fashion. Although the exact details of how to achieve this are beyond the scope of this paper, we believe that it could be achieved using the online variant of regularized

stochastic dual averaging (Xiao 2010). In recent follow-up work to our paper, Gao and Kroer (2021, theorem 7) showed how to do this in a related Fisher market model.

7.3. Multiple Slots

In practice, each auction may allocate more than a single item. Typically, a visit of a user triggers a multi-item auction (for example, several ad slots are auctioned off simultaneously both for Facebook and Instagram feed ads and for Google search auctions). Similarly to other authors, we abstract away this aspect of the problem and assume that each item can be sold in isolation (Balseiro et al. 2015, 2021; Balseiro and Gur 2019; Conitzer et al. 2021). Modeling multiple slots would require substantial changes in the analysis of the auction mechanism, and it would be harder to leverage the relationship to market equilibria.

8. Conclusion

In an advertiser 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 ensuing aggregate behavior. In the case of feed-based advertising platforms, advertisers commonly control their campaigns through targeting criteria and budgets, with no ability to directly control individual auctions. This substantially changes the advertising platform's system design problem compared with other settings. In this paper, we have seen that the steady-state properties of a pacing-based budget management system for such scenarios are in fact quite good when a first price auction is used to sell each impression. We have showed that first price auctions enjoy several properties that second price auctions do not: uniqueness, monotonicity, and computability among others. At the same time, SPPE enjoys other advantages, such as no static shading incentives (although our experiments on real ad auction data suggest that this may not be a major concern in FPPE either).

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 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 advertising 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, core membership, shill proofness, false name proofness, etc.—suggesting that first price auctions

may, in fact, have a serious role to play in today's advertising platforms.

Taking a broader perspective, we recognize that there is generally no one-size-fits-all answer to the design of a budget management system because no class of systems strictly dominates the others, and trade-offs remain. This is readily evidenced by the fact that different large platforms have adopted different solutions in terms of both which budget smoothing method to adopt and which auction format to run. Many factors will be weighed by any real-world system. For example, comparing multiplicative pacing systems with throttling systems, the former might be appealing because they encapsulate bidders' natural incentive to shade their bids, whereas the latter makes it easier for bidders to get a somewhat more representative sample of available impressions. Similarly, we see different solutions in use in the real world between first price and second price mechanisms. What we have done in this paper is to put forward a framework where the two auction formats can be compared from an end-to-end perspective that captures several design choices of real-world advertising platforms.

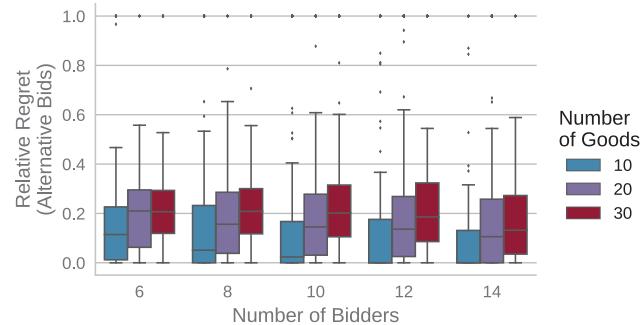
Acknowledgments

The authors thank Santiago Balseiro, Ozan Candogan, Roberto Cominetti, Nikhil Devanur, Yoni Gur, Renato Paes Leme, Marco Scarsini, and Marc Schroder for useful discussions on the technical content of this paper. In addition, the authors extend thanks to the review team of *Management Science*; conference program committee members and reviewers at the ACM Conference on Economics and Computation; participants at Duke Computer Science-Economics seminar, Michigan Ross, Massachusetts Institute of Technology Operations Research Center, and Stanford Graduate School of Business; and attendees of the AAMAS-IJCAI workshop on Agents and Incentives in Artificial Intelligence, the Artificial Intelligence and Marketing Science workshop at the AAAI Conference on Artificial Intelligence, the Algorithmic Game Theory and Data Science workshop at the ACM Conference on Economics and Computation, the Barbados Algorithmic Game Theory workshop, the International Conference on Game Theory, and the Schloss Dagstuhl seminar on Traffic Models. Their remarks and questions allowed the authors to significantly improve the manuscript and its presentation. This work was done while C. Kroer and C. A. Wilkens were working full time at the Core Data Science team at Meta.

Appendix A. Supplemental Experiments on Realistic Instances

In addition to the regret analysis in the body of the paper (which is with respect to the ex post best pacing multiplier given a fixed competition), in this section we also consider a much stronger model of deviations where bidders can optimally place individual bids in auctions corresponding to each good. This allows us to put the regret analysis discussed in

Figure A.1. (Color online) The Relative Regret of Advertisers in Compared with a Hypothetical Ability to Directly Control Bids in Each Individual Auction



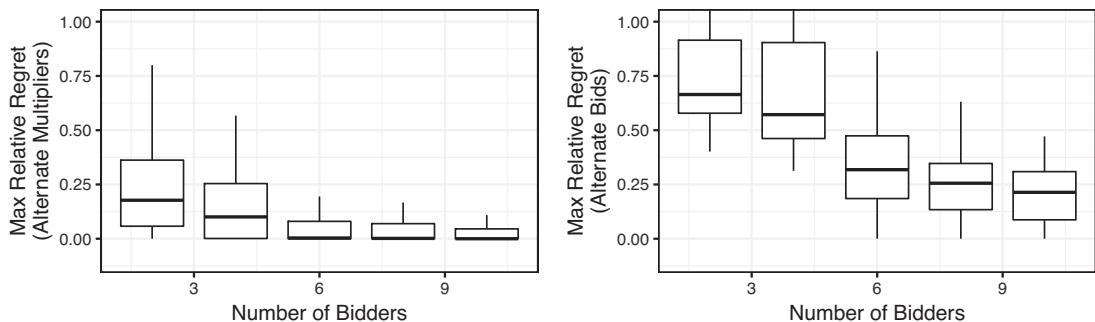
the body of the paper in perspective compared with what can be considered as an upper bound on regret. Recall that the platform computes the effective bid as conversion rate \times value per conversion \times pacing multiplier for the bidder-good pair. Although in practice, the advertiser only sets the value per conversion, in this section we assume that bidders can choose the optimal bid by influencing the platform's conversion rate estimates for the bidder-good pair. Consequently, the bidder can effectively control the individual bid placed in each of the auctions.

The resulting relative regret is shown in Figure A.1. As in the body, the middle line in each box represents the median; the lower and upper ends of the box show the first and third quartiles, respectively; the lines extending from the box show outliers within 1.5 times the interquartile range; and the dots represent individual outliers outside that range. For all parameter settings (number of bidders and number of goods), 75% of bidders in all instances have less than 30% relative regret. Considering how much control bidders have on the deviations in this analysis, it is rather surprising that in virtually all instances, the FPPE outcome yields more than 70% of the optimal utility achievable. Moreover, we can see that the distribution of regret (in particular, the three quartiles) gets smaller as the competition in the market increases (e.g., when the number of goods decreases or the number of bidders increases). For 12 or more advertisers and 10 users, the median regret is zero, indicating that more than half of all the buyers get their optimal utility in the equilibrium.

Appendix B. Supplemental Experiments on Synthetic Instances

In the body of the paper, we have shown that based on realistic data, advertisers have low regret in FPPE compared with alternative inputs or pacing multipliers and that FPPE generates more revenue than SPPE, with roughly similar welfare. To understand whether these results are specific to the data that we used or if these results continue to hold for very different data, we repeat the analysis on synthetic data where bidders have valuations drawn from a uniform distribution. Because realistic data may have correlated valuations and valuation distributions with a long tail, these synthetic data represent a significant departure from realistic data. Although the data are quite different, this is a robustness check where we find qualitatively similar insights, as described.

Figure B.1. Summary Statistics over the Maximum Relative Ex Post Regret (Maximum Taken over Bidders in a Given Auction; the Statistic Taken over the Maximum Regret Across Instances)



Notes. (Left panel) The bidder can choose a single alternative pacing multiplier. (Right panel) The bidder can choose an individual bid for every auction.

We generated instances according to the complete-graph model of Conitzer et al. (2021). In this model, every bidder is interested in every good, and each valuation is drawn independently from $\text{unif}(0, 1)$. We generated five instances for each point in the Cartesian product of $\{2, 4, 6, 8\}$ bidders and $\{4, 6, \dots, 14\}$ goods. For each instance, we computed (a) an FPPE as the solution to CP using the solver CVX (Grant and Boyd 2014) and (b) an SPPE for every objective function using the mixed-integer program (MIP) given by Conitzer et al. (2021). We considered at most eight bidders and 14 goods because of the limited scalability of the SPPE MIP; we were able to solve all CP instances in less than 2 ms.

B.1. Incentive Properties

B.1.1. Analysis of Regret. Similarly to what we have done for the realistic instances, first we look at the ex post regret of each bidder under an FPPE compared with being able to unilaterally deviate while keeping the FPPE multipliers fixed for all other bidders. We consider two sets of deviations: jointly setting bids through a best-response pacing multiplier and the more powerful model of individually setting bids in each auction.

The resulting relative regrets are shown in Figure B.1. The figure shows summary statistics over the maximum 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. 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, respectively. The line extending from the box shows outliers within 1.5 times the interquartile range.

Compared with the results on realistic data, the regret experienced by bidders is higher, but qualitatively, we see similar effects; the regret decreases when market density increases. In particular for 6, 8, and 10 bidders, the median regret for alternative multipliers is zero. Strikingly, the max relative regret under the model where bidders can set individual bids does very poorly for two and four bidder instances. This is likely because the expected difference between the first- and second-order statistics of n draws from a uniform distribution (which shrinks as $O(\frac{1}{n})$), which is the difference between the highest and second highest bids.

B.1.2. Misreporting to Influence Equilibria. Figure B.2 shows the resulting regrets when a single bidder can misreport by scaling the valuations and budgets with factors $(\lambda_v, \lambda_b) \in [0.1, 0.2, \dots, 1.1]^2$, so the inputs provided to the proxy bidder are the true figures times those scalars. Similar

Figure B.2. Summary Statistics When a Bidder Can Misreport Their Values and Budget to Influence the FPPE

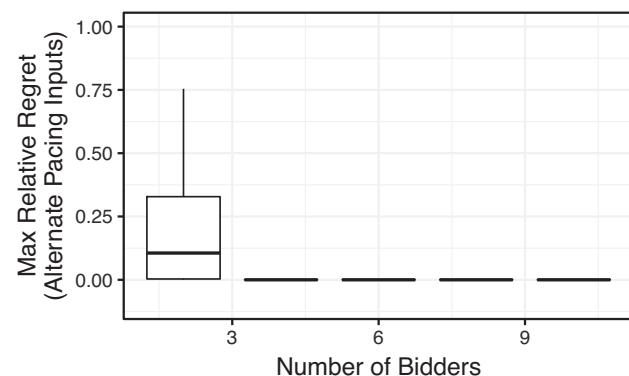


Figure B.3. (Color online) Relative Regret as a Function of the FPPE Pacing Multiplier Run on Instances with Budgets Scaled by a Factor of Two

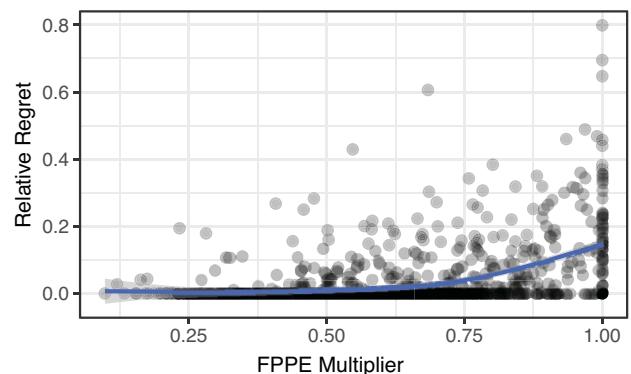
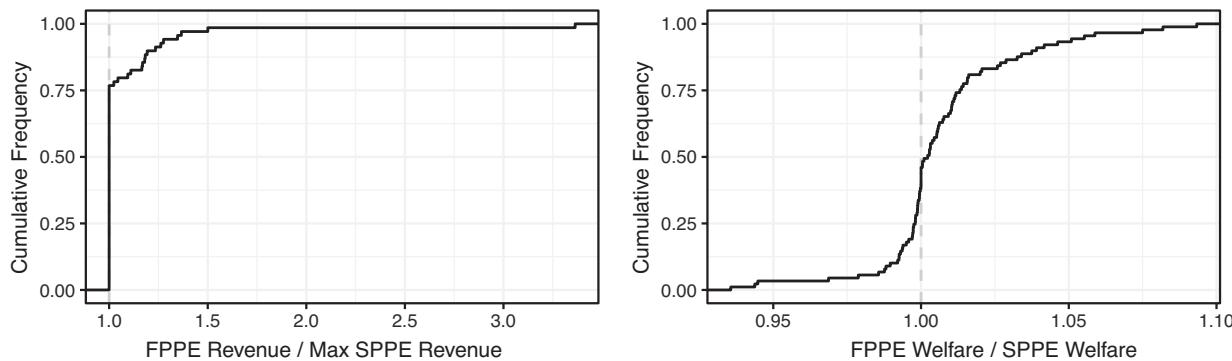


Figure B.4. Distribution of the Ratio of FPPE to SPPE Revenues (Left Panel) and Social Welfare (Right Panel)

to the realistic instances, the regret is generally lower than for alternative pacing multipliers.

B.1.3. Regret as a Function of the Tightness of Budget Constraints.

We consider the effect of the pacing multiplier on regret. To generate instances with larger pacing multipliers, we compute FPPE on a modified set of problem instances, which are identical to instances in the preceding experiments, except that all budgets were scaled up by a factor of two. We show the relative regret for these instances in Figure B.3. Each point represents a bidder in one of the instances. The uptick in regret is slightly higher than for realistic data, although qualitatively similar.

B.2. Revenue and Welfare

In this section, we compare revenue and social welfare under FPPE and SPPE on an instance by instance basis. The comparison is shown in Figure B.4. The left panel shows the cdf of the distribution of ratios of FPPE and SPPE revenues, whereas the right panel does the same for the ratios of welfare. Putting in perspective the curves corresponding to realistic instances shown in Figure 5, we see that FPPE revenue is now the same for about 75% of instances (compared with 30% before) and almost never more than 1.5 times as high (compared with 4.5 times before). For welfare, we observe that the cdf of the ratio between FPPE and SPPE looks exactly the same compared with the realistic instances, albeit the x axes representing the ratios are at a different scale.

Endnotes

¹ Remnant inventory in display comprised ad opportunities that were not otherwise sold by a salesperson through a guaranteed contract.

² Specifying the bid is available as an advanced feature offered for certain types of campaigns. We refer the reader to Facebook ad product documentation, which includes a list of campaign types in a table where it can be seen in the second column that most of the campaign types are *autobid*. See <https://www.facebook.com/business/m/one-sheets/facebook-bid-strategy-guide>.

³ In 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 after the winner of an auction is selected, she pays her entire (unscaled) bid for the impression. In contrast, in our work, the individual bids are scaled using the respective pacing multipliers, and the winner pays the scaled bid instead of her

original bid for the impression. Another difference is that in the budgeted allocation literature, the scaling multiplier changes from one impression to the next because it is only a tool for winner selection, whereas we choose a single pacing multiplier for a bidder that is used to scale her bids for all the impressions in the problem instance.

⁴ More precisely, pacing multipliers are unique, but there may be different equivalent allocations because of tie breaking. Tie breaking does not impact revenue, social welfare, or individual utilities.

⁵ To 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 two. Changing a budget will yield the same allocation for both instances (and pacing multipliers are precisely a factor of two off), but the change in social welfare will be twice as large in the second instance.

⁶ CP is easily seen to always be feasible, and the feasible set is compact. Thus, CP always attains an optimal solution. By Theorem 3, that solution will be an FPPE.

⁷ There is a public notebook on synthetic data for these experiments available on GitHub at <https://github.com/facebookresearch/FPPE>.

⁸ The campaigns used in this study were themselves frequently budget paced. In the logs, we have both the original bid and the pacing multiplier that were used. We use the original bid in this step.

⁹ Scaling budgets might yield different results. When, for example, we scale budgets up by a factor of two, the fraction of budget-constrained buyers changes, and the quantitative results change as a result. We refer the reader to the publicly available notebook mentioned earlier to try out different budget levels in the synthetic experiments and explore the effects.

¹⁰ This deviation model is unrealistic in ad markets because a bidder would have to change the platform’s estimated conversion rate for each user in such a way that it leads precisely to the right paced bid for each auction. The reason we consider deviations of this kind is to show that even under unreasonably powerful assumptions of what a bidder can do, the relative regret is still low.

¹¹ The LOWESS regression is a nonparametric method that is frequently used for data analysis (Cleveland 1979). For instance, it is the default way to run a regression with smoothed data in R using the function `GEOM_SMOOTH`. At each point, it fits a local low-degree polynomial, where the input data are weighted by distance to that point.

References

Abrams Z, Sathiya Keerthi S, Mendelevitch O, Tomlin J (2008) Ad delivery with budgeted advertisers: A comprehensive LP approach. *J. Electronic Commerce Res.* 9(1):16–32.

Akbarpour M, Li S (2020) Credible auctions: A trilemma. *Econometrica* 88(2):425–467.

Ausubel LM, Milgrom P (2006) *The Lovely but Lonely Vickrey Auction. Combinatorial Auctions* (MIT Press, Cambridge, MA).

Azar Y, Birnbaum B, Karlin AR, Nguyen CT (2009) On revenue maximization in second-price ad auctions. Fiat A, Sanders P, eds. *Algorithms—ESA 2009, Lecture Notes in Computer Science*, vol. 5757 (Springer, Berlin), 155–166.

Balseiro SR, Gur Y (2019) Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Sci.* 65(9):3952–3968.

Balseiro SR, Besbes O, Weintraub GY (2015) Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Sci.* 61(4):864–884.

Balseiro S, Kim A, Mahdian M, Mirrokni V (2021) Budget-management strategies in repeated auctions. *Oper. Res.* 69(3): 859–876.

Bigler J (2019) Rolling out first price auctions to google ad manager partners. Accessed July 13, 2021, <https://www.blog.google/products/admanager/rolling-out-first-price-auctions-google-ad-manager-partners/>.

Borgs C, Chayes J, Immorlica N, Jain K, Etesami O, Mahdian M (2007) Dynamics of bid optimization in online advertisement auctions. *Proc. 16th Internat. Conf. World Wide Web (WWW)* (Association for Computing Machinery), 531–540.

Cary M, Das A, Edelman B, Giotis I, Heimerl K, Karlin AR, Mathieu C, Schwarz M (2007) Greedy bidding strategies for keyword auctions. *Proc. 8th ACM Conf. Electronic Commerce (EC)* (Association for Computing Machinery), 262–271.

Chen Y (2017) Programmatic advertising is preparing for the first-price auction era. Accessed July 13, 2021, <https://digiday.com/marketing/programmatic-advertising-readying-first-price-auction-era/>.

Chen X, Kroer C, Kumar R (2021) The complexity of pacing for second-price auctions. Preprint, submitted March 25, <https://arxiv.org/abs/2103.13969>.

Cleveland WS (1979) Robust locally weighted regression and smoothing scatterplots. *J. Amer. Statist. Assoc.* 74(368):829–836.

Cole R, Devanur N, Gkatzelis V, Jain K, Mai T, Vazirani VV, Yazdanbod S (2017) Convex program duality, Fisher markets, and Nash social welfare. *Proc. 2017 ACM Conf. Econom. Comput. (EC)* (Association for Computing Machinery), 459–460.

Conitzer V, Kroer C, Sodomka E, Stier-Moses NE (2021) Multiplicative pacing equilibria in auction markets. *Oper. Res.*, ePub ahead of print December 13, <https://doi.org/10.1287/opre.2021.2167>.

Dahl J, Andersen ED (2021) A primal-dual interior-point algorithm for nonsymmetric exponential-cone optimization. *Math. Programming*, ePub ahead of print March 9.

Dütting P, Fischer F, Parkes DC (2019) Expressiveness and robustness of first-price position auctions. *Math. Oper. Res.* 44(1):196–211.

Edelman B, Ostrovsky M (2007) Strategic bidder behavior in sponsored search auctions. *Decision Support Systems* 43(1):192–198.

Feldman J, Muthukrishnan S, Pal M, Stein C (2007) Budget optimization in search-based advertising auctions. *Proc. 8th ACM Conf. Electronic Commerce (EC)* (Association for Computing Machinery), 40–49.

Gao Y, Kroer C (2020) First-order methods for large-scale market equilibrium computation. *Adv. Neural Inform. Processing Systems (NeurIPS)* 33:21738–21750.

Gao Y, Kroer C (2021) Infinite-dimensional Fisher markets: Equilibrium, duality and optimization. *Proc. AAAI Conf. Artificial Intelligence* 35(6):5432–5439.

Goel A, Mahdian M, Nazerzadeh H, Saberi A (2010) Advertisement allocation for generalized second-pricing schemes. *Oper. Res. Lett.* 38(6):571–576.

Grant M, Boyd S (2008) Graph implementations for nonsmooth convex programs. Blondel V, Boyd S, Kimura H, eds. *Recent Advances in Learning and Control, Lecture Notes in Control and Information Sciences* (Springer, Berlin), 95–110.

Grant M, Boyd S (2014) CVX: Matlab software for disciplined convex programming, version 2.1. Accessed March 6, 2022, <http://cvxr.com/cvx>.

Hosanagar K, Cherepanov V (2008) Optimal bidding in stochastic budget constrained slot auctions. *Proc. 9th ACM Conf. Electronic Commerce (EC)* (Association for Computing Machinery), 20.

Jordan CA (2020) Simplifying mobile app advertising: Moving to a first price auction. Accessed July 13, 2021, <https://www.mopub.com/en/blog/first-price-auction>.

Karande C, Mehta A, Srikant R (2013) Optimizing budget constrained spend in search advertising. *Proc. Sixth ACM Internat. Conf. Web Search Data Mining (WSDM)* (Association for Computing Machinery), 697–706.

Kroer C, Peysakhovich A, Sodomka E, Stier-Moses NE (2021) Computing large market equilibria using abstractions. *Oper. Res.* 70(1):329–351.

Mas-Colell A, Whinston MD, Green JR (1995) *Microeconomic Theory* (Oxford University Press, Oxford, United Kingdom).

Mehta A (2013) Online matching and ad allocation. *Foundations Trends Theoret. Comput. Sci.* 8(4):265–368.

Mehta A, Saberi A, Vazirani U, Vazirani V (2007) Adwords and generalized online matching. *J. ACM* 54(5):22–es.

Milgrom P (2004) *Putting Auction Theory to Work* (Cambridge University Press, Cambridge, United Kingdom).

Nesterov Y, Shikhman V (2018) Computation of Fisher–Gale equilibrium by auction. *J. Oper. Res. Soc. China* 6(3):349–389.

Nisan N, Roughgarden T, Tardos E, Vazirani VV (2007) *Algorithmic Game Theory* (Cambridge University Press, Cambridge, United Kingdom).

Paes Leme R, Sivan B, Teng Y (2020) Why do competitive markets converge to first-price auctions? *Proc. Web Conference 2020*, 596–605.

Peysakhovich A, Kroer C (2019) Fair division without disparate impact. *Mechanism Design Social Good Workshop EC 2019 (MD4SG)*.

Powell M (2018) Economics 2010b: Microeconomic Theory II. Accessed July 13, 2021, <http://www.pareto-optimal.com/s/EC-2010B-Teaching-Notes-Full-GE-Feb-9-2018-blbz.pdf>.

Rusmevichientong P, Williamson DP (2006) An adaptive algorithm for selecting profitable keywords for search-based advertising services. *Proc. 7th ACM Conf. Electronic Commerce (EC)* (Association for Computing Machinery), 260–269.

Sluis S (2017) Big changes coming to auctions, as exchanges roll the dice on first-price. Accessed July 13, 2021, <https://adexchanger.com/platforms/big-changes-coming-auctions-exchanges-roll-dice-first-price/>.

Xiao L (2010) Dual averaging methods for regularized stochastic learning and online optimization. *J. Machine Learn. Res.* 11(3): 2543–2596.