

Lecture Note 10: Auctions with Budgets

Christian Kroer*

March 2, 2022

1 Introduction

The previous lecture note introduced novel aspects associated with internet advertising auctions, and focused on how to manage the multi-item aspect via position auctions. This lecture note studies a second major practical aspect of internet advertising auctions: budgets. In these auctions, a large fraction of advertisers specify a budget constraint that must hold in aggregate across all of the payments made by the advertiser. At an intuitive level, these budget constraints bind together all the auctions, and thus we need to worry about strategic aspects that may not occur in the analysis of each individual auction in isolation. Most prominently, the budget constraints break the very nice property of second-price auctions which says that each buyer should simply bid their true value.

2 Auctions Markets

Throughout the rest of this lecture note, we will consider settings where each individual auction is a single-item auction, using either first or second-price rules. This itself is also a simplification: in practice each individual auction would be more complicated, but even just for single-item individual auctions it turns out that there are a lot of interesting problems.

In this setting we have n buyers and m goods. Buyer i has value v_{ij} for good j , and each buyer has some budget B_i . Each good j will be sold via sealed-bid auction, using either first or second-price. We assume that for all buyers i , there exists some item j such that $v_{ij} > 0$, and similarly for all j there exists i such that $v_{ij} > 0$. Let $x \in \mathbb{R}^{n \times m}$ be an allocation of items to buyers, with associated prices $p \in \mathbb{R}^m$. The utility that a buyer i derives from this allocation is

$$u_i(x_i, p) = \begin{cases} \langle v_i, x_i \rangle - \langle p, x_i \rangle & \text{if } \langle p, x_i \rangle \leq B_i \\ -\infty & \text{otherwise} \end{cases}.$$

We call this setting an *auction market*. If SP (FP) auctions are used then we call it an SP (FP) auction market.

3 Second-Price Auction Markets

In the previous lecture we saw that the second-price auction is strategyproof. However, this relied on there being a single auction, and no budgets. It's easy to construct an example showing that

*Department of Industrial Engineering and Operations Research, Columbia University. Email: christian.kroer@columbia.edu.




Figure 1: Comparison of pacing methods. Left: no pacing, middle: probabilistic pacing, right: multiplicative pacing.

this is no longer true in SP auction markets. Consider a setting with two buyers and two items, with valuations $v_1 = (100, 100)$, $v_2 = (1, 1)$ and budgets $B_1 = B_2 = 1$. If both buyers submit their true valuations then buyer 1 wins both items, pays 2, and gets $-\infty$ utility.

Instead, each buyer needs to somehow smooth out their spending across auctions. For large-scale Internet auctions this is typically achieved via some sort of *pacing rule*. Here we will mention two that have been used in practice:

1. *Probabilistic pacing*: each buyer i is given a parameter $\alpha_i \in [0, 1]$ denoting the probability that they should participate in each auction. For each auction j , an independent coin is flipped which comes up heads with probability α_i , and if it comes up heads then the buyer submits a bid $b_{ij} = v_{ij}$ to that auction.
2. *Multiplicative pacing*: each buyer i is given a parameter $\alpha_i \in [0, 1]$, which acts as a scalar multiplier on their truthful bids. In particular, for each auction j , buyer i submits a bid $b_{ij} = \alpha_i v_{ij}$.

Both methods have been applied in real-life large-scale Internet ad markets.

Figure 1 shows a comparison of pacing methods for a simplified setting where time is taken into account. Here we assume that we are considering some buyer i whose value is the same for every item, but other bidders are causing the items to have different prices. On the x-axis we plot time, and on the y-axis we plot the price of each item. On the left is the outcome from naive bidding: the buyer spends their budget much too fast, and ends up running out of budget when there are many high-value items left for them to buy. In practice, many buyers also prefer to smoothly spend their budget throughout the day. In the middle we show probabilistic pacing, where we do get smooth budget expenditure. However, the buyer ends up buying some very expensive item, while missing out on much cheaper items that have the same value to them. Finally, on the right is the result from probabilistic pacing, where the buyer picks an optimal threshold to buy at, and thus buys item optimally in order of bang-per-buck.

In this note we will focus on multiplicative pacing, but see the historical notes section for some references to papers that also consider probabilistic pacing.

The intuition given in Figure 1 can be shown to hold more generally when items have different values to the buyer. Generally, it turns out that given a set of bids by all the other bidders, a buyer can always specify a best response by choosing an optimal pacing multiplier:

Proposition 1. *Suppose we allow arbitrary bids in each auction. If we hold all bids for buyers $k \neq i$ fixed, then buyer i has a best response that consists of multiplicatively-paced bids (assuming that if a buyer is tied for winning an auction, they can specify the fraction that they win).*

Proof. Since every other bid is held fixed, we can think of each item as having some price $p_j = \max_{k \neq i} b_{kj}$, which is what i would pay if they bid $b_{ij} \geq b_{kj}$. Now we may sort the items in decreasing

order of bang-per-buck $\frac{v_{ij}}{p_j}$. An optimal allocation for i clearly consists of buying items in this order, until they reach some index j such that if they buy every item with index $l < j$ and some fraction x_{ij} of item j , they either spend their whole budget, or j is the first item with $\frac{v_{ij}}{p_j} \geq 1$ (if $\frac{v_{ij}}{p_j} > 1$ then $x_{ij} = 0$). Now set $\alpha_i = \frac{p_j}{v_{ij}}$. With this bid, i gets exactly this optimal allocation: for all items $l \leq j$ (which are the items in the optimal allocation), we have $\alpha_i v_{il} = \frac{p_j}{v_{ij}} v_{il} \geq \frac{p_l}{v_{il}} v_{il} = p_l$. \square

The goal will be to find a *pacing equilibrium*:

Definition 1. A *second-price pacing equilibrium (SPPE)* is a vector of pacing multipliers $\alpha \in [0, 1]^n$, a fractional allocation x_{ij} , and a price vector such that for every buyer i :

- For all j , $\sum_i x_{ij} = 1$, and if $x_{ij} > 0$ then i is tied for highest bid on item j .
- If $x_{ij} > 0$ then $p_j = \max_{k \neq i} \alpha_k v_{kj}$.
- For all i , $\sum_j p_j x_{ij} \leq B_i$. Additionally, if the inequality is strict then $\alpha_i = 1$.

The first and second conditions of pacing equilibrium simply enforce that the item always goes to winning bids at the second-price rule. The third condition ensures that a buyer is only paced if their budget constraint is binding. It follows (almost) immediately from Proposition 1 that every buyer is best responding in SPPE.

A nice property of SPPE is that it is always guaranteed to exist (this is not immediate from the existence of, say, a Nash equilibrium in a standard game, since an SPPE corresponds to a specific type of pure-strategy Nash equilibrium):

Theorem 1. An SPPE of a pacing game is always guaranteed to exist.

We won't cover the whole proof here, but we will state the main ingredients, which are useful to know more generally.

- First, a smoothed pacing game is constructed. In the smoothed game, the allocation is smoothed out among all bids that are within ϵ of the maximum bid, thus making the allocation a deterministic function of the pacing multipliers α . Several other smooth approximations are also introduced to deal with other discontinuities. In the end, a game is obtained, where each player simply has as their action space the interval $[0, 1]$ and utilities are nice continuous and quasi-concave functions.
- Secondly, the following fixed-point theorem is invoked to guarantee existence of a pure-strategy Nash equilibrium in the smoothed game.

Theorem 2. Consider a game with n players, strategy space A_i , and utility function $u_i(a_i, a_{-i})$. If the following conditions are satisfied:

- A_i is convex and compact for all i
- $u_i(s_i, \cdot)$ is continuous in s_{-i}
- $u_i(\cdot, s_{-i})$ is continuous and quasi-concave in s_i (quasi-concavity of a function $f(x)$ means that for all x, y and $\lambda \in [0, 1]$ it holds that $f(\lambda x + (1 - \lambda)y) \geq \min(f(x), f(y))$)

then a pure-strategy Nash equilibrium exists.

- Finally, the limit point of smoothed games as the smoothing factor ϵ tends to zero is shown to yield an equilibrium in the original pacing problem.

i	v_{i1}	v_{i2}	v_{i3}	v_{i4}	B_i
1	100	1	99	100	1
2	1	100	99		1
3				100	100

α_i	b_{i1}	b_{i2}	b_{i3}	b_{i4}	spend
1	100	1	99	100	1
0.01	0.01	1	0.99		1
1				100	100

α_i	b_{i1}	b_{i2}	b_{i3}	b_{i4}	spend
0.01	1	0.01	0.99	1	1
1	1	100	99		1
1				100	1

Figure 2: Multiplicity of SPPE. On the left is shown a problem instance, and on the right is shown two possible second-price pacing equilibria.

Unfortunately, while SPPE is guaranteed to exist, it turns out that sometimes there are several SPPE, and they can have large differences in revenue, social welfare, and so on. An example is shown in Figure 2. In practice this means that we might need to worry about whether we are in a good and fair equilibrium.

Another positive property of SPPE is that every SPPE is also a market equilibrium, if we consider a market equilibrium setting where each buyer has a quasi-linear demand function that respects the total supply as follows:

$$D_i(p) = \operatorname{argmax}_{0 \leq x_i \leq 1} \langle v_i - p, x_i \rangle \text{ s.t. } \langle p, x_i \rangle \leq B_i.$$

This follows immediately by simply using the allocation x and prices p from the SPPE as a market equilibrium. Proposition 1 tells us that $x_i \in D_i(p)$, and the market clears by definition of SPPE. This means that SPPE has a number of nice properties such as no envy and Pareto optimality (although Pareto optimality requires considering the seller as an agent too).

Finally we turn to the question of computing an SPPE. Unfortunately the news there is bad. It was shown recently that computing an SPPE is a PPAD-complete problem. This means that there exists a polynomial-time reduction between the problem of computing a Nash equilibrium in a general-sum game and that of computing an SPPE, and thus the two problems are equally hard, from the perspective of computing a solution in polynomial time. Moreover, it was also shown that we cannot hope for iterative methods to efficiently compute an approximate SPPE. Beyond merely computing *any* SPPE, we could also try to find one that maximizes revenue or social welfare. This problem turns out to be NP complete.

There is a mixed-integer program for computing SPPE, but unfortunately it is not very scalable.

4 First-Price Auction Markets

Next we consider what happens if we instead sell each item by first-price auction as part of an auction market.

First we start by defining what we call *budget-feasible pacing multipliers*. Intuitive, this is simply a set of pacing multipliers such that everything is allocated according to first-price auction, and everybody is within budget.

Definition 2. A set of budget-feasible pacing multipliers (BFPM) is a vector of pacing multipliers $\alpha \in [0, 1]^n$ and a fractional allocation x_{ij} such that for every buyer i :

- Prices are defined to be $p_j = \max_k \alpha_i v_{kj}$.
- For all j , $\sum_i x_{ij} = 1$, and if $x_{ij} > 0$ then i is tied for highest bid on item j .
- For all i , $\sum_j p_j x_{ij} \leq B_i$.

Again, the goal will be to find a *pacing equilibrium*. This is simply a BFPM that satisfied the complementarity condition on the budget constraint and pacing multiplier.

Definition 3. A first-price pacing equilibrium (FPPE) is a BFPM (α, x) such that for every buyer i :

- For all i , if $\sum_j p_j x_{ij} < B_i$ then $\alpha_i = 1$.

Notably, the only difference to SPPE is the pricing condition, which now uses first price.

A very nice property of the first-price setting is that BFPMs satisfy a monotonicity condition: if (α', x') and (α'', x'') are both BFPM, then the pacing vector $\alpha = \max(\alpha', \alpha'')$ (where the max is taken componentwise) is also a BFPM. The associated allocation is that for each item j , we first identify whether the highest bid comes from α' or α'' , and use the corresponding allocation of j (breaking ties towards α').

Intuitively, the reason that (α, x) is also BFPM is that for every buyer i , their bids are the same as in one of the two previous BFPMs (say (α', x') WLOG.), and so the prices they pay are the same as in (α', x') . Furthermore, since every other buyer is bidding at least as much as in (α', x') , they win weakly less of each item (using the tie-breaking scheme described above). Since (α', x') satisfied budgets, (α, x) must also satisfy budgets. The remaining conditions are easily checked.

In addition to componentwise maximality, there is also a *maximal* BFPM (α, x) (there could be multiple x compatible with α) such that $\alpha \geq \alpha'$ for all α' that are part of any BFPM. Consider $\alpha_i^* = \sup\{\alpha_i \mid \alpha \text{ is part of a BFPM}\}$. For any ϵ and i , we know that there must exist a BFPM such that $\alpha_i > \alpha_i^* - \epsilon$. For a fixed ϵ we can take componentwise maxima to conclude that there exists $(\alpha^\epsilon, x^\epsilon)$ that is a BFPM. This yields a sequence $\{(\alpha^\epsilon, x^\epsilon)\}$ as $\epsilon \rightarrow 0$. Since the space of both α and x is compact, the sequence has a limit point (α^*, x^*) . By continuity (α^*, x^*) is a BFPM.

We can use this maximality to show existence and uniqueness (of multipliers) of FPPE:

Theorem 3. An FPPE always exists and the set of pacing multipliers $\{\alpha\}$ that are part of an FPPE is a singleton.

Proof. Here we give a high-level proof, a more explicit proof can be found in the paper listed in the notes.

Consider the component-wise maximal α and an associated allocation x such that they form a BFPM.

Since α, x is a BFPM, we only need to check that it has no unnecessarily paced bidders. Suppose some buyer i is spending strictly less than B_i and $\alpha_i < 1$. If i is not tied for any items, then we can increase α_i for some sufficiently small ϵ and retain budget feasibility, contradicting the maximality of α . If i is tied for some item, consider the set $N(i)$ of all bidders tied with i . Now take the transitive closure of this set by repeatedly adding any bidder that is tied with any bidder in $N(i)$. We can now redistribute all the tied items among bidders in $N(i)$ such that no bidder in $N(i)$ is budget constrained (this can be done by slightly increasing i 's share of every item they are tied on, then slightly increasing the share of every other buyer in $N(i)$ who is now below budget, and

so on). But now there must exist some small enough $\delta > 0$ such that we can increase the pacing multiplier of every bidder in $N(i)$ by δ while retaining budget feasibility and creating no new ties. This contradicts α being maximal. We get that there can be no unnecessarily paced bidders under α .

Finally, to show uniqueness, consider any alternative BFPM α', x' . Consider the set I of buyers such that $\alpha'_i < \alpha$; Since $\alpha \geq \alpha'$ and $\alpha \neq \alpha'$ this set must have size at least one. Since all buyers in I were spending less than their budget under α , and their collective spending strictly decreased, at least one buyer in I must not be spending their whole budget. But $\alpha'_i < \alpha_i \leq 1$ for all $i \in I$, so that buyer must be unnecessarily paced. \square

4.1 Sensitivity

FPPE enjoys several nice monotonicity and sensitivity properties that SPPE does not. Several of these follow from the maximality property of FPPE: the unique FPPE multipliers α are such that $\alpha \geq \alpha'$ for any other BFPM (α', x') .

The following are all guaranteed to weakly increase revenue of the FPPE:

1. Adding a bidder i : the old FPPE (α, x) is still BFPM by setting $\alpha_i = 0, x_i = 0$. By α monotonicity prices increase weakly.
2. Adding an item: The new FPPE α' satisfies $\alpha' \leq \alpha$ (for contradiction, consider the set of bidders whose multipliers increased, since they win weakly more and prices went up, somebody must break their budget). Now consider the bidders such that $\alpha'_i < \alpha_i$. Those bidders spend their whole budget by the FPPE “no unnecessary pacing” condition. For bidders such that $\alpha'_i = \alpha_i$, they pay the same as before, and win weakly more.
3. Increasing a bidder i ’s budget: the old FPPE (α, x) is still BFPM, so this follows by α maximality.

It is also possible to show that revenue enjoys a Lipschitz property: increasing a single buyer’s budget by Δ increases revenue by at most Δ . Similarly, social welfare can be bounded in terms of Δ , though multiplicatively, and it does not satisfy monotonicity.

4.2 Convex Program

Next we consider how to compute an FPPE. This turns out to be easier than for SPPE. This is due to a direct relationship between FPPE and market equilibrium: FPPE solutions are exactly the set of solutions to the *quasi-linear* variant of the Eisenberg-Gale convex program for computing a market equilibrium:

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

$$\begin{aligned}
 \min_{p \geq 0, \beta \geq 0} \quad & \sum_j p_j - \sum_i B_i \log(\beta_i) \\
 \forall i, p_j \geq \quad & v_{ij} \beta_i \\
 \beta_i \leq \quad & 1
 \end{aligned} \tag{3}$$

On the left is shown the primal convex program, and on the right is shown the dual convex program. The variables x_{ij} denote the amount of item j that bidder i wins. The leftover budget is denoted by δ_i , it arises from the dual program: it is the primal variable for the dual constraint $\beta_i \leq 1$, which constrains bidder i to paying at most a price-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: $\min_{j.s.t. x_{ij} > 0} \frac{p_j}{v_{ij}}$ for buyer i , and p_j is the price of good j .

We may use the following basic fact from convex optimization to conclude that strong duality holds and get optimality conditions:

Theorem 4. *Consider a convex program and its dual*

$$\begin{aligned} \min_x \quad & f(x) \\ g_i(x) \leq 0, \quad & \forall i \\ x \geq 0 \end{aligned} \quad (4) \qquad \begin{aligned} \max_{\lambda \geq 0} \quad & q(\lambda) \\ q(\lambda) := \min_{x \geq 0} \quad & L(x, \lambda) \\ L(x, \lambda) := f(x) + \sum_i \lambda_i g_i(x) \end{aligned} \quad (5)$$

with Lagrange multipliers λ_i for each constraint i . Assume that the following Slater constraint qualification is satisfied: there exists some $x \geq 0$ such that $g_i(x) < 0$ for all i . If (4) has a finite optimal value f^* then (5) has a finite optimal value q^* and $f^* = q^*$. Furthermore a solution pair x^*, λ^* is optimal if and only if the following Karush-Kuhn-Tucker (KKT) conditions hold:

- (primal feasibility) x^* is a feasible solution of (4)
- (dual feasibility) $\lambda^* \geq 0$
- (complementary slackness) $\lambda_i^* g_i(x^*) = 0$ for all i
- (stationarity) $x^* \in \operatorname{argmin}_{x \geq 0} L(x, \lambda^*)$

We can use the strong duality theorem above, and in particular the KKT conditions, to show that FPPE and EG are equivalent.

Informally, the correspondence between FPPE and solutions to the convex program follows because β_i specifies a single price-per-utility rate per bidder which exactly yields the pacing multiplier $\alpha_i = \beta_i$. Complementary slackness then guarantees that if $p_j > v_{ij}\beta_i$ then $x_{ij} = 0$, so any item allocated to i has exactly rate β_i . Similarly, complementary slackness on $\beta_i \leq 1$ and the associated primal variable δ_i guarantees that bidder i is only paced if they spend their whole budget.

Theorem 5. *An optimal solution to the quasi-linear Eisenberg-Gale convex program corresponds to an FPPE with pacing multiplier $\alpha_i = \beta_i$ and allocation x_{ij} , and vice versa.*

Proof. Clearly the quasi-linear Eisenberg-Gale convex program satisfies the Slater constraint qualification: we may use the proportional allocation where every buyers gets $\frac{1}{n}$ of every item to see this. Thus the optimal solution must satisfy the following 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: the primal program has the exact packing constraints. Budgets are also satisfied (here we may assume $u_i > 0$ since otherwise budgets are

satisfied since the bidder wins no items): by KKT condition 1 and KKT condition 7 we have that for any item j that bidder i is allocated part of:

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

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

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

This part of the budget argument is exactly the same as for the standard Eisenberg-Gale proof [12]. Note that (1) always holds exactly since the objective is strictly increasing in u_i . Thus $\delta_i = 0$ denotes full budget expenditure. If $\delta_i > 0$ then (1) implies that $u_i > \sum_j v_{ij} x_{ij}$ which gives:

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

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

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

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

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

It follows that an FPPE can be computed in polynomial time, and that we can apply various first-order methods to compute large-scale FPPE.

5 Conclusion

There are interesting differences in the properties satisfied by SPPE and FPPE. We summarize them quickly here (these are all covered in the literature noted in the Historical Notes):

- FPPE is unique (can be shown from the convex program, or directly from the monotonicity property of BFPM), SPPE is not
- FPPE can be computed in polynomial time, computing an SPPE is a PPAD-complete problem

- FPPE is less sensitive to perturbation (e.g. revenue increases smoothly as budgets are increased)
- SPPE corresponds to a pure-strategy Nash equilibrium, and thus buyers are best responding to each other
- Both correspond to different market equilibria (but SPPE requires buyer demands to be “supply aware”)
- Neither of them are strategyproof
- Due to the market equilibrium connection, both can be shown strategyproof in an appropriate “large market” sense

FPPE and SPPE have also been studied experimentally, both via random instances, as well as instances generated from real ad auction data. The most interesting takeaways from those experiments are:

- In practice SPPE multiplicity seems to be very rare
- Manipulation is hard in both SPPE and FPPE if you can only lie about your value-per-click
- FPPE dominates SPPE on revenue
- Social welfare can be higher in either FPPE or SPPE. Experimentally it seems to largely be a toss-up on which solution concept has higher social welfare.

6 Historical Notes

The multiplicative pacing equilibrium results shown in this lecture note were developed by Conitzer et al. [7] for SP auction markets, and Conitzer et al. [8] for FP auction markets. Another strand of literature has studied models where items arrive stochastically and valuations are then drawn independently. Balseiro et al. [1] show existence of pacing equilibrium for multiplicative pacing as well as several other pacing rules for such a setting; they also give a very interesting comparison of revenue and social welfare properties of the various pacing option in the unique symmetric equilibrium of their setting. Most notably, multiplicative pacing achieves strong social welfare properties, while probabilistic pacing achieves higher revenue properties. Balseiro, Besbes, and Weintraub [2] show that when bidders get to select their bids individually, multiplicative pacing equilibrium arises naturally via Lagrangian duality on the budget constraint, under a fluid-based mean-field market model. The PPAD-completeness of computing an SPPE was given by Chen, Kroer, and Kumar [5]

The quasi-linear variant of Eisenberg-Gale was given by Chen, Ye, and Zhang [4] and independently by Cole et al. [6] (an unpublished note from one of the authors in Cole et al. [6] was in existence around a decade before the publication of Cole et al. [6]). Theorem 4 is a specialization to the FPPE setting. In reality much stronger statements can be made: For a more general statement of the strong duality theorem and KKT conditions used here, see Bertsekas, Nedic, and Ozdaglar [3] Proposition 6.4.4. The KKT conditions can be significantly generalized beyond convex programming.

The fixed-point theorem that is invoked to guarantee existence of a pure-strategy Nash equilibrium in the smoothed game is by Debreu [9], Glicksberg [11], and Fan [10].

References

- [1] Santiago Balseiro, Anthony Kim, Mohammad Mahdian, and Vahab Mirrokni. Budget management strategies in repeated auctions. In *Proceedings of the 26th International Conference on World Wide Web*, pages 15–23, 2017.
- [2] Santiago R Balseiro, Omar Besbes, and Gabriel Y Weintraub. Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Science*, 61(4):864–884, 2015.
- [3] Dimitri P Bertsekas, A Nedic, and A Ozdaglar. Convex analysis and optimization. 2003. *Athena Scientific*, 2003.
- [4] Lihua Chen, Yinyu Ye, and Jiawei Zhang. A note on equilibrium pricing as convex optimization. In *International Workshop on Web and Internet Economics*, pages 7–16. Springer, 2007.
- [5] Xi Chen, Christian Kroer, and Rachitesh Kumar. The complexity of pacing for second-price auctions. In *Proceedings of the 2021 ACM Conference on Economics and Computation*, 2021.
- [6] Richard Cole, Nikhil R Devanur, Vasilis Gkatzelis, Kamal Jain, Tung Mai, Vijay V Vazirani, and Sadra Yazdanbod. Convex program duality, fisher markets, and Nash social welfare. In *18th ACM Conference on Economics and Computation, EC 2017*. Association for Computing Machinery, Inc, 2017.
- [7] Vincent Conitzer, Christian Kroer, Eric Sodomka, and Nicolás E Stier-Moses. Multiplicative pacing equilibria in auction markets. In *International Conference on Web and Internet Economics*, 2018.
- [8] Vincent Conitzer, Christian Kroer, Debmalya Panigrahi, Okke Schrijvers, Eric Sodomka, Nicolas E Stier-Moses, and Chris Wilkens. Pacing equilibrium in first-price auction markets. In *Proceedings of the 2019 ACM Conference on Economics and Computation*. ACM, 2019.
- [9] Gerard Debreu. A social equilibrium existence theorem. *Proceedings of the National Academy of Sciences*, 38(10):886–893, 1952.
- [10] Ky Fan. Fixed-point and minimax theorems in locally convex topological linear spaces. *Proceedings of the National Academy of Sciences of the United States of America*, 38(2):121, 1952.
- [11] Irving L Glicksberg. A further generalization of the kakutani fixed point theorem, with application to nash equilibrium points. *Proceedings of the American Mathematical Society*, 3(1): 170–174, 1952.
- [12] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani. *Algorithmic game theory*. Cambridge University Press, 2007.