

A Note on the Efficiency of Position Mechanisms with Budget Constraints*

Alexandros A. Voudouris

Department of Computer Science, University of Oxford, UK

Abstract

We study the social efficiency of several well-known mechanisms for the allocation of a set of available (advertising) positions to a set of competing budget-constrained users (advertisers). Specifically, we focus on the Generalized Second Price auction (GSP), the Vickrey-Clarke-Groves mechanism (VCG) and the Expressive Generalized First Price auction (EGFP). Using liquid welfare as our efficiency benchmark, we prove a tight bound of 2 on the liquid price of anarchy and stability of these mechanisms for pure Nash equilibria.

Keywords: position mechanisms, liquid welfare, budget constraints, liquid price of anarchy

1. Introduction

Position mechanisms have been extensively used for the allocation of advertising positions (with different click-through rates) when keywords are queried in search engines. Such mechanisms auction off the available positions to the interested advertisers, who in turn compete with each other by submitting bids, expressing how much they value the available advertising positions (per user click).

There have been numerous papers analyzing the properties of position mechanisms. Edelman et al. [11] (see also [18]) studied the *generalized first price auction* (GFP) as well as the *generalized second price auction* (GSP). According to these mechanisms, the advertisers are sorted in terms of the scalar bids that they submit, and each of them pays her own bid or the next highest bid, respectively. These definitions allow the advertisers to strategize over their bids and engage as players into a strategic game. Edelman et al. [11] proved that the games induced by GFP are not guaranteed to have pure Nash equilibria, while the games induced by GSP always have socially efficient pure Nash equilibria with respect to the *social welfare* benchmark (the total value of the players for the positions they are given); consequently, the *price of stability* [1] of GSP is equal to 1. Caragiannis et al. [3] (see also [17]) focused on worst-case equilibria and proved several bounds on the *price of anarchy* [13] of GSP with respect to a variety of equilibrium concepts, ranging from pure Nash and coarse-correlated equilibria in the full information setting to Bayes-Nash equilibria in the incomplete information setting. Dütting et al. [10] proved bounds on the revenue and exploited more expressive input formats as a remedy for the non-existence of pure Nash equilibria in games induced by GFP.

*This work has been partially supported by a PhD scholarship from the Onassis Foundation, and by the European Research Council (ERC) under grant number 639945 (ACCORD).

Email address: alexandros.voudouris@cs.ox.ac.uk (Alexandros A. Voudouris)

They designed the *expressive generalized first price auction* (EGFP) according to which each player submits a bid per position, the positions are auctioned off sequentially, and each player pays her bid for the position she is given.

All of the aforementioned papers studied the *no-budget* setting, where the players are assumed to be able to afford any payments, no matter how large these can get. However, in reality, the players have hard budget constraints that upper-bound the payments that they can afford. Following a series of recent papers that focus on such budget-constrained settings, we also study the social efficiency of position mechanisms by bounding the (pure) price of anarchy and stability in terms of the *liquid welfare* benchmark that takes budgets into account. Liquid welfare was first introduced by Dobzinski and Paes Leme [9] who focused on the design of truthful mechanisms for the allocation of multiple units of a single divisible item (see also [14, 15] for extensions of this setting). One of the very first results of Dobzinski and Paes Leme [9] is the observation that the celebrated VCG mechanism [19, 7, 12] is no longer truthful, which is the case for VCG in our setting as well.

The liquid price of anarchy has been considered in a few related papers so far. Syrgkanis and Tardos [17] considered the liquid welfare benchmark under the term *effective welfare* and bounded the ratio between the optimal liquid welfare and the worst-case social welfare at equilibrium, in various strategic auction settings, including position mechanisms. Caragiannis and Voudouris [4] and Christodoulou et al. [6] were the first to provide constant bounds on the liquid price of anarchy (ratio of optimal liquid welfare over worst-case liquid welfare at equilibrium) of the proportional mechanism for the allocation of divisible resources. These results are based on the now-standard unilateral deviations technique (see also [16]) and can be extended to more general equilibrium concepts, given a specific definition of the liquid welfare for randomized allocations. Our upper bounds follow this technique as well, but it seems non-trivial to extend them to more general equilibrium concepts due to the particular form of the deviating bids used. For pure equilibria in particular, by exploiting the structure of worst-case equilibria, Caragiannis and Voudouris [5] were able to characterize the liquid price of anarchy of all mechanisms for the allocation of a single divisible resource, leading to tight bounds. Finally, Azar et al. [2] refined the definition of the liquid welfare for randomized allocations and proved constant liquid price of anarchy bounds over general equilibrium concepts for simultaneous first price auctions.

In Section 2, we formally describe the setting considered in this paper and the mechanisms that we are interested in. Then, in Section 3 we prove our main result: the liquid price of anarchy and stability of GSP, VCG and EGFP is exactly 2. Consequently, in contrast to the no-budget setting, when we consider players with budget constraints and the liquid welfare benchmark, these mechanisms do not have socially efficient equilibria. Such a phenomenon was first observed by Caragiannis and Voudouris [5] for all single divisible resource allocation mechanisms, and it might be the case that this holds for any position mechanism as well. We conclude with a short discussion of possible extensions of our work in Section 4.

2. Preliminaries

There are n available *positions* such that position j has associated click-through-rate (CTR) $\alpha_j \in \mathbb{R}_{>0}$ such that $\alpha_j \geq \alpha_{j+1}$ for $j \in [n-1]$; let $\boldsymbol{\alpha} = (\alpha_j)_{j \in [n]}$ be the vector containing the CTRs of all positions. Furthermore, there are n players that compete over these positions. Player i has a valuation v_i and a total private budget c_i ; let $\mathbf{v} = (v_i)_{i \in [n]}$ and $\mathbf{c} = (c_i)_{i \in [n]}$ be the vectors containing the valuations and budgets of all players. The valuation v_i indicates the value that player i has per click and, therefore, if player i is assigned to some position j , then her total value is $\alpha_j v_i$. The budget c_i can be thought of as an upper bound to the payment that the player can afford in order to buy some position.

We consider several greedy mechanisms for the allocation of positions to players, which generally work as follows. Let M be a greedy mechanism. Each player i submits a bid b_i that can either be a real non-negative scalar or a vector of scalars per position, depending on the input format that M requires; let \mathbf{b} be the vector (or matrix) of bids submitted by all players. Then, the players are sorted in non-increasing order in terms of their bids and the induced ranking $\sigma(\mathbf{b})$ indicates the position $\sigma_i(\mathbf{b})$ that is assigned to each player i ; therefore, we call $\sigma(\mathbf{b})$ an *assignment* that is induced by \mathbf{b} . Also, let $\pi_j(\mathbf{b})$ denote the player that is assigned to position j such that $\pi_{\sigma_i(\mathbf{b})}(\mathbf{b}) = i$. The mechanism charges player i an amount of money $p_i(\mathbf{b}, M)$ that depends on \mathbf{b} , and may or may not depend on $\alpha_{\sigma_i(\mathbf{b})}$. Given a bid vector \mathbf{b} , each player i has utility $u_i(\mathbf{b}, M) = \alpha_{\sigma_i(\mathbf{b})} v_i - p_i(\mathbf{b}, M)$ if $p_i(\mathbf{b}, M) \leq c_i$, and $-\infty$ otherwise. We focus on three important greedy allocation mechanisms that function as follows.

Generalized Second Price (GSP). Each player i submits a scalar $b_i \in \mathbb{R}_{\geq 0}$. The players are sorted in non-increasing order in terms of these bids and are assigned to the corresponding positions. Each player i is charged the next highest bid per click, that is, the bid of player $\pi_{\sigma_i(\mathbf{b})+1}(\mathbf{b})$ who is assigned to the next position $\sigma_i(\mathbf{b}) + 1$. Hence, the payment of player i is $p_i(\mathbf{b}, \text{GSP}) = \alpha_{\sigma_i(\mathbf{b})} b_{\pi_{\sigma_i(\mathbf{b})+1}(\mathbf{b})}$, and her utility can be written as $u_i(\mathbf{b}, \text{GSP}) = \alpha_{\sigma_i(\mathbf{b})} (v_i - b_{\pi_{\sigma_i(\mathbf{b})+1}(\mathbf{b})})$.¹

Vickrey-Clarke-Groves (VCG). Again, each player i submits a scalar $b_i \in \mathbb{R}_{\geq 0}$, and the players are sorted in non-increasing order in terms of their bids. Each player i is charged the difference between the social welfare (based on the bids) of the players ranked below i if i did not participate and their actual social welfare when i participates. In other words, the payment of player i is $p_i(\mathbf{b}, \text{VCG}) = \sum_{j=\sigma_i(\mathbf{b})+1}^n b_{\pi(j)} (\alpha_{j-1} - \alpha_j)$, and her utility can be written as $u_i(\mathbf{b}, \text{VCG}) = \alpha_{\sigma_i(\mathbf{b})} \left(v_i - \frac{1}{\alpha_{\sigma_i(\mathbf{b})}} \sum_{j=\sigma_i(\mathbf{b})+1}^n b_{\pi(j)} (\alpha_{j-1} - \alpha_j) \right)$.

Expressive Generalized First Price (EGFP). Each player i submits a vector $\mathbf{b}_i \in \mathbb{R}_{\geq 0}^n$ containing a bid per position. The positions are assigned to the players sequentially so that the next available position gets assigned to the player with the maximum bid for it, among the players that have not yet been allocated a position. In other words, let S_j

¹Interestingly, Díaz et al. [8] proved that GSP may not have any equilibria when the number of players exceeds the number of available positions and proposed alternative mechanisms; we here consider the same number of players and positions.

be the set of players that are competing for positions $\ell \geq j$; initially, S_1 contains all players. Then, $\pi_j(\mathbf{b}, \text{EGFP}) = \arg \max_{i \in S_j} b_{i,j}$. Each player i is charged (in total) her bid for the position that she is allocated, i.e., $p_i(\mathbf{b}, \text{EGFP}) = b_{i, \sigma_i(\mathbf{b})}$, and her utility is $u_i(\mathbf{b}, \text{EGFP}) = \alpha_{\sigma_i(\mathbf{b})} v_i - b_{i, \sigma_i(\mathbf{b})}$.

The game. Let $M \in \{\text{GSP}, \text{VCG}, \text{EGFP}\}$ be any of the aforementioned position mechanisms. Mechanism M induces a strategic position game $\mathcal{G}(M)$ among the players who act as utility maximizers; this is true even for VCG as we will see in the next section. A bid vector (or matrix) \mathbf{b} is called a *pure Nash equilibrium* (or, simply, equilibrium) for $\mathcal{G}(M)$ if all players simultaneously maximize their utilities and have no incentive to deviate to any different bid in order to increase their personal utility, i.e., $u_i(\mathbf{b}, M) \geq u_i((y, \mathbf{b}_{-i}), M)$, for all players i and bids $y \neq b_i$. Here, the notation (y, \mathbf{b}_{-i}) is used to denote the vector (or matrix) that is obtained by \mathbf{b} when player i bids y (and all other players bid according to \mathbf{b}). Let $\text{eq}(\mathcal{G}(M))$ be the set of all equilibria of the position game $\mathcal{G}(M)$.

Liquid welfare, price of anarchy and price of stability. We measure the social efficiency of an assignment $\sigma(\mathbf{b})$ by the *liquid welfare* benchmark, which is defined as the total value of the players, with the value of each player capped by her budget, i.e.,

$$\text{LW}(\sigma(\mathbf{b})) = \sum_i \min\{\alpha_{\sigma_i(\mathbf{b})} v_i, c_i\}.$$

The *liquid price of anarchy* (*liquid price of stability*) of a position game $\mathcal{G}(M)$ that is induced by a position mechanism M is defined as the ratio between the optimal liquid welfare achieved by any assignment to the minimum (maximum) liquid welfare achieved at any equilibrium assignment. In other words, the liquid price of anarchy and the liquid price of stability of $\mathcal{G}(M)$ are, respectively, equal to

$$\text{LPoA}(\mathcal{G}(M)) = \frac{\max_{\mathbf{y}} \text{LW}(\sigma(\mathbf{y}))}{\min_{\mathbf{b} \in \text{eq}(\mathcal{G}(M))} \text{LW}(\sigma(\mathbf{b}))}$$

and

$$\text{LPoS}(\mathcal{G}(M)) = \frac{\max_{\mathbf{y}} \text{LW}(\sigma(\mathbf{y}))}{\max_{\mathbf{b} \in \text{eq}(\mathcal{G}(M))} \text{LW}(\sigma(\mathbf{b}))}.$$

Then, the liquid price of anarchy and stability of a mechanism M are respectively defined as the worst-case liquid price of anarchy and stability among all position games that are induced by M , i.e., $\text{LPoA}(M) = \sup_{\mathcal{G}(M)} \text{LPoA}(\mathcal{G}(M))$ and $\text{LPoS}(M) = \sup_{\mathcal{G}(M)} \text{LPoS}(\mathcal{G}(M))$.

The no-over assumption: no-overbidding and no-overbudgeting. For the GSP and VCG mechanisms, in order to have meaningful bounds on the liquid price of anarchy, we assume that $\alpha_{\sigma_i(\mathbf{b})} b_i \leq \min\{\alpha_{\sigma_i(\mathbf{b})} v_i, c_i\}$ for every player i . This is a combination of the well-known no-overbidding assumption that demands that $b_i \leq v_i$ and a no-overbudgeting assumption that demands that $\alpha_{\sigma_i(\mathbf{b})} b_i \leq c_i$. This assumption is necessary as it is easy (like in the case of the classic price of anarchy literature that deals with the social welfare objective) to construct position games that have arbitrarily bad liquid price of anarchy when the players overbid. For the EGFP mechanism such an assumption is of course not necessary due to the definition of the payment function.

3. Bounds on the liquid price of anarchy and stability

We begin with Theorem 1, where we show that the LPoA and LPoS of GSP, VCG and EGFP are at least 2; notice that the example that we present in the following proof also proves that VCG is no longer truthful when the players have budget constraints. Then, in Theorem 2 we prove that this bound of 2 on the LPoA and LPoS is tight.

Theorem 1. *The liquid price of anarchy and stability of GSP, VCG (under the no-over assumption) and EGFP are at least 2.*

PROOF. Let $M \in \{\text{GSP}, \text{VCG}, \text{EGFP}\}$, $\lambda > 2$ and $\varepsilon \in (0, 1/2)$. Consider a position game $\mathcal{G}(M)$ among two players with valuations $\mathbf{v} = (\lambda, 1)$ and budgets $\mathbf{c} = (1 + \varepsilon, 1)$, for two positions with CTRs $\boldsymbol{\alpha} = (1, 1/\lambda)$. Observe that, for the two possible assignments (1, 2) and (2, 1), the liquid welfare is $\text{LW}(1, 2) = \min\{\lambda, 1 + \varepsilon\} + \min\{\frac{1}{\lambda}, 1\} = \frac{(1+\varepsilon)\lambda+1}{\lambda}$ and $\text{LW}(2, 1) = \min\{1, 1 + \varepsilon\} + \min\{1, 1\} = 2$. Therefore, since $\lambda > 2$ and $\varepsilon < 1/2$, we have that $\text{LW}(2, 1) > \text{LW}(1, 2)$, and the optimal assignment is (2, 1). The ratio

$$\frac{\text{LW}(2, 1)}{\text{LW}(1, 2)} = \frac{2\lambda}{(1 + \varepsilon)\lambda + 1}$$

tends to 2 as λ becomes arbitrarily large and ε becomes arbitrarily small. In order to prove the theorem, it suffices to show that there exists an equilibrium bid vector that induces the assignment (1, 2), while there exists no equilibrium bid vector that induces the assignment (2, 1).

GSP. First, consider the bid vector $\mathbf{b} = (1 + \varepsilon, 1)$ which induces the assignment (1, 2). The utilities of the two players are $u_1(\mathbf{b}, \text{GSP}) = \lambda - 1$ and $u_2(\mathbf{b}, \text{GSP}) = \frac{1}{\lambda}$. Player 2 has no incentive to deviate as, by the no-over assumption, she cannot bid above her budget (which coincides with her value), while any other bid would not change the assignment. Player 1 obviously has no incentive to deviate to any other bid $b'_1 \geq b_2$ as the assignment as well as her payment would not change. So, consider the deviation of player 1 to the bid $b'_1 = b_2 - \gamma$, for some $\gamma > 0$. Then, the induced assignment would be (2, 1) and player 1 would have utility $u_1((b'_1, b_2), \text{GSP}) = 1 < u_1(\mathbf{b}, \text{GSP})$ since $\lambda > 2$. Therefore, \mathbf{b} is an equilibrium, and the price of anarchy bound follows.

Now, assume that there exists an equilibrium bid vector $\mathbf{b} = (b_1, b_2)$ with $b_1 \leq b_2 \leq 1$ so that the assignment (2, 1) is induced, while the no-over assumption is satisfied (for player 2). The utilities of the two players at this equilibrium are $u_1(\mathbf{b}, \text{GSP}) = 1$ and $u_2(\mathbf{b}, \text{GSP}) = 1 - b_1$. Consider the deviation of player 1 to the bid $b'_1 = c_1 = 1 + \varepsilon > b_2$. Then, the utility of this player would be $u_1((b'_1, b_2), \text{GSP}) = \lambda - b_2 \geq \lambda - 1 > 1$, since $b_2 \leq 1$ and $\lambda > 2$. Hence, player 1 has incentive to deviate to b'_1 and \mathbf{b} cannot be an equilibrium. The price of stability bound follows.

VCG. Like in the case of GSP, consider the bid vector $\mathbf{b} = (1 + \varepsilon, 1)$ which induces the assignment (1, 2). The payments of the players are $p_1(\mathbf{b}, \text{VCG}) = 1 - \frac{1}{\lambda}$ and $p_2(\mathbf{b}, \text{VCG}) = 0$, yielding utilities of $u_1(\mathbf{b}, \text{VCG}) = \lambda - 1 + \frac{1}{\lambda}$ and $u_2(\mathbf{b}, \text{VCG}) = \frac{1}{\lambda}$. Obviously, again player 2 has no incentive to deviate, while player 1 has no incentive to deviate to any bid $b'_1 \geq b_2$. So, consider the deviation of player 1 to the bid $b'_1 = b_2 - \gamma$, for some $\gamma > 0$. Then, the

induced assignment would be $(2, 1)$, the payment of player 1 would be $p_1((b'_1, b_2), \text{VCG}) = 0$ and her utility would be $u_1((b'_1, b_2), \text{VCG}) = 1 < \lambda - 1 + \frac{1}{\lambda} = u_1(\mathbf{b}, \text{VCG})$ since $(\lambda - 1)^2 > 0$, for any $\lambda > 2$. Therefore, \mathbf{b} is an equilibrium, and the price of anarchy bound follows.

Now, assume that there exists an equilibrium bid vector $\mathbf{b} = (b_1, b_2)$ with $b_1 \leq b_2 \leq 1$ so that the assignment $(2, 1)$ is induced, while the no-over assumption is satisfied (for player 2). The payments of the players at this equilibrium are $p_1(\mathbf{b}, \text{VCG}) = 0$ and $p_2(\mathbf{b}, \text{VCG}) = b_1(1 - \frac{1}{\lambda})$, yielding utilities of $u_1(\mathbf{b}, \text{VCG}) = 1$ and $u_2(\mathbf{b}, \text{VCG}) = 1 - b_1(1 - \frac{1}{\lambda})$. Consider the deviation of player 1 to the bid $b'_1 = c_1 = 1 + \varepsilon > b_2$. Then, the induced assignment would be $(1, 2)$, while the payment and the utility of player 1 would be $p_1((b'_1, b_2), \text{VCG}) = b_2(1 - \frac{1}{\lambda}) \leq 1 - \frac{1}{\lambda}$ and $u_1((b'_1, b_2), \text{VCG}) = \lambda - p_1(b'_1, b_2) \geq \lambda - 1 + \frac{1}{\lambda} > 1$, respectively; the last inequality follows since $(\lambda - 1)^2 > 0$, for any $\lambda > 2$. Hence, since player 1 has incentive to deviate to b'_1 , \mathbf{b} cannot be an equilibrium, and the price of stability bound follows.

EGFP. To show that there exists an equilibrium bid vector \mathbf{b} that induces the assignment $(1, 2)$, consider the bids $\mathbf{b}_1 = (1 + \delta, 0)$, where $\delta > 0$ is arbitrarily small, and $\mathbf{b}_2 = (1, 0)$ of the two players for the two available positions, respectively. Observe that after the allocation of the first position, the second one is given without any competition to the only remaining player. Therefore, at equilibrium, no player has any incentive to submit a bid that is greater than zero for the second position. Player 2 has no incentive to change her bid for the first position since she simply cannot bid any higher, while bidding any lower would not change the assignment. Player 1 has no incentive to deviate to any other bid $b'_{1,1} \geq b_{2,1}$ as the allocation and her payment would not change, and δ is assumed to be arbitrarily small. So, consider the deviating bid $b'_{1,1} < b_{2,1}$ which would change the assignment to $(2, 1)$ and the utility of player 1 would be $u'_1(((b'_{1,1}, 0), \mathbf{b}_2), \text{EGFP}) = 1 < \lambda - 1 = u_1(\mathbf{b}, \text{EGFP})$. Therefore, \mathbf{b} is indeed an equilibrium, and the price of anarchy bound follows.

For the price of stability bound, assume that there exists an equilibrium bid matrix \mathbf{b} such that $b_{1,1} \leq b_{2,1} \leq 1$ so that the allocation $(2, 1)$ is induced; again the two players must bid zero for the second position which is, basically, for free. The utilities of the two players at this equilibrium are $u_1(\mathbf{b}, \text{EGFP}) = 1$ and $u_2(\mathbf{b}) = 1 - b_{2,1}$. Consider the deviation of player 1 to the bid $b'_{1,1} = 1 + \delta > b_{2,1}$ for the first position, where $\delta > 0$ is arbitrarily small. Then, player 1 would be allocated the first position and her utility would be $u_1(((b'_{1,1}, 0), \mathbf{b}_2), \text{EGFP}) = \lambda - b_{1,1} = \lambda - 1 - \delta > 1$, since $\lambda > 2$ and δ is arbitrarily small. Hence, player 1 has no incentive to deviate to $\mathbf{b}'_1 = (b'_{1,1}, 0)$, \mathbf{b} cannot be an equilibrium, and the proof is complete. \square

The proof of the upper bounds exploits the well-known technique (for proving welfare guarantees in games) of deviating bids. However, it is more complicated since the selected deviating bids must be such that the payments of the players are within their budgets. In fact, this is the main barrier in proving LPoA bounds for more general equilibrium concepts, like Bayes-Nash equilibria in the incomplete information model, where the bids of the players are random variables.

Theorem 2. *The liquid price of anarchy and stability of GSP, VCG (under the no-over assumption) and EGFP are at most 2.*

PROOF. Let $M \in \{\text{GSP}, \text{VCG}, \text{EGFP}\}$ and consider any n -player position game $\mathcal{G}(M)$. Let v_i and c_i be the value and budget of player $i \in [n]$, and let α_j be the CTR of position $j \in [n]$. Let \mathbf{b} be an equilibrium bid vector that induces an assignment $\sigma(\mathbf{b})$; recall that $\pi_j(\mathbf{b})$ denotes the player that is assigned to position j . Moreover, let o_i denote the position given to player i at an optimal allocation, and $\text{OPT} = \sum_i \min\{\alpha_{o_i} v_i, c_i\}$.

Now, consider the following partition of the players: $A = \{i : \alpha_{\sigma_i(\mathbf{b})} v_i \leq c_i\}$. Then, for every player $i \notin A$, we have that $\min\{\alpha_{\sigma_i(\mathbf{b})} v_i, c_i\} = c_i \geq \min\{\alpha_{o_i} v_i, c_i\}$, and by summing over all such players, we obtain

$$\sum_{i \notin A} \min\{\alpha_{\sigma_i(\mathbf{b})} v_i, c_i\} \geq \sum_{i \notin A} \min\{\alpha_{o_i} v_i, c_i\}. \quad (1)$$

The rest of the proof is dedicated to showing that, for any player $i \in A$ and some $\gamma > 0$, it holds that

$$u_i(\mathbf{b}) \geq \min\{\alpha_{o_i} v_i, c_i\} - \min\{\alpha_{o_i} v_{\pi_{o_i}(\mathbf{b})}, c_{\pi_{o_i}(\mathbf{b})}\} - \gamma. \quad (2)$$

Then, since $\min\{\alpha_{\sigma_i(\mathbf{b})} v_i, c_i\} = \alpha_{\sigma_i(\mathbf{b})} v_i \geq u_i(\mathbf{b})$, by summing over all players $i \in A$, and by the fact that $|A| \leq n$, we obtain

$$\begin{aligned} \sum_{i \in A} \min\{\alpha_{\sigma_i(\mathbf{b})} v_i, c_i\} &\geq \sum_{i \in A} \min\{\alpha_{o_i} v_i, c_i\} - \sum_{i \in A} \min\{\alpha_{o_i} v_{\pi_{o_i}(\mathbf{b})}, c_{\pi_{o_i}(\mathbf{b})}\} - \gamma|A| \\ &\geq \sum_{i \in A} \min\{\alpha_{o_i} v_i, c_i\} - \text{LW}(\sigma(\mathbf{b})) - \gamma n. \end{aligned} \quad (3)$$

Hence, the theorem will follow by combining inequalities (1) and (3), and by choosing γ to be arbitrarily small, since we have that

$$\begin{aligned} \text{LW}(\sigma(\mathbf{b})) &= \sum_{i \notin A} \min\{\alpha_{\sigma_i(\mathbf{b})} v_i, c_i\} + \sum_{i \in A} \min\{\alpha_{\sigma_i(\mathbf{b})} v_i, c_i\} \\ &\geq \sum_{i \notin A} \min\{\alpha_{o_i} v_i, c_i\} + \sum_{i \in A} \min\{\alpha_{o_i} v_i, c_i\} - \text{LW}(\sigma(\mathbf{b})) - \gamma n \\ &\geq \text{OPT} - \text{LW}(\sigma(\mathbf{b})) - \gamma n. \end{aligned}$$

We now distinguish between cases depending on which mechanism is used. In the following, since the mechanism under consideration is clear from context, we drop it from our notation.

GSP. For any player $i \in A$ consider the deviating bid $b_{\pi_{o_i}(\mathbf{b})} + \tilde{\gamma}$, where $b_{\pi_{o_i}(\mathbf{b})}$ is the bid of the player that is given position o_i at equilibrium and $\tilde{\gamma} = \frac{\gamma}{\alpha_{o_i}}$ is such that $b_{\pi_{o_i}(\mathbf{b})} \leq b_{\pi_{o_i}(\mathbf{b})} + \tilde{\gamma} \leq b_{\pi_{o_i-1}(\mathbf{b})}$; notice that player i can choose such a $\tilde{\gamma}$ as she has full information about the bids of the other players, and there exists a tie-breaking assigning the position o_i to player i after the deviation (in case of equality). With this deviating bid, player i essentially plays only for her optimal position o_i , if she can afford to do so.

If the deviating bid $b_{\pi_{o_i}(\mathbf{b})} + \tilde{\gamma}$ satisfies the no-over assumption, then player i is guaranteed to be given position o_i in the new allocation and pay $b_{\pi_{o_i}(\mathbf{b})}$ per click. By the equilibrium condition, the fact that $\gamma > 0$, and since $\alpha_{o_i} b_{\pi_{o_i}(\mathbf{b})} \leq \min\{\alpha_{o_i} v_{\pi_{o_i}(\mathbf{b})}, c_{\pi_{o_i}(\mathbf{b})}\}$ (by the no-over assumption for player $\pi_{o_i}(\mathbf{b})$), we have that

$$u_i(\mathbf{b}) \geq u_i(b_{\pi_{o_i}(\mathbf{b})} + \tilde{\gamma}, \mathbf{b}_{-i}) = \alpha_{o_i}(v_i - b_{\pi_{o_i}(\mathbf{b})}) \geq \min\{\alpha_{o_i} v_i, c_i\} - \min\{\alpha_{o_i} v_{\pi_{o_i}(\mathbf{b})}, c_{\pi_{o_i}(\mathbf{b})}\} - \gamma.$$

If the deviating bid does not satisfy the no-over assumption, then we have that $b_{\pi_{o_i}(\mathbf{b})} + \tilde{\gamma} > v_i$ or $\alpha_{o_i}(b_{\pi_{o_i}(\mathbf{b})} + \tilde{\gamma}) > c_i$. Due to the no-over assumption for player $\pi_{o_i}(\mathbf{b})$, both of these inequalities imply that $\min\{\alpha_{o_i} v_{\pi_{o_i}(\mathbf{b})}, c_{\pi_{o_i}(\mathbf{b})}\} + \gamma > \min\{\alpha_{o_i} v_i, c_i\}$. Since player i has non-negative utility at equilibrium, we conclude that

$$u_i(\mathbf{b}) \geq 0 > \min\{\alpha_{o_i} v_i, c_i\} - \min\{\alpha_{o_i} v_{\pi_{o_i}(\mathbf{b})}, c_{\pi_{o_i}(\mathbf{b})}\} - \gamma,$$

as desired.

VCG. The proof is similar to that for GSP. The main difference here is that when the deviating bid $b_{\pi_{o_i}(\mathbf{b})} + \tilde{\gamma}$ of player $i \in A$ satisfies the no-over assumption, then player i is again guaranteed to be given position o_i , but now has to pay $\sum_{j=o_i+1}^n b_{\pi_j(\mathbf{b})}(\alpha_{j-1} - \alpha_j)$ in total. Observe that, since VCG is a greedy mechanism, at equilibrium we have that $b_{\pi_{o_i}(\mathbf{b})} \geq b_{\pi_j(\mathbf{b})}$ for every $j \in \{o_i + 1, \dots, n\}$. This implies that

$$\sum_{j=o_i+1}^n b_{\pi_j(\mathbf{b})}(\alpha_{j-1} - \alpha_j) \leq b_{\pi_{o_i}(\mathbf{b})} \sum_{j=o_i+1}^n (\alpha_{j-1} - \alpha_j) = b_{\pi_{o_i}(\mathbf{b})}(\alpha_{o_i} - \alpha_n) \leq \alpha_{o_i} b_{\pi_{o_i}(\mathbf{b})}.$$

Using this, we can follow the proof template for GSP and show the desired inequality.

EGFP. For any player $i \in A$ consider the deviating bid vector \mathbf{y} so that $y_{o_i} = b_{\pi_{o_i}(\mathbf{b}), o_i} + \gamma$ and $y_j = 0$ for any other position $j \neq o_i$. Again, player i plays only for her optimal position o_i , if she can afford to do so. If $y_{o_i} > c_i$, then since the utility of player i is non-negative at equilibrium, we obtain

$$u_i(\mathbf{b}) \geq 0 > c_i - b_{\pi_{o_i}(\mathbf{b}), o_i} - \gamma \geq \min\{\alpha_{o_i} v_i, c_i\} - \min\{\alpha_{o_i} v_{\pi_{o_i}(\mathbf{b})}, c_{\pi_{o_i}(\mathbf{b})}\} - \gamma,$$

where the last inequality follows by the fact that player $\pi_{o_i}(\mathbf{b})$ has non-negative utility at equilibrium and her payment is within her budget, which imply that $b_{\pi_{o_i}(\mathbf{b}), o_i} \leq \min\{\alpha_{o_i} v_{\pi_{o_i}(\mathbf{b})}, c_{\pi_{o_i}(\mathbf{b})}\}$.

Otherwise, the deviating bid is such that player i is allocated position o_i and her payment y_{o_i} is within her budget. Therefore, by the equilibrium condition, and by the fact that $b_{\pi_{o_i}(\mathbf{b}), o_i} \leq \min\{\alpha_{o_i} v_{\pi_{o_i}(\mathbf{b})}, c_{\pi_{o_i}(\mathbf{b})}\}$, we have that

$$u_i(\mathbf{b}) \geq u_i(\mathbf{y}, \mathbf{b}_{-i}) \geq \alpha_{o_i} v_i - b_{\pi_{o_i}(\mathbf{b}), o_i} - \gamma \geq \min\{\alpha_{o_i} v_i, c_i\} - \min\{\alpha_{o_i} v_{\pi_{o_i}(\mathbf{b})}, c_{\pi_{o_i}(\mathbf{b})}\} - \gamma$$

and inequality (2) follows. \square

4. Possible extensions

In this letter, we studied the efficiency of several well-known mechanisms for the allocation of (advertising) positions to strategic budget-constrained users, and proved that their liquid price of anarchy and stability for pure equilibria is exactly 2. Of course, there are multiple interesting open questions that one could attempt to answer here, like exploring *all* position mechanisms and bounding their liquid price of anarchy and stability. In particular, is there any position mechanism with liquid price of anarchy strictly smaller than 2, even for the fundamental case of two players?

Another important direction for future research is to consider more general settings, possibly with incomplete information where both the values and the budgets of the players are randomly drawn from a prior distribution, and bound the liquid price of anarchy of position mechanisms for more general equilibrium notions, like coarse-correlated and Bayes-Nash equilibria. Finally, it might be interesting to study scenarios where the budgets of the players are assumed to be common knowledge (or they can be inferred in some way), and design mechanisms with improved social efficiency guarantees, by exploiting this information.

References

References

- [1] Anshelevich, E., Dasgupta, A., Kleinberg, J. M., Tardos, E., Wexler, T., Roughgarden, T., 2008. The price of stability for network design with fair cost allocation. *SIAM Journal of Computing* 38 (4), 1602–1623.
- [2] Azar, Y., Feldman, M., Gravin, N., Roytman, A., 2017. Liquid price of anarchy. In: *Proceedings of the 10th International Symposium on Algorithmic Game Theory (SAGT)*. pp. 3–15.
- [3] Caragiannis, I., Kaklamani, C., Kanellopoulos, P., Kyropoulou, M., Lucier, B., Leme, R. P., Tardos, E., 2015. Bounding the inefficiency of outcomes in generalized second price auctions. *Journal of Economic Theory* 156, 343–388.
- [4] Caragiannis, I., Voudouris, A. A., 2016. Welfare guarantees for proportional allocations. *Theory of Computing Systems* 59 (4), 581–599.
- [5] Caragiannis, I., Voudouris, A. A., 2018. The efficiency of resource allocation mechanisms for budget-constrained users. In: *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*. pp. 681–698.
- [6] Christodoulou, G., Sgouritsa, A., Tang, B., 2016. On the efficiency of the proportional allocation mechanism for divisible resources. *Theory of Computing Systems* 59 (4), 600–618.
- [7] Clarke, E., 1971. Multipart pricing of public goods. *Public Choice* 11 (1), 17–33.
- [8] Díaz, J., Giotis, I., Kirousis, L. M., Markakis, E., Serna, M. J., 2016. On the stability of generalized second price auctions with budgets. *Theory of Computing Systems* 59 (1), 1–23.

- [9] Dobzinski, S., Paes Leme, R., 2014. Efficiency guarantees in auctions with budgets. In: Proceedings of the 41st International Colloquium on Automata, Languages and Programming (ICALP). pp. 392–404.
- [10] Dütting, P., Fischer, F. A., Parkes, D. C., 2014. Expressiveness and robustness of first-price position auctions. In: Proceedings of the 15th ACM Conference on Economics and Computation (EC). pp. 57–74.
- [11] Edelman, B., Ostrovsky, M., Schwarz, M., 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review* 97 (1), 242–259.
- [12] Groves, T., 1973. Incentives in teams. *Econometrica: Journal of the Econometric Society*, 617–631.
- [13] Koutsoupias, E., Papadimitriou, C. H., 1999. Worst-case equilibria. In: Proceedings of the 16th Symposium on Theoretical Aspects of Computer Science (STACS). pp. 404–413.
- [14] Lu, P., Xiao, T., 2015. Improved efficiency guarantees in auctions with budgets. In: Proceedings of the 16th ACM Conference on Economics and Computation (EC). pp. 397–413.
- [15] Lu, P., Xiao, T., 2017. Liquid welfare maximization in auctions with multiple items. In: Proceedings of the 10th International Symposium on Algorithmic Game Theory (SAGT). pp. 41–52.
- [16] Roughgarden, T., Syrgkanis, V., Tardos, E., 2017. The price of anarchy in auctions. *Journal of Artificial Intelligence Research* 59, 59–101.
- [17] Syrgkanis, V., Tardos, E., 2013. Composable and efficient mechanisms. In: Proceedings of the 45th Annual ACM Symposium on Theory of Computing (STOC). pp. 211–220.
- [18] Varian, H., 2007. Position auctions. *International Journal of Industrial Organization* 25 (6), 1163–1178.
- [19] Vickrey, W., 1961. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance* 16 (1), 8–37.