Liquid Price of Anarchy

Yossi Azar[1*], Michal Feldman[1**], Nick Gravin[2], and Alan Roytman[3***]

[1] Tel Aviv University, Tel Aviv, Israel
email: azar@tau.ac.il, mfeldman@tau.ac.il
[2] Massachusetts Institute of Technology, Cambridge, USA
email: ngravin@mit.edu
[3] University of Copenhagen, Copenhagen, Denmark
email: alanr@di.ku.dk

Abstract. Incorporating budget constraints into the analysis of auctions has become increasingly important, as they model practical settings more accurately. The social welfare function, which is the standard measure of efficiency in auctions, is inadequate for settings with budgets, since there may be a large disconnect between the value a bidder derives from obtaining an item and what can be liquidated from her. The *Liquid Welfare* objective function has been suggested as a natural alternative for settings with budgets. Simple auctions, like simultaneous item auctions, are evaluated by their performance at equilibrium using the Price of Anarchy (PoA) measure – the ratio of the objective function value of the optimal outcome to the worst equilibrium. Accordingly, we evaluate the performance of simultaneous item auctions in budgeted settings by the *Liquid Price of Anarchy* (LPoA) measure – the ratio of the optimal Liquid Welfare to the Liquid Welfare obtained in the worst equilibrium. For pure Nash equilibria of simultaneous first price auctions, we obtain a bound of 2 on the LPoA for additive buyers. Our results easily extend to the larger class of fractionally-subadditive valuations. Next we show that the LPoA of mixed Nash equilibria for first price auctions with additive bidders is bounded by a constant. Our proofs are robust, and can be extended to achieve similar bounds for Bayesian Nash equilibria. To derive our results, we develop a new technique in which some bidders deviate (surprisingly) toward a non-optimal solution. In particular, this technique goes beyond the smoothness-based approach.

* Supported in part by the Israel Science Foundation (grant No. 1506/16), by the I-CORE program (Center No. 4/11), and by the Blavatnik Fund.

** This work was partially supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement number 337122.

*** This work was partially supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement number 337122, by Thorup's Advanced Grant DFF-0602-02499B from the Danish Council for Independent Research, by grant number 822/10 from the Israel Science Foundation, and by the Israeli Centers for Research Excellence (ICORE) program.

1 Introduction

Budget constraints have become an important practical consideration in most existing auctions, as reflected in recent literature (see, e.g., [4, 6, 20, 37]), because they model reality more accurately. The issue of limited liquidity of buyers arises when transaction amounts are large and may exhaust bidders' liquid assets, as is the case for privatization auctions in Eastern Europe and FCC spectrum auctions in the U.S. (see, e.g., [5]). As another example, advertisers in Google Adword auctions are instructed to specify their budget even before specifying their bids and keywords. Many other massive electronic marketplaces have a large number of participants with limited liquidity, which impose budget constraints. Buyers would not borrow money from a bank to partake in multiple auctions on eBay, and even with available credit, they only have a limited amount of attention, so that in aggregate they cannot spend too much money by participating in every auction online. Finally, budget constraints also arise in small scale systems, such as the reality TV show Storage Wars, where people participate in cash-only auctions to win the content of an expired storage locker with an unknown asset.

Maximizing social welfare is a classic objective function that has been extensively studied within the context of resource allocation problems, and auctions in particular. The *social welfare* of an allocation is the sum of agents' valuations for their allocated bundles. Unfortunately, in settings where agents have limited budgets (hereafter, *budgeted settings*), the social welfare objective fails to accurately capture what happens in practice. Consider, for example, an auction in which there are two bidders and one item to be allocated among the bidders. One bidder has a high value but a very small budget, while the second bidder has a medium value along with a medium budget. In this case, a high social welfare is achieved by allocating the item to the bidder who values the item highly. In contrast, most Internet advertising and electronic marketplaces (such as Google and eBay) would allocate the item in the opposite way, namely to the bidder with a medium value and budget. Indeed, it seems reasonable to favor participants with substantial investments and engagement in the economical system to maintain a healthy economy regardless of the marketplace intermediary's personal gains. Hence, the social welfare objective is a poor model for how auctions are executed in reality.

In this work, we study the efficiency of simultaneous first price auctions in budgeted settings. Following Dobzinski and Paes Leme [21] (see also [11, 23, 32, 38]), we measure the efficiency of outcomes in budgeted settings according to their *Liquid Welfare* objective, motivated as follows. In the mechanism design literature, a buyer i with additive values for items v_{ij} (where v_{ij} denotes buyer i's value for item j) and a budget cap B_i is usually modeled with *budget additive* valuations $v_i(S) = \min(B_i, \sum_{j \in S} v_{ij})$, where S is the set of items that player i receives (see, e.g., Lehmann et al. [31] and many follow-up works). Budget additive valuations are convenient to work with (they form a simple subclass of submodular valuations) and, from the designer's perspective, are a natural proxy for the contribution of each bidder to the economical system.

However, in reality such valuations do not capture the real preferences of the buyers, since each buyer usually prefers to get as many items of high value as possible $v_i(S) = \sum_{j \in S} v_{ij}$, with the only concern being that her total payments for the received set of items S, denoted by $p_i(S)$, should not exceed her budget constraint $p_i(S) \leq B_i$. To reconcile this discrepancy, Dobzinski and Paes Leme [21] proposed to evaluate the welfare of buyers in budgeted settings according to their *admissibility-to-pay*; that is, the minimum between the buyer's value for the allocated bundle and the buyer's budget. The aggregate welfare according to this definition is termed the *Liquid Welfare* (LW). Hence, the Liquid Welfare objective can be seen as a natural analogue to social welfare in budgeted settings, as it simultaneously captures the health of an economic system while still modeling buyers as preferring items of high value, despite budget constraints.

For simultaneous first price auctions, we use the following natural *item-clearing mechanism* for each individual item. Each player submits a bid they are willing to pay for the whole item, along with the maximal fraction of the item they are willing to purchase. Then, in decreasing order of the bids and as long as some fraction of the item remains to be allocated, each buyer receives their requested fraction of the item (or whatever remains), and pays their bid multiplied by the fraction they received. In the context of additive values, we model players' utilities for each item as their value for the item minus their submitted bid (both of which are scaled by the fraction of the item they receive).

Our model is closely related to a prominent simultaneous item auction format with heterogeneous items, which has been extensively studied recently. In such auctions, buyers submit bids simultaneously on all items, and the allocation and prices are determined separately for each individual item, based only on the bids submitted for that item. This format is similar to auctions used in practice (e.g., eBay auctions). The standard measure for quantifying efficiency in such settings is the *Price of Anarchy* (PoA) [29,35,38], defined as the ratio of the optimal social welfare to the social welfare of the worst equilibrium. In budgeted settings, it is thus natural to quantify the efficiency of such auctions by the *Liquid Price of Anarchy* (LPoA), defined as the ratio of the optimal Liquid Welfare to the Liquid Welfare of the worst equilibrium.

New Techniques. The most common framework for analyzing the Price of Anarchy of games and auctions is the *smoothness* framework (see, e.g., [35,38]). Such techniques usually involve a thought experiment in which each player deviates toward some strategy related to the optimal solution, and hence the total utility of all players can be bounded appropriately. One important and necessary condition for applying the smoothness framework is that the objective function must dominate the sum of utilities (which holds for social welfare). However, this technique falls short in the case of Liquid Welfare, since a bidder's utility can be arbitrarily higher than their admissibility-to-pay, and in aggregate, bidders may achieve a total utility that is much larger than the Liquid Welfare at equilibrium. To overcome this issue, we develop new techniques to bound the LPoA in budgeted settings. Our techniques include a novel type of hypothetical deviation that is used to *upper bound* the aggregate utility of bidders (in addition to the traditional deviation that is used to lower bound it), and the consideration of a special set of carefully chosen bidders to engage in these hypothetical deviations (see more details in the full version of our paper [3]). To the best of our knowledge, most prior techniques, including those that depart from the smoothness framework (e.g., [25]), examine the utility derived when *every* player deviates toward the optimal solution.

With our new techniques at hand, we address the following question: *What is the Liquid Price of Anarchy of simultaneous first price item auctions in settings with budgets?*

Clarifying Remarks and Examples

Settings where agents have additive valuations and are constrained by budgets (as in our setting) should not be confused with settings with *budget additive* valuations. The latter assumes quasilinear utilities, while the former does not[4]. The class of budget additive valuations is a proper subclass of submodular valuations, a setting for which the Price of Anarchy of simultaneous combinatorial auctions is well understood, and known to be bounded by a constant[5]. However, these results do not apply to the budgeted setting, since a bidder's perspective and consequently their behavior at equilibrium is very different from the budget additive setting (see the full version of our paper [3] for an example and a more detailed discussion).

The budgeted simultaneous item bidding setting has also been studied by [38], where a different approach was taken. They measured the social welfare at equilibrium against the optimal Liquid Welfare. Note that according to their measure, the benchmark (i.e., optimal Liquid Welfare) may be lower than the measured welfare. Please see the full version of our paper [3] for an example illustrating the difference between their measure and our LPoA measure.

Our Contributions

We show that simultaneous first price item auctions achieve nearly optimal performance, i.e., a constant Liquid Price of Anarchy. Our main result concerns the case in which agent valuations are additive (i.e., agent i's value for item j is v_{ij} and the value for a set of items is the sum of the individual valuations, each of which is scaled by the corresponding fraction received).

Main theorem: For simultaneous first price auctions with additive bidders and divisible items, the LPoA with respect to mixed Nash equilibria and Bayesian Nash equilibria is constant.

We also show that for pure Nash equilibria in simultaneous first price auctions, our results hold for more general settings.

[4] The difference is also pointed out in the literature on the design of truthful combinatorial auctions [20, 21].

[5] In particular, there are tight PoA bounds of $\frac{e}{e-1}$ for submodular bidders, and 2 for subadditive bidders.

Theorem: For fractionally-subadditive bidders, the LPoA of pure Nash equilibria in simultaneous first price auctions is 2. Moreover, this bound is tight.

The following remarks are in order:

1. In settings without budgets, simultaneous first price item auctions for additive bidders reduce to m independent auctions (where m is the number of items). In contrast, when agents have budget constraints, the separate auctions exhibit non-trivial dependencies even under additive valuations.
2. Since fractionally-subadditive valuations are not typically defined over divisible items, we discretize the bidding space so that requested fractions of items can only be multiples of a fixed small size in our fractionally-subadditive results. This essentially induces an indivisible setting with discrete items, and hence fractionally-subadditive valuations are well-defined.

Related Work

There is a vast literature in algorithmic game theory that incorporates budgets into the design of incentive compatible mechanisms. The paper of [6] showed that, in the case of one divisible good, the adaptive clinching auction is incentive compatible under some assumptions. Moreover, the work of [37] initiated the design of incentive compatible mechanisms in the context of reverse auctions, where the payments of the auctioneer cannot exceed a hard budget constraint (follow-up works include [1,4,12,15,22]). A great deal of work focused on designing incentive compatible mechanisms that approximately maximize the auctioneer's revenue in various settings with budget-constrained bidders [9,13,30,33,34]. Some works analyzed how budgets affect markets and non-truthful mechanisms [5,14].

Earlier work on multi-unit auctions with budgets deals with designing incentive compatible mechanisms that always produce Pareto-optimal allocations [20]. The results in this line of work are mostly negative with a notable exception of mechanisms based on Ausubel's adaptive clinching auction framework [2].

Some recent results concern the design of incentive compatible mechanisms with respect to the Liquid Welfare objective, introduced by [21]. They gave a constant approximation for the auction that sells a single divisible good to additive buyers with budgets. In a follow-up work, [32] gave an $O(1)$-approximation for bidders with general valuations in the single-item setting. The work of [23] extended the notion of a combinatorial Walrasian equilibrium (see [26]) to settings with budgets. They showed that their generalization, termed a lottery pricing equilibrium, achieves high Liquid Welfare. They also argued how to efficiently compute randomized allocations that have near-optimal Liquid Welfare for large classes of valuation functions (including subadditive valuations).

A large body of literature is concerned with simultaneous item bidding auctions. These simple auctions have been studied from a computational perspective [10,19]. There is also extensive work addressing the Price of Anarchy of such simple auctions (see [36] for more general Price of Anarchy results). The work of [16] initiated the study of simultaneous item auctions within the Price of Anarchy framework. The authors showed that, for second price auctions, the social welfare of every Bayesian Nash equilibrium is a 2-approximation to the optimal social welfare, even for players with fractionally-subadditive valuation functions. A large amount of follow-up work [7,8,17,24,25,28,38] made significant progress in understanding simultaneous item auctions, but all of these works measure inefficiency only with respect to the social welfare objective.

Much less is known about the Price of Anarchy in auctions for objectives other than social welfare. In fact, we are aware of only one such work [27], which studies the revenue of simultaneous auctions with reserve prices for *single-parameter* bidders with regular distributions. This work essentially reduces the revenue maximization problem to the welfare maximization problem for virtual values in single-parameter settings and then employs smoothness analysis to bound virtual value welfare. We note that this approach fails for multi-parameter settings such as simultaneous multi-item auctions with additive valuations.

The work of [38] considered Liquid Welfare when measuring the inefficiency of equilibria. They gave various Price of Anarchy results, developed a smoothness framework for broad solution concepts such as correlated and Bayesian Nash equilibria, and explored composition properties of various mechanisms. They extended their results to the setting where players are budget-constrained, achieving similar approximation guarantees when comparing the *social welfare achieved at equilibrium* to the *optimal Liquid Welfare*. In particular, their results imply an $\frac{e}{e-1}$-approximation for simultaneous first price auctions, and a 2-approximation for all-pay auctions and simultaneous second price auctions under the no-overbidding assumption. While [38] show that the social welfare at equilibrium cannot be much worse than the optimal Liquid Welfare, one should note that the social welfare at equilibrium can be arbitrarily better than the optimal Liquid Welfare (e.g., if all budgets are small, the optimal Liquid Welfare is small). It is useful to note that, in general, the ratio between the Liquid Welfare at equilibrium and the social welfare at equilibrium can be arbitrarily bad (if all budgets are small, then the Liquid Welfare of any allocation is small, while players' values for received goods can be arbitrarily large).

The works of [11,18] also considered the setting where players have budgets and studied the same ratio we consider, namely the *Liquid Welfare at equilibrium* to the *optimal Liquid Welfare*. In [11], they studied the proportional allocation mechanism, which concerns auctioning off one divisible item proportionally according to the bids that players submit. They showed that, assuming players have concave non-decreasing valuation functions, the Liquid Welfare at coarse-correlated equilibria and Bayesian Nash equilibria achieve at least a constant fraction of the optimal Liquid Welfare. It should be noted that, for random allocations, they measure the benchmark at equilibrium ex-ante over the randomness of the allocation, i.e., $\sum_{i=1}^{n} \min\{\mathbb{E}_{\mathbf{v}_{\text{-}i},B_{\text{-}i}}[v_i(x_i)], B_i\}$, where v_i is player i's valuation, B_i is player i's budget, and x_i denotes the allocation player i receives. In contrast, for random allocations, we use the stronger ex-post measure of the expected Liquid Welfare at equilibrium given by $\sum_{i=1}^{n} \mathbb{E}[\min\{v_i(x_i), B_i\}]$. The work of [18] studied a similar setting, except that multiple divisible items were considered. They gave improved bounds for coarse-correlated equilibria even with multiple items and also an improved bound for Bayesian Nash equilibria with one item. In addition, they studied the polyhedral environment and showed an exact bound of 2 for pure Nash equilibria when agents have subadditive valuations.

2 Model and Preliminaries

We consider *simultaneous first price item auctions*, in which m heterogeneous items are sold to n bidders (or players) in m independent auctions. We first describe our notation in the context of indivisible items, and then describe the divisible model. A bidder's *strategy* is a bid vector $b_i \in \mathbb{R}^m_{\geq 0}$, where b_{ij} represents player i's bid for item j. We use $\mathbf{b}$ to denote the bid profile $\mathbf{b} = (b_1, \ldots, b_n)$, and we will often use the notation $\mathbf{b} = (b_i, \mathbf{b}_{\text{-}i})$ to denote the strategy profile where player i bids b_i and the remaining players bid according to $\mathbf{b}_{\text{-}i} = (b_1, \ldots, b_{i-1}, b_{i+1}, \ldots, b_n)$.

The outcome of an auction consists of an allocation rule $\mathbf{x}$ and payment rule $\mathbf{p}$. The allocation rule $\mathbf{x}$ maps bid profiles to an allocation vector for each bidder i, where $x_i(\mathbf{b}) = (x_{i1}, \ldots, x_{im})$ denotes the set of items won by player i ($x_{ij} = 1 \Leftrightarrow$ player i wins item j). In a *simultaneous first price auction*, each item is allocated to the highest bidder (breaking ties according to some rule) and the winner pays their bid. The total payment of bidder i is $p_i(\mathbf{b}) = \sum_{j \in x_i(\mathbf{b})} b_{ij}$.

Each player i has a *valuation function* v_i, which maps sets of items to $\mathbb{R}_{\geq 0}$ (v_i captures how much player i values item bundles), and a budget B_i. We assume that all valuations are normalized and monotone, i.e., $v_i(\emptyset) = 0$ and $v_i(S) \leq v_i(T)$ for any $i \in [n]$ and $S \subseteq T \subseteq [m]$. We mostly consider bidders with additive valuations, i.e., $v_i(S) = \sum_{j \in S} v_{ij}$ (where v_{ij} denotes agent i's value for item j). The *utility* $u_i(x_i(\mathbf{b}))$ of each player i is $v_i(x_i(\mathbf{b})) - p_i(\mathbf{b}) = \sum_j v_{ij} \cdot x_{ij} - p_i(\mathbf{b})$ if $p_i(\mathbf{b}) \leq B_i$; and $u_i(x_i(\mathbf{b})) = -\infty$ if $p_i(\mathbf{b}) > B_i$. Buyers select their bids strategically in order to maximize utility.

For divisible items, agent's i bid $\mathbf{b}_{ij}$ consists of two parameters for each item j: a price b_{ij} and a desired fraction δ_{ij}. For each item j, in decreasing order of b_{ij}, and as long as some fraction of item j remains, each buyer i receives a fraction of item j given by $x_{ij} = \delta_{ij}$ (or whatever remains). If the agent receives an x_{ij}-fraction of item j, then their value is given by $v_i(x_{ij}) = v_{ij} \cdot x_{ij}$ and they pay $b_{ij} \cdot x_{ij}$. We write all individual bids $\mathbf{b}_{ij}$ on item j as a vector $\mathbf{b}_{\cdot j}$ and bids for all items as $\mathbf{b}$.

Definition 1 (Pure Nash Equilibrium). *A bid profile* $\mathbf{b}$ *is a Pure Nash Equilibrium if, for any player* i *and any deviating bid* b'_i*:* $u_i(b_i, \mathbf{b}_{\text{-}i}) \geq u_i(b'_i, \mathbf{b}_{\text{-}i})$.

A mixed Nash equilibrium is defined similarly, except that bidding strategies can be randomized $b_i \sim s_i$ and utility is measured in expectation over the joint bid distribution $\mathbf{s} = s_1 \times \cdots \times s_n$.

Definition 2 (Mixed Nash Equilibrium). *A bid profile* $\mathbf{s}$ *is a mixed Nash equilibrium if, for any player* i *and any deviating bid* b'_i*:* $\mathbb{E}_{\mathbf{b} \sim \mathbf{s}}[u_i(b_i, \mathbf{b}_{\text{-}i})] \geq \mathbb{E}_{\mathbf{b}_{\text{-}i} \sim \mathbf{s}_{\text{-}i}}[u_i(b'_i, \mathbf{b}_{\text{-}i})]$.

Note that, in general, we assume the bidding space is discretized (i.e., each player can only bid in multiples of a sufficiently small value ε). This is done to ensure that there always exists a mixed Nash equilibrium, as otherwise we do not have a finite game.

Definition 3 (Liquid Welfare). *The Liquid Welfare, denoted by* LW*, of an allocation* $\mathbf{x}$ *is given by* $\mathrm{LW}(\mathbf{x}) = \sum_{i\in[n]} \min\{v_i(x_i), B_i\}$*. For random allocations, we use the measure given by* $\mathrm{LW}(\mathbf{x}) = \sum_{i\in[n]} \mathbb{E}[\min\{v_i(x_i), B_i\}]$.

For a given vector of valuations $\mathbf{v} = (v_1, \dots, v_n)$, we use $\mathrm{OPT}(\mathbf{v})$ to denote the Liquid Welfare of an optimal outcome given by the expression $\mathrm{OPT}(\mathbf{v}) = \max_{x_1,\dots,x_n} \sum_i \min\{v_i(x_i), B_i\}$, where the bundles x_i form a fractional partition of $[m]$ (i.e., $\sum_i x_{ij} = 1$ for any j). We often use OPT instead of $\mathrm{OPT}(\mathbf{v})$ when the context is clear.

Definition 4 (Liquid Price of Anarchy). *Given a fixed valuation profile* $\mathbf{v}$*, the* Liquid Price of Anarchy *(LPoA) is the worst-case ratio between the optimal Liquid Welfare and the expected Liquid Welfare at an equilibrium (pure, mixed, or Bayesian Nash) and is given by* $LPoA(\mathbf{v}) = \sup_{\mathbf{s}} \left\{ \frac{\mathrm{OPT}(\mathbf{v})}{\mathrm{LW}(\mathbf{s}(\mathbf{v}))} \;\middle|\; \mathbf{s} \in \textit{Equilibria} \right\}$.

3 Simultaneous First Price Auctions

In this section, we prove our main theorem that, for mixed Nash equilibria, the Liquid Price of Anarchy of simultaneous first price auctions is constant. In what follows we build up notation and intuition toward the proof. Recall that agents have additive valuations and submit separate bids on each item. We assume that the buyers bid according to a mixed Nash equilibrium $\mathbf{b} \sim \mathbf{s}$. For all items we can define an "expected price per item" at equilibrium or just a "price per item" as $\mathbf{p} = (p_1, \dots, p_m)$, where $p_j = \sum_{i=1}^{n} \mathbb{E}_{\mathbf{b}\sim\mathbf{s}}[b_{ij} \cdot x_{ij}(\mathbf{b}_{\cdot j})]$.

Each bidder i has a good chance of winning a particular item if they bid above the expected price of this item. To this end we define boosted prices $\overline{\mathbf{p}} = (\overline{p}_1, \dots, \overline{p}_m)$, where $\overline{p}_j = \alpha p_j$ for some $\alpha > 1$ ($\alpha = 2$ will be sufficient for us). One simple observation about $\overline{\mathbf{p}}$ is the following:

Observation 1 *Revenue is related to prices, namely:* $\mathrm{REV}(\mathbf{s}) = \frac{1}{\alpha} \sum_{j=1}^{m} \overline{p}_j$*, where* $\mathrm{REV}(\mathbf{s})$ *denotes the expected revenue at the equilibrium profile* $\mathbf{s}$.

We next show that if players bid on a fraction of an item j according to $\overline{p}_j$, then they win a large fraction of j in expectation. The proof is in the full version of our paper [3].

Proposition 1. *For any item* j*, if a player bids on a* δ*-fraction of item* j *at price* $\overline{p}_j$*, then the player receives in expectation at least a* $\delta \cdot \left(1 - \frac{1}{\alpha}\right)$*-fraction of item* j.

When relating prices to Liquid Welfare we notice that

Observation 2 *Revenue is bounded by the Liquid Welfare:* $\text{REV}(\mathbf{s}) \leq \text{LW}(\mathbf{s})$, *where* $\text{LW}(\mathbf{s})$ *denotes the expected Liquid Welfare at the equilibrium profile* $\mathbf{s}$.

We consider the following Linear Program for the allocation problem with the goal of optimizing Liquid Welfare.

$$\text{Maximize} \quad \sum_{i=1}^{n}\sum_{j=1}^{m} v_{ij} \cdot z_{ij} \qquad\qquad \text{Liquid linear program (LLP)}$$

$$\text{Subject to} \quad \sum_{j} v_{ij} \cdot z_{ij} \leq B_i \quad \forall i; \qquad \sum_{i} z_{ij} \leq 1 \quad \forall j; \qquad z_{ij} \geq 0 \quad \forall i,j.$$

We denote by $\mathbf{y} = (y_{ij})$ the optimal solution to LLP. Notice that the solution for the Liquid Welfare never benefits from allocating a set of items to a player such that their value for the set exceeds their budget. Thus

Observation 3 *The solution to LLP is equal to the optimal allocation, namely:* $\sum_{i=1}^{n}\sum_{j=1}^{m} v_{ij} \cdot y_{ij} = \text{OPT}$.

We now define some notation that will be useful in order to obtain our result. We let $q_{ij} \overset{\text{def}}{=} \mathbb{E}_{\mathbf{b}\sim\mathbf{s}}[x_{ij}(\mathbf{b}_{\cdot j})]$ be the expected fraction of item j that player i receives at an equilibrium strategy $\mathbf{s}$. In addition, for each agent i, we consider a set of high value items $J_i \overset{\text{def}}{=} \left\{ j \mid v_{ij} \geq \overline{p}_j \right\}$. In particular, in our deviations, we ensure that each player i only bids on items in J_i, since this guarantees that they attain nonnegative utility if they win such items. We further define $Q_i \overset{\text{def}}{=} \mathbf{Pr}_{\mathbf{b}\sim\mathbf{s}}[v_i(x_i(\mathbf{b})) \geq B_i]$ to be the probability that $v_i(x_i) \geq B_i$ at equilibrium (recall that x_i denotes the random set that player i receives in the mixed Nash equilibrium). We also define two sets of bidders, the first one is for budget feasibility reasons and the second is for bidders that often fall under their budget in equilibrium (these sets need not be disjoint). In particular, for a fixed parameter $\gamma > 1$ ($\gamma = 4$ will be sufficient for us), we define sets $\mathcal{I}_1$ and $\mathcal{I}_2$: $\mathcal{I}_1 \overset{\text{def}}{=} \left\{ i \;\middle|\; \gamma \sum_{j \in J_i} \overline{p}_j \cdot q_{ij} \leq B_i \right\}$ and $\mathcal{I}_2 \overset{\text{def}}{=} \left\{ i \;\middle|\; Q_i \leq \frac{1}{2\gamma} \right\}$.

Throughout our proof, we focus on bidders in the set $\mathcal{I} \overset{\text{def}}{=} \mathcal{I}_1 \cap \mathcal{I}_2$. As mentioned, the way we achieve our main result is to consider two types of deviations for all players in $\mathcal{I}$, the first of which is a deviation towards an optimal solution, and the second of which is a γ-boosting deviation (where players essentially bid on a larger fraction of items by a factor γ relative to what they receive at equilibrium, namely $\gamma \cdot q_{ij}$). We only consider deviations for players in set $\mathcal{I}$ for the following reasons. Players in $\mathcal{I}_1$ are guaranteed to respect their budgetary constraints in our γ-boosting deviation. Players that do not belong to $\mathcal{I}_2$ have the property that the value they receive at equilibrium often exceeds their budgets, and hence such players already contribute a lot to the Liquid Welfare at equilibrium. In particular, whenever a player receives a value that exceeds their budget, their contribution to the Liquid Welfare at equilibrium is at least as much as their contribution in the optimal Liquid Welfare, which is always bounded above by their budget. Hence, we need only consider players in $\mathcal{I}_2$ as far as deviations are concerned. We define sets $\overline{\mathcal{I}}_1 \stackrel{\text{def}}{=} [n] \setminus \mathcal{I}_1$, $\overline{\mathcal{I}}_2 \stackrel{\text{def}}{=} [n] \setminus \mathcal{I}_2$, and $\overline{\mathcal{I}} \stackrel{\text{def}}{=} [n] \setminus \mathcal{I}$. To this end, we relate the total budget of bidders outside of the set $\mathcal{I}$ to the revenue at equilibrium $\mathbf{s}$. The proof is in the full version of our paper [3].

Proposition 2. *The total budget of players in $\overline{\mathcal{I}}$ is small:* $\sum_{i\in\overline{\mathcal{I}}} B_i < \alpha \cdot \gamma \cdot \mathrm{REV}(\mathbf{s}) + \sum_{i\in\overline{\mathcal{I}}_2} B_i$.

To achieve our result, we essentially consider two main ideas for player deviations in set $\mathcal{I}$. The first idea is to use the solution to LLP as guidance to claim that players can extract a large amount of value relative to the optimal solution. Define the first LLP deviation to be $\mathbf{b}_1 = (b'_i, \mathbf{b}_{\text{-}i})$, where in b'_i buyer i bids on a y_{ij}-fraction of each item $j \in J_i$ with price $\overline{p}_j$. We note that the LLP deviation $\mathbf{b}_1$ is feasible, since $v_{ij} \geq \overline{p}_j$ for every $j \in J_i$, and $\sum_j v_{ij} \cdot y_{ij} \leq B_i$ as $\mathbf{y}$ is a solution to LLP.

Lemma 1 (LLP deviations). *Buyers in $\mathcal{I}$ at equilibrium $\mathbf{s}$ derive large value:*

$$\sum_{i\in\mathcal{I}} \sum_j v_{ij} \cdot q_{ij} \geq \left(1 - \frac{1}{\alpha}\right) \left(\mathrm{OPT} - \alpha\,(1+\gamma)\,\mathrm{REV}(\mathbf{s}) - \sum_{i\in\overline{\mathcal{I}}_2} B_i\right).$$

We defer the proof of Lemma 1 to the full version of our paper [3]. We now turn to our second type of deviation, but we need to further restrict the set of items that players bid on. In particular, we let $\Gamma_i = \left\{ j \;\middle|\; q_{ij} \leq \frac{1}{\gamma} \right\}$, and define $G_i = J_i \cap \Gamma_i$. The set Γ_i is defined to ensure that each player i that deviates never bids on a fraction of an item j that exceeds 1. We now define the γ-boosting deviation as $\mathbf{b}_2 = (b'_i, \mathbf{b}_{\text{-}i})$, where in b'_i buyer i bids on a $\gamma \cdot q_{ij}$-fraction of each item $j \in G_i$ with price $\overline{p}_j$, where $\gamma > 1$ is a constant. Note that each $\mathbf{b}_2$ deviation for every $i \in \mathcal{I}$ is feasible since $\mathcal{I} \subseteq \mathcal{I}_1$.

Lemma 2 (γ-boosting deviation). *The value derived by buyers in $\mathcal{I}$ is comparable to the Liquid Welfare obtained at equilibrium:*

$$\left(1 - \frac{\alpha}{\gamma(\alpha-1)}\right) \sum_{i\in\mathcal{I}} \sum_j v_{ij} \cdot q_{ij} \leq \alpha \cdot \mathrm{REV}(\mathbf{s}) + 2 \cdot \mathrm{LW}(\mathbf{s}) - \frac{1}{\gamma} \sum_{i\in\overline{\mathcal{I}}_2} B_i.$$

We defer the proof of Lemma 2 to the full version of our paper [3]. Now we have all necessary components to conclude the proof of our main theorem and show that the Liquid Price of Anarchy of any mixed Nash equilibrium is bounded. The proof is given in the full version of our paper [3].

Theorem 1 *For mixed Nash equilibria, the Liquid Price of Anarchy of simultaneous first price auctions is constant (at most* 26*).*

The above reasoning also extends to Bayesian Nash equilibria with the same LPoA bound.

Theorem 2 *For Bayesian Nash equilibria, the Liquid Price of Anarchy of simultaneous first price auctions is constant (at most* 26*).*

We omit the proof as it is very similar to the proof of Theorem 1. We also study pure Nash equilibria of simultaneous first price auctions. The proof of the next theorem is given in the full version of our paper [3].

Theorem 3 *Consider a simultaneous first price auction where budgeted bidders have fractionally-subadditive valuations*[6]*. If* $\mathbf{b}$ *is a pure Nash equilibrium, then* $\mathrm{LW}(\mathbf{b}) \geq \frac{\mathrm{OPT}}{2}$.

A complementary tightness result for Theorem 3 is given in the full version of our paper [3]. Unfortunately, this result is not quite satisfying compared to mixed Nash equilibria, as pure Nash equilibria might not even exist (see the full version of our paper [3]).

References

1. Anari, N., Goel, G., Nikzad, A.: Mechanism design for crowdsourcing: An optimal 1-1/e competitive budget-feasible mechanism for large markets. In: FOCS. pp. 266–275 (2014)
2. Ausubel, L.M.: An efficient ascending-bid auction for multiple objects. American Economic Review 94(5), 1452–1475 (2004)
3. Azar, Y., Feldman, M., Gravin, N., Roytman, A.: Liquid price of anarchy. CoRR abs/1511.01132 (2015)
4. Bei, X., Chen, N., Gravin, N., Lu, P.: Budget feasible mechanism design: From prior-free to bayesian. In: STOC. pp. 449–458 (2012)
5. Benoit, J.P., Krishna, V.: Multiple-object auctions with budget constrained bidders. Review of Economic Studies 68(1), 155–179 (2001)
6. Bhattacharya, S., Conitzer, V., Munagala, K., Xia, L.: Incentive compatible budget elicitation in multi-unit auctions. In: SODA. pp. 554–572 (2010)
7. Bhawalkar, K., Roughgarden, T.: Welfare guarantees for combinatorial auctions with item bidding. In: SODA. pp. 700–709 (2011)
8. Bhawalkar, K., Roughgarden, T.: Simultaneous single-item auctions. In: WINE. pp. 337–349 (2012)
9. Borgs, C., Chayes, J., Immorlica, N., Mahdian, M., Saberi, A.: Multi-unit auctions with budget-constrained bidders. In: EC. pp. 44–51 (2005)
10. Cai, Y., Papadimitriou, C.: Simultaneous bayesian auctions and computational complexity. In: EC. pp. 895–910 (2014)
11. Caragiannis, I., Voudouris, A.A.: Welfare guarantees for proportional allocations. In: SAGT. pp. 206–217 (2014)
12. Chan, H., Chen, J.: Truthful multi-unit procurements with budgets. In: WINE. pp. 89–105 (2014)
13. Chawla, S., Malec, D., Malekian, A.: Bayesian mechanism design for budget-constrained agents. In: EC. pp. 253–262 (2011)

[6] Valuation v is fractionally-subadditive or equivalently XOS if there is a set of additive valuations $A = \{a_1, \ldots, a_\ell\}$ such that $v_i(S) = \max_{a \in A} a(S)$ for every $S \subseteq [m]$. XOS is a super class of submodular and additive valuations.

14. Che, Y.K., Gale, I.: Standard auctions with financially constrained bidders. Review of Economic Studies 65(1), 1–21 (1998)
15. Chen, N., Gravin, N., Lu, P.: On the approximability of budget feasible mechanisms. In: SODA. pp. 685–699 (2011)
16. Christodoulou, G., Kovács, A., Schapira, M.: Bayesian combinatorial auctions. In: ICALP. pp. 820–832 (2008)
17. Christodoulou, G., Kovács, A., Sgouritsa, A., Tang, B.: Tight bounds for the price of anarchy of simultaneous first price auctions. CoRR abs/1312.2371 (2013)
18. Christodoulou, G., Sgouritsa, A., Tang, B.: On the efficiency of the proportional allocation mechanism for divisible resources. In: SAGT. pp. 165–177 (2015)
19. Dobzinski, S., Fu, H., Kleinberg, R.: On the complexity of computing an equilibrium in combinatorial auctions. In: SODA. pp. 110–122 (2015)
20. Dobzinski, S., Lavi, R., Nisan, N.: Multi-unit auctions with budget limits. In: FOCS. pp. 260–269 (2008)
21. Dobzinski, S., Paes Leme, R.: Efficiency guarantees in auctions with budgets. In: ICALP. pp. 392–404 (2014)
22. Dobzinski, S., Papadimitriou, C., Singer, Y.: Mechanisms for complement-free procurement. In: EC. pp. 273–282 (2011)
23. Dughmi, S., Eden, A., Feldman, M., Fiat, A., Leonardi, S.: Lottery pricing equilibria. In: EC. pp. 401–418 (2016)
24. Dütting, P., Henzinger, M., Starnberger, M.: Valuation compressions in vcg-based combinatorial auctions. In: WINE. pp. 146–159 (2013)
25. Feldman, M., Fu, H., Gravin, N., Lucier, B.: Simultaneous auctions are (almost) efficient. In: STOC. pp. 201–210 (2013)
26. Feldman, M., Gravin, N., Lucier, B.: Combinatorial walrasian equilibrium. In: STOC. pp. 61–70 (2013)
27. Hartline, J., Hoy, D., Taggart, S.: Price of anarchy for auction revenue. In: EC. pp. 693–710 (2014)
28. Hassidim, A., Kaplan, H., Mansour, Y., Nisan, N.: Non-price equilibria in markets of discrete goods. In: EC. pp. 295–296 (2011)
29. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. In: STACS. pp. 404–413 (1999)
30. Laffont, J.J., Robert, J.: Optimal auction with financially constrained buyers. Economics Letters 52(2), 181–186 (1996)
31. Lehmann, B., Lehmann, D., Nisan, N.: Combinatorial auctions with decreasing marginal utilities. Games and Economic Behavior 55(2), 270–296 (2006)
32. Lu, P., Xiao, T.: Improved efficiency guarantees in auctions with budgets. In: EC. pp. 397–413 (2015)
33. Malakhov, A., Vohra, R.V.: Optimal auctions for asymmetrically budget constrained bidders. Review of Economic Design 12(4), 245–257 (2008)
34. Pai, M., Vohra, R.V.: Optimal auctions with financially constrained bidders. Discussion papers, Northwestern University, Center for Mathematical Studies in Economics and Management Science (Aug 2008)
35. Roughgarden, T.: Intrinsic robustness of the price of anarchy. In: STOC. pp. 513–522 (2009)
36. Roughgarden, T., Tardos, E.: Introduction to the inefficiency of equilibria. In: Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V. (eds.) Algorithmic Game Theory. Cambridge University Press, New York, NY, USA (2007)
37. Singer, Y.: Budget feasible mechanisms. In: FOCS. pp. 765–774 (2010)
38. Syrgkanis, V., Tardos, E.: Composable and efficient mechanisms. In: STOC. pp. 211–220 (2013)