

A Bridge between Liquid and Social Welfare in Combinatorial Auctions with Submodular Bidders

Dimitris Fotakis*

Yahoo Research, NY, USA
National Technical University of Athens
Athens, Greece
dfotakis@oath.com

Kyriakos Lotidis*

National Technical University of Athens
Athens, Greece
klot@corelab.ntua.gr

Chara Podimata*

Harvard University
Cambridge, USA
tpodimata@g.harvard.edu

Abstract

We study incentive compatible mechanisms for Combinatorial Auctions where the bidders have submodular (or XOS) valuations and are budget-constrained. Our objective is to maximize the *liquid welfare*, a notion of efficiency for budget-constrained bidders introduced by Dobzinski and Paes Leme (2014). We show that some of the known truthful mechanisms that best-approximate the social welfare for Combinatorial Auctions with submodular bidders through demand query oracles can be adapted, so that they retain truthfulness and achieve asymptotically the same approximation guarantees for the liquid welfare. More specifically, for the problem of optimizing the liquid welfare in Combinatorial Auctions with submodular bidders, we obtain a universally truthful randomized $O(\log m)$ -approximate mechanism, where m is the number of items, by adapting the mechanism of Krysta and Vöcking (2012).

Additionally, motivated by large market assumptions often used in mechanism design, we introduce a notion of competitive markets and show that in such markets, liquid welfare can be approximated within a constant factor by a randomized universally truthful mechanism. Finally, in the Bayesian setting, we obtain a truthful $O(1)$ -approximate mechanism for the case where bidder valuations are generated as independent samples from a known distribution, by adapting the results of Feldman, Gravin and Lucier (2014).

1 Introduction

Imagine that you are a social planner wanting to auction-off the seats of a local stadium at an extremely wealthy neighborhood (i.e., people have *no* budget constraints for the seats) for a big concert. As a social planner, your goal is to allocate the seats in a way that maximizes (or, at least, approximates as closely as possible) the happiness of the people interested in these seats. However, different people have different seat preferences; some people are happy with two consecutive seats anywhere in the stadium, and some might want a whole row. Phrased in mechanism design language, this is a *Combinatorial Auction*, where you seek to optimize the *social welfare* by a truthful mechanism. Combinatorial Auctions, like the one above, appear in many AI-centric contexts (e.g., spectrum auctions, network routing

auctions (Hershberger and Suri 2001), airport time-slot auctions (Rassenti, Smith, and Bulfin 1982), etc.) and have been a central topic in the study of Multi-Agent Systems. They have also experienced a recent interest in the AI community with works employing ML algorithms to overcome standard complexity problems (e.g., (Brero and Lahaie 2018; Brero, Lubin, and Seuken 2018)).

As if this problem was not hard enough to solve, imagine that you find out two unfortunate events; the stadium is in fact at a working-middle class neighborhood (i.e., people *do* have budget constraints) and your boss is concerned about the effect of these budget constraints on the potential revenue. Now, the objective function should balance between the willingness and the ability of the people to pay for their seats. Motivated by usual discrepancies between the auction participants' ability and willingness to pay, (Dobzinski and Paes Leme 2014) introduced the notion of *liquid welfare*, which is the minimum of an agent's budget and valuation for a bundle of goods. As such, maximizing the liquid welfare achieves a reasonable compromise between social efficiency and potential for revenue extraction (which is constrained by the budgets).

Problem Definition. More formally, a Combinatorial Auction (CA) consists of a set U of m items to be allocated to n bidders. Each bidder i has a valuation function $v_i : 2^U \rightarrow \mathbb{R}_{\geq 0}$. Valuation functions, v , are assumed to be non-decreasing, i.e., $v(S) \leq v(T)$, for all $S \subseteq T \subseteq U$, and normalized $v(\emptyset) = 0$. For the objective of social welfare (SW), the goal is to compute a partitioning $\mathcal{S} = (S_1, \dots, S_n)$ of the set of items, U , that maximizes $v(\mathcal{S}) = \sum_{i=1}^n v_i(S_i)$. For the objective of liquid welfare (LW), we assume that each bidder i also has a budget $B_i \in \mathbb{R}_{\geq 0}$ and the liquid welfare that can be extracted from agent i for each set of items $S \subseteq U$ is $\bar{v}_i(S) = \min\{v_i(S), B_i\}$ ¹. Under this objective, the goal is to compute a partitioning $\mathcal{S} = (S_1, \dots, S_n)$ of U that maximizes $\bar{v}(\mathcal{S}) = \sum_{i=1}^n \bar{v}_i(S_i)$.

We focus on CAs with submodular or XOS bidders. A set function $v : 2^U \rightarrow \mathbb{R}_{\geq 0}$ is *submodular* if for every $S, T \subseteq U$, $v(S) + v(T) \geq v(S \cap T) + v(S \cup T)$ and *subadditive* if $v(S) + v(T) \geq v(S \cup T)$. A set function v is *XOS* (a.k.a. *fractionally subadditive*, see (Feige 2009)) if there ex-

*The names of the authors are in alphabetical order.
Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹Slightly abusing the terminology, we refer to $\bar{v}_i(S)$ as agent i 's *liquid valuation*.

ist additive functions $w_k : 2^U \rightarrow \mathbb{R}_{\geq 0}$ such that for every $S \subseteq U$, $v(S) = \max_k \{w_k(S)\}$. The class of submodular functions is a proper subset of the class of XOS functions, which in turn is a proper subset of the class of subadditive functions.

Since bidder valuations have exponential size, a polynomial (in m and n) algorithm must have oracle access to them. A *value query* specifies a set $S \subseteq U$ and receives the value $v(S)$. A *demand query*, denoted by $\text{DQ}(v, U, \vec{p})$, specifies a valuation function v , a set U of available items and a price p_j for each available item $j \in U$, and receives the set (or bundle) $S \subseteq U$ maximizing $v(S) - \sum_{j \in S} p_j$, i.e., the set of available items that maximizes bidder’s utility at these prices. For brevity, we often write $p(S) = \sum_{j \in S} p_j$ to denote the price of a bundle S . Demand queries are strictly more powerful than value queries. Value queries can be simulated by polynomially many demand queries, and in terms of communication cost, demand queries are exponentially stronger than value queries (Blumrosen and Nisan 2009). Our mechanisms are polynomial-time, given access to demand oracles, which in general can be NP-hard to compute.

1.1 Previous Work on Social Welfare

Truthful maximization of SW in CAs with submodular or XOS bidders has been a central problem in Algorithmic Mechanism Design, with many powerful results. Due to space restrictions, we only discuss results most relevant to our work. While discussing previous work below, we assume XOS bidders and polynomial-time randomized truthful mechanisms that approximate the SW, by accessing valuations through demand queries, unless mentioned otherwise.

In the *worst-case* setting, where we do not make any further assumptions on bidder valuations, (Dobzinski, Nisan, and Schapira 2006) presented the first truthful mechanism with a non-trivial approximation guarantee of $O(\log^2 m)$. (Dobzinski 2007) improved the approximation ratio to $O(\log m \log \log m)$ for the more general class of subadditive valuations. Subsequently, (Krysta and Vöcking 2012) provided an elegant randomized online mechanism that achieves an approximation ratio of $O(\log m)$ for XOS valuations. (Dobzinski 2016) broke the logarithmic barrier for XOS valuations, by providing an approximation guarantee of $O(\sqrt{\log m})$. We highlight that accessing valuations through demand queries is essential for these strong positive results. (Dobzinski 2011) proved that any truthful mechanism for submodular CAs with approximation ratio better than $m^{\frac{1}{2}-\epsilon}$ must use *exponentially* many value queries.

In the *Bayesian* setting, bidder valuations are drawn as independent samples from a known distribution. (Feldman, Gravin, and Lucier 2014) showed how to obtain item prices that provide a constant approximation ratio for XOS valuations. These results were significantly extended and strengthened in the recent work of (Düetting et al. 2017), and a (truly) polynomial algorithm was provided as well.

1.2 Intuition, Main Ideas, and Contribution

Our aim is to extend these results to the objective of LW. To this end, we exploit the fact that most of the mech-

anisms above follow a simple pattern: first, by exploring either part of the instance in (Krysta and Vöcking 2012) or the knowledge about the valuation distribution in (Feldman, Gravin, and Lucier 2014), the mechanism computes appropriate (a.k.a. *supporting*) prices for all items. Then, these prices are “posted” to the bidders, who arrive one-by-one and select their utility-maximizing bundle, through a demand query, from the set of available items (see Algorithm 1).

The technical intuition behind the high level approach above is nicely explained in (Dobzinski 2016, Section 1.2). Let $\mathcal{O} = (O_1, \dots, O_n)$ be an optimal solution for the SW (in fact, any constant factor approximation suffices). The supporting price of item i in \mathcal{O} is $q_i = w_k(\{j\})$, where w_k is the additive valuation determining the value $v_i(O_i)$ (recall that valuation functions are XOS). Intuitively, q_i is how much item j contributes to the social welfare of \mathcal{O} . Then, a price of $p_j = q_j/2$ for each item j is appropriate in the sense that a constant approximation to $v(\mathcal{O})$ can be obtained by letting the bidders arrive one-by-one, in an arbitrary order, and allocating to each bidder i her utility maximizing bundle, chosen from the set of available items by a demand query (see (Dobzinski 2016, Lemma 4.2)).

Hence, approximating the SW by demand queries boils down to computing such prices p_j . In the Bayesian setting, prices p_j can be obtained by drawing n samples from the valuation distribution and computing the expected contribution of each item j to a constant factor approximation of the optimal allocation (see Section 3 and Lemma 3.4 in (Feldman, Gravin, and Lucier 2014)). Similarly, the idea of estimating the contribution of the items would work under some market uniformity assumption, as the one introduced in Definition 5.1. In the worst-case setting, if we assume integral and polynomially-bounded valuations (i.e., that $\max_i \{v_i(U)\} \leq m^d$, for some constant d), a uniform price for all items selected at random from $1, 2, 4, 8, \dots, 2^{d \log m}$ results in an logarithmic approximation ratio. (Krysta and Vöcking 2012) show how to estimate supporting prices online, by combining binary search and randomized rounding. Importantly, as long as each bidder does not affect the prices offered to her, this general approach results in (randomized, universally) truthful mechanisms.

Towards extending the above approach and results to the LW, our first observation (Lemma 3.1) is that if a valuation function v is submodular (resp. XOS), then the corresponding liquid valuation function $\bar{v} = \min\{v, B\}$ is also submodular (resp. XOS). Then, one can directly use the mechanisms of e.g., (Krysta and Vöcking 2012; Dobzinski 2016; Feldman, Gravin, and Lucier 2014) with valuation functions $\bar{v} = \min\{v, B\}$ and demand queries of the form: $\text{DQ}(\min\{v, B\}, U, \vec{p})$ (i.e., wrt. the liquid valuation of the bidders) and obtain the same approximation guarantees but now for the LW. However, the resulting mechanisms are no longer truthful; bidders still seek to maximize their *utility* (i.e., value minus price) from the bundle that they get, subject to their budget constraint, rather than their *liquid utility* (i.e., liquid value minus price). Specifically, given a set of items U available at prices p_j , $j \in U$, a budget-constrained bidder i wants to receive the bundle $S_i =$

$\arg \max_{S \subseteq U} \{v_i(S) - p(S) \mid p(S) \leq B_i\}$, and might not be happy with the bundle $S'_i = \arg \max_{S \subseteq U} \{\bar{v}_i(S) - p(S)\}$ computed by the demand query for the liquid valuation².

To restore truthfulness, we replace demand queries with *budget-constrained demand queries*. A budget-constrained demand query, denoted by $\text{BCDQ}(v, U, \vec{p}, B)$, specifies a valuation function v , a set of available items U , a price p_j for each $j \in U$ and a budget B , and receives the set $S \subseteq U$ maximizing $v(S) - p(S)$, subject to $p(S) \leq B$, i.e., the set of available items that maximizes bidder’s utility subject to her budget constraint.

To establish the approximation ratio, we first observe that the fact that liquid valuations are XOS suffices for estimating supporting prices, as in previous work on the SW. Additionally, we show that the bundles allocated by $\text{BCDQ}(v, U, \vec{p}, B)$ approximately satisfy the efficiency guarantees on the liquid welfare and the liquid utility of the allocated bundles (see Lemma 3.3). Specifically, we observe that the approximation guarantees of mechanisms for the SW mostly follow from the fact that a demand query $\text{DQ}(v, U, \vec{p})$ guarantees that for the allocated bundle S and for any $T \subseteq U$, (i) $v(S) - p(S) \geq v(T) - p(T)$, and (ii) $v(S) \geq v(T) - p(T)$. In Lemma 3.3, we show that a budget-constrained demand query, $\text{BCDQ}(v, U, \vec{p}, B)$, guarantees that for the allocated bundle S and any $T \subseteq U$, (i) $2\bar{v}(S) - p(S) \geq \bar{v}(T) - p(T)$, and (ii) $\bar{v}(S) \geq \bar{v}(T) - p(T)$. Using this property, we can prove the equivalent of (Dobzinski 2016, Lemma 4.2) and also the approximation guarantees of the mechanisms in (Krysta and Vöcking 2012), (Feldman, Gravin, and Lucier 2014) but for the LW.

Contribution. Formalizing the intuition above, we obtain (i) a randomized universally truthful mechanism that approximates the LW within a factor of $O(\log m)$ (Section 4), and (ii) a posted-price mechanism that approximates the LW within a constant factor when bidder valuations are drawn as independent samples from a known distribution (Section 6). Both mechanisms assume XOS bidder valuations; the former is based on the mechanism of (Krysta and Vöcking 2012) and the latter on the mechanism of (Feldman, Gravin, and Lucier 2014).

Motivated by large market assumptions often used in Algorithmic Mechanism Design (see e.g., (Borgs et al. 2005; Eden, Feldman, and Vardi 2017; Lu and Xiao 2017) and the references therein), we introduce a competitive market assumption in Section 5. Competitive Markets are closer to practice, since they stand in between the stochastic and the worst-case settings, in terms of the assumptions made. The

²For a concrete example, consider a bidder with budget $B = 2$ and two items a and b available at prices $p_a = 2$ and $p_b = 1$. Assume that the bidder’s valuation function is $v(\{a\}) = v(\{a, b\}) = 10$, $v(\{b\}) = 2$ (and therefore, her liquid valuation is $\bar{v}(\{a\}) = \bar{v}(\{b\}) = \bar{v}(\{a, b\}) = 2$). The bidder wants to get item a at price 2, which gives her utility 8. However, the demand query for her liquid valuation function \bar{v} allocates item b , which gives her a utility of 1. Clearly, in this example, the bidder would have incentive to misreport her preferences to the demand query oracle.

main idea is that when there is an abundance of bidders, even if we remove a random half of them, the optimal LW does not decrease by much. Then, computing supporting prices for all items based on a randomly chosen half of the bidders, and offering these prices through budget-constrained demand queries to the other half, yields a universally truthful mechanism that approximates LW within a constant factor (Theorem 5.5).

Conceptually, in this work, we present a general approach through which known truthful approximations to the SW, that access valuations through demand queries, can be adapted so that they retain truthfulness and achieve similar approximation guarantees for the LW. The important properties required are that liquid valuation functions \bar{v} belong to the same class as valuation functions v (proven for submodular, XOS and subadditive valuations), and that the efficiency guarantees of budget-constrained demand queries on liquid welfare and liquid utility are similar to the corresponding efficiency guarantees of standard demand queries for liquid valuations (proven for all classes of valuations functions). Indeed, applying this approach to the mechanism of (Dobzinski 2016), we obtain a universally truthful mechanism that approximates the LW for CAs with XOS bidders within a factor of $O(\sqrt{\log m})$ (the details are omitted due to space constraints). Similarly, we can take advantage of the improved results of (Düetting et al. 2017) in the Bayesian setting. All the missing proofs can be found in the full version of the paper on (Fotakis, Lotidis, and Podimata 2018).

1.3 Previous Work on Liquid Welfare

Liquid welfare was introduced as an efficiency measure for auctions when bidders are budget constrained in (Dobzinski and Paes Leme 2014) (since it was known that getting any non-trivial approximation for the SW in these cases is impossible) and it corresponds to the optimal revenue an omniscient seller could extract from the set of the bidders, had he known their valuations and their budgets. Moreover, (Dobzinski and Paes Leme 2014) proved a $O(\log n)$ (resp. $(\log^2 n)$)-approximation to the optimal LW for the case of a single divisible item and submodular (resp. subadditive) bidders. (Dobzinski and Paes Leme 2014) and (Lu and Xiao 2015) proved that the optimal LW can be approximated truthfully within constant factor for a single divisible good, additive bidder valuations and public budgets. Closer to our setting, (Lu and Xiao 2017) provided a truthful mechanism that achieves a constant factor approximation to the LW for multi-item auctions with divisible auctions, under a large market assumption. Under similar large market assumptions, (Eden, Feldman, and Vardi 2017) obtained mechanisms that approximate the optimal revenue within a constant factor for multi-unit online auctions with divisible and indivisible items, and a mechanism that achieves a constant approximation to the optimal LW for general valuations over divisible items. However, prior to our work, there was no work on approximating the LW in CAs (in fact, that was one of the open problems in (Dobzinski and Paes Leme 2014)).

Our work is remotely related to the literature of *budget feasible mechanism design*, a topic introduced by (Singer 2010) and studied in e.g., (Dobzinski, Papadimitriou, and

Singer 2011; Chen, Gravin, and Lu 2011; Bei et al. 2012; Balkanski and Hartline 2016; Wu, Liu, and Li 2018). Budget feasible mechanism design focuses on payment optimization in reverse auctions, a setting almost orthogonal to the setting we consider in this work.

2 Notation and Preliminaries

The problem and most of the terminology and the notation are introduced in Section 1. In this section, we introduce some additional notation required for the technical part.

We use $\mathbb{E}[X]$ to denote the expectation of a random variable X and $\mathbb{P}[E]$ to denote the probability of an event E . Let OPT (resp. $\overline{\text{OPT}}$) denote the optimal SW (resp. LW)³. For some $\rho > 1$, which may depend on n and m , we say that a mechanism is ρ -approximate for the optimal SW (resp. LW) if it produces an allocation \mathcal{S} with $\rho \cdot v(\mathcal{S}) \geq \text{OPT}$ (resp. $\rho \cdot \bar{v}(\mathcal{S}) \geq \overline{\text{OPT}}$).

Let a social choice function $f : \bar{V}^n \rightarrow A$, which maps the set of liquid valuation functions of the bidders, $\bar{V} : V \times B$, to an allocation, A , and a payment scheme $q = (q_1, \dots, q_n)$ for this allocation. A deterministic mechanism is defined by the pair (f, q) . Our mechanisms in this work are going to be *randomized*, i.e., they are probability distributions over *deterministic* mechanisms. The incentives desiderata for randomized mechanisms are usually either *universal truthfulness* (when for all the deterministic mechanisms, the bidders' dominant strategy is truthfulness) or *truthfulness in expectation* (Dobzinski, Fu, and Kleinberg 2010; Dughmi, Roughgarden, and Yan 2011) (when bidders' expected utilities are maximized under truthful reporting of their private information). In this work, we are focusing on the former, stronger notion; the one of *universal truthfulness*, under the bidders' budget constraints.

Definition 2.1 (Universal Truthfulness under Budget Constraints). *Let (\tilde{f}, \tilde{q}) be a randomized mechanism over a set of deterministic mechanisms $\{(f^1, q^1), \dots, (f^k, q^k)\}$. Mechanism (\tilde{f}, \tilde{q}) is universally truthful if for all $i \in [n], \kappa \in [k]$ and for any v'_i and any B'_i , such that $q^\kappa(v'_i, v_{-i}) \leq B'_i$ and $q^\kappa(v_i, v_{-i}) \leq B_i$, it holds that:*

$$\begin{aligned} & v_i(f^\kappa(v_i, B_i, v_{-i}, B_{-i})) - q^\kappa(v_i, B_i, v_{-i}, B_{-i}) \\ & \geq v_i(f^\kappa(v'_i, B'_i, v_{-i}, B_{-i})) - q^\kappa(v'_i, B'_i, v_{-i}, B_{-i}) \end{aligned}$$

3 Approach

First, we show that if the bidder valuations are submodular (resp. XOS, subadditive), then their liquid valuations are submodular (resp. XOS, subadditive) as well.

Lemma 3.1. *Let v be a non-negative monotone submodular (resp. XOS, subadditive) function. Then, for any $B \in \mathbb{R}_{\geq 0}$, $\bar{v} = \min\{v, B\}$ is also monotone submodular (resp. XOS, subadditive).*

In Algorithm 1, we present a universally truthful (since the prices offered to each bidder do not depend on her declaration and demand queries maximize bidders' utility) mechanism, which is a simplified version of the mechanism in

Algorithm 1 Core Mechanism

- 1: Fix an ordering π of bidders and set $U_1 = U$.
 - 2: Set initial prices for the items: $\bar{p}^{(1)} = (p_1^{(1)}, \dots, p_m^{(1)})$.
 - 3: **for** each bidder $i = 1, \dots, n$ according to π **do**
 - 4: Let $S_i = \text{DQ}(v_i, U_i, \bar{p}^{(i)})$,
 - 5: With probability q , allocate S_i to i and set $U_{i+1} = U_i \setminus S_i$. Otherwise, set $U_{i+1} = U_i$.
 - 6: Update item prices to $\bar{p}^{(i+1)} = (p_1^{(i+1)}, \dots, p_m^{(i+1)})$.
 - 7: **end for**
-

(Krysta and Vöcking 2012) for approximating SW in CAs. Since for the LW, bidders have budgets, we replace the demand queries $\text{DQ}(v_i, U_i, \bar{p}^{(i)})$ in line 4 with budget constrained demand queries $\text{BCDQ}(v_i, U_i, \bar{p}^{(i)}, B_i)$ ⁴. As explained in Section 1.2, Algorithm 1 with BCDQs remains universally truthful for budget-constrained bidders.

Lemma 3.2 (Truthfulness of BCDQs). *For budget-constrained bidders, Algorithm 1 with BCDQs in line 4, is universally truthful.*

The lemma follows directly from Definition 2.1. Nevertheless, universal truthfulness is not our sole desideratum; in each of the three settings analyzed in the following sections, we show why mechanisms similar in spirit to Algorithm 1 with BCDQs, yield good approximation guarantees for the LW. Before the setting-specific analyses, we relate the efficiency of BCDQ to the efficiency of standard DQs for liquid valuations.

Lemma 3.3. *Let $S \subseteq U$ be the set allocated by the BCDQ for a bidder with valuation v and budget B . Then, for every other $T \subseteq U$, the following hold: i) $\bar{v}(S) \geq \bar{v}(T) - p(T)$ ii) $2\bar{v}(S) - p(S) \geq \bar{v}(T) - p(T)$.*

Proof. We will prove each claim of the lemma separately. For claim i), notice that if $p(T) > B$, then the right hand side of the inequality will be negative and thus, the inequality trivially holds. So, we will focus on the case where $p(T) \leq B$. We distinguish the following cases:

1. ($\bar{v}(S) = v(S)$ and $\bar{v}(T) = v(T)$) Hence, $B \geq v(T)$. Bundle T was considered at the time of the query and yet, the query returned set S . Thus: $\bar{v}(S) \geq \bar{v}(S) - p(S) = v(S) - p(S) \geq v(T) - p(T) = \bar{v}(T) - p(T)$.
2. ($\bar{v}(S) = B$ and $\bar{v}(T) = B$) Then, the inequality trivially holds since: $B \geq B - p(T)$ and prices are non-negative.
3. ($\bar{v}(S) = B$ and $\bar{v}(T) = v(T)$) The inequality holds since: $B \geq B - p(T) \geq v(T) - p(T) = \bar{v}(T) - p(T)$.
4. ($\bar{v}(S) = v(S)$ and $\bar{v}(T) = B$) Hence, $B \leq v(T)$. Bundle T was considered at the time of the query and yet, the query returned set S . Thus, $\bar{v}(S) \geq \bar{v}(S) - p(S) = v(S) - p(S) \geq v(T) - p(T) \geq B - p(T) = \bar{v}(T) - p(T)$.

⁴In all our mechanisms, if budgets are larger than the valuations of the allocated bundles, the mechanism with BCDQ behaves identically to the mechanism with DQ (i.e., revenue and SW are not affected by the change in the objective.)

³The instance is clear from the context.

Algorithm 2 KV-Mechanism for Liquid Welfare

- 1: Fix an ordering π of bidders and set $U_1 = U$.
 - 2: Set initial prices $p_1^{(1)} = \dots = p_m^{(1)} = \frac{L}{4m}$.
 - 3: **for** each bidder $i = 1, \dots, n$ according to π **do**
 - 4: Let $S_i = \text{BCDQ}(v_i, U_i, \vec{p}^{(i)}, B_i)$
 - 5: With probability q , allocate $R_i = S_i$ to i and set $U_{i+1} = U_i \setminus S_i$. Otherwise, set $U_{i+1} = U_i, R_i = \emptyset$.
 - 6: Update prices $\forall j \in S_i: p_j^{(i+1)} = 2p_j^{(i)}$.
 - 7: **end for**
-

This concludes our proof for claim i).

For claim ii), notice that since S is the set received from the BCDQ, then it is *affordable*: $\bar{v}(S) \geq p(S)$. Adding this inequality to the inequality of claim i), we have that: $2\bar{v}(S) - p(S) \geq \bar{v}(T) - p(T)$. \square

4 Worst-Case Setting

For the worst-case instances, adapting appropriately our Core Mechanism, we present Algorithm 2 (based again, on the mechanism of (Krysta and Vöcking 2012)). The only difference is that budget-constrained bidders in Algorithm 2 are restricted to using BCDQs, instead of DQs, thus making the mechanism universally truthful (see Section 3). Resembling the analysis of (Krysta and Vöcking 2012), we show that for $1/q = \Theta(\log m)$, Algorithm 2 achieves an approximation ratio of $O(\log m)$ for the LW. First, we note that parameter⁵ L is selected so that there exists only one bidder whose liquid valuation for U (weakly) exceeds it.

Theorem 4.1. *Algorithm 2 is universally truthful and for $q = 1/\Theta(\log m)$, achieves an approximation ratio of $O(\log m)$ for the LW.*

We present a series of Lemmas that will lead us naturally to the proof of the Theorem. Let $S = (S_1, \dots, S_n)$ and $\mathcal{R} = (R_1, \dots, R_n)$ the provisional and the final allocation of Algorithm 2 respectively. First, we provide two useful bounds on $\bar{v}(S)$. We find it important to also discuss an overselling variant of Algorithm 2. In the *Overselling* variant, allow us to assume that for Step 5 of Algorithm 2, $q = 1$ (i.e., S_i is allocated to bidder i with certainty) and $U_{i+1} = U_i = U$ (thus the name of the variant). The *Overselling* variant allocates at most $k = \log(4m) + 2$ copies of each item and collects a liquid welfare within a constant factor of the optimal LW. To see that, observe that for $q = 1$, after allocating $k - 1$ copies of some item j , j 's price becomes $\frac{L}{4m} 2^{\log(4m)+1} = 2L$. Then, there is only one agent with liquid valuation larger than L who can get a copy of j .

Lemma 4.2. *Let p_j denote the final price of each item j . Then, for any sets $U_1, \dots, U_n \subseteq U$ of items available when the bidders arrive, Algorithm 2 with $q = 1$ satisfies $\bar{v}(S) \geq \sum_{j \in U} p_j - L/4$.*

Proof. Since bidders are individually rational and do not violate their budget constraints, for every bidder i it holds that

⁵ L can be computed with standard techniques, as explained in (Krysta and Vöcking 2012).

$B_i \geq \sum_{j \in S_i} p_j^{(i)}$ and $v_i(S_i) \geq \sum_{j \in S_i} p_j^{(i)}$. The rest of the proof is identical to the proof of (Krysta and Vöcking 2012, Lemma 2) for $b = 1$. Specifically, let $\ell_j^{(i)}$ be the number of copies of item j allocated just before bidder i arrives, and let ℓ_j be the total number of copies of item j allocated by Algorithm 2 with $q = 1$. Then, using the fact that $p_j = L \cdot 2^{\ell_j} / 4m$: $\bar{v}(S) \geq \sum_{i=1}^n \sum_{j \in S_i} p_j^{(i)} = \frac{L}{4m} \sum_{i=1}^n \sum_{j \in S_i} 2^{\ell_j^{(i)}} = \frac{L}{4m} \sum_{j \in U} (2^{\ell_j} - 1) = \sum_{j \in U} p_j - L/4$ \square

Lemma 4.3. *For sets $U_1 = \dots = U_n \subseteq U$, the Overselling variant of Algorithm 2 with $q = 1$ satisfies $\bar{v}(S) \geq \overline{\text{OPT}} - \sum_{j \in U} p_j$.*

Proof. Let $\mathcal{O} = (O_1, \dots, O_n)$ be the optimal allocation. From Lemma 3.3, we get that for each bidder i , $\bar{v}(S_i) \geq \bar{v}(O_i) - \sum_{j \in O_i} p_j^{(i)} \geq \bar{v}(O_i) - \sum_{j \in O_i} p_j$, where we use that the final price of each item is the largest one. Summing over all bidders, we have that $\bar{v}(S) \geq \bar{v}(\mathcal{O}) - \sum_{i=1}^n \sum_{j \in O_i} p_j \geq \overline{\text{OPT}} - \sum_{j \in U} p_j$, where the last inequality uses the fact that the optimal solution is feasible and thus, each item is allocated at most once in \mathcal{O} . \square

Lemma 4.4. *The Overselling variant of Algorithm 2 with $q = 1$ allocates at most $\log(4m) + 2$ copies of each item and computes an allocation S with liquid welfare $\bar{v}(S) \geq \frac{3}{8} \overline{\text{OPT}}$.*

Proof. Follows directly from Lemma 4.2, Lemma 4.3 and the fact that $\overline{\text{OPT}} \geq L$. \square

Of course, the allocation S in Lemma 4.4 is infeasible, since it allocates a logarithmic number of copies of each item. The remedy is to use an allocation probability $q = 1/\Theta(\log m)$. For such values of q , we can plug in the proof of (Krysta and Vöcking 2012, Lemma 6) (which just uses that the valuation functions are fractionally subadditive) and show that for each agent i and for all $A \subseteq U$, $\mathbb{E}[\bar{v}_i(A \cap U_i)] \geq \bar{v}_i(A)/2$. We are now ready to conclude the proof of Theorem 4.1.

Lemma 4.5. *For Algorithm 2 with $q^{-1} = 4(\log(4m) + 1)$, it holds that $\mathbb{E}[\bar{v}(S)] \geq \overline{\text{OPT}}/8$ and $\mathbb{E}[\bar{v}(\mathcal{R})] \geq q \overline{\text{OPT}}/8$.*

Proof. Let $\mathcal{O} = (O_1, \dots, O_n)$ be the optimal allocation. For each bidder i , Lemma 3.3 implies that the response S_i of BCDQ satisfies $\bar{v}_i(S_i) \geq \bar{v}_i(O_i \cap U_i) - \sum_{j \in O_i \cap U_i} p_j^{(i)}$, for any U_i resulted from the outcome of the random coin flips. Therefore, $\mathbb{E}[\bar{v}_i(S_i)] \geq \mathbb{E}[\bar{v}_i(O_i \cap U_i)] - \mathbb{E}[\sum_{j \in O_i \cap U_i} p_j^{(i)}]$. By the choice of q , for any bidder i , $\mathbb{E}[\bar{v}_i(O_i \cap U_i)] \geq \bar{v}_i(O_i)/2$. Then, working with the expectations as in the proofs of Lemma 4.2, Lemma 4.3 and Lemma 4.4, we can show that $\mathbb{E}[\bar{v}(S)] \geq \overline{\text{OPT}}/8$. Finally, one can use linearity of expectation and show that $\mathbb{E}[\bar{v}(\mathcal{R})] = q \mathbb{E}[\bar{v}(S)]$. The details are omitted, due to lack of space, and can be found in (Krysta and Vöcking 2012, Lemma 4). \square

5 Competitive Market

(Borgs et al. 2005) were the first ones to define a *budget dominance parameter* that corresponded to the ratio of the maximum budget of all the bidders to the value of the optimum SW in the context of multi-unit auctions with budget-constrained bidders. More recently, (Eden, Feldman, and Vardi 2017) and (Lu and Xiao 2017) used similar notions of budget dominance⁶ (termed *large market assumptions*) as a means to achieve constant factor approximation to the LW in multi-unit auctions and auctions with divisible items respectively. However, for the case of *divisible* items, it is clear that the definition of a large market used in the previous cases, becomes almost void⁷. Below, we first introduce our definition of Competitive Markets for indivisible goods and then, show how one can obtain a constant factor approximation of the optimal LW, when bidders have XOS liquid valuations.

Definition 5.1 ((ε, δ) - Competitive Market). *Let $0 \leq \varepsilon < 2$ and a constant $0 \leq \delta \leq 1/2$. A market is called (ε, δ) - Competitive Market, if for any randomly removed set of bidders, \mathbb{S} , with cardinality $n/2$, then for the remaining set of bidders, \mathbb{T} , it holds that:*

$$\mathbb{P}\left[\overline{OPT}_{\mathbb{T}} \geq \left(1 - \frac{\varepsilon}{2}\right) \cdot \overline{OPT}\right] \geq 1 - \delta \quad (1)$$

where by $\overline{OPT}_{\mathbb{T}}$ we denote the optimal LW achieved by bidders in set \mathbb{T} .

Proposition 5.2. *In an (ε, δ) - Competitive Market, let $\mathbb{S} \subseteq [n]$ be randomly chosen s.t. $|\mathbb{S}| = \frac{n}{2}$ and let $\mathbb{T} = [n] \setminus \mathbb{S}$. Then:*

$$\mathbb{P}\left\{\left\{\overline{OPT}_{\mathbb{T}} \geq \left(1 - \frac{\varepsilon}{2}\right) \overline{OPT}\right\} \cap \left\{\overline{OPT}_{\mathbb{S}} \geq \left(1 - \frac{\varepsilon}{2}\right) \overline{OPT}\right\}\right\} \geq 1 - 2\delta$$

Proof. Let $X_{\mathbb{S}}$ the event that $\overline{OPT}_{\mathbb{S}} \geq \left(1 - \frac{\varepsilon}{2}\right) \overline{OPT}$ and $X_{\mathbb{T}}$ the event that $\overline{OPT}_{\mathbb{T}} \geq \left(1 - \frac{\varepsilon}{2}\right) \overline{OPT}$. Then, we have:

$$\mathbb{P}[X_{\mathbb{S}} \cap X_{\mathbb{T}}] = 1 - \mathbb{P}[\overline{X}_{\mathbb{S}} \cup \overline{X}_{\mathbb{T}}] \geq 1 - 2\delta$$

where the inequality follows from the Union Bound. \square

We are now ready to state our Competitive Market mechanism that will be used for approximating the optimal LW. We note here that the greedy algorithm \mathcal{A} is due to (Lehmann, Lehmann, and Nisan 2006).

As usual, we denote $\mathcal{S} = (S_1, \dots, S_n)$ the final allocation from mechanism presented in Algorithm 3. Valuations of bidders are XOS (and so are the liquid valuations (Lemma 3.1)); let a_i be the maximizing clause of S_i in the liquid valuation \bar{v}_i of bidder i . Since a_i 's are additive, for each bidder i and $j \in S_i$ let $q_j = a_i(\{j\})$. Notice that $\sum_{i \in [n]} \bar{v}(S_i) = \sum_{j \in \cup_{i \in [n]} S_i} q_j$. We denote by $\overline{OPT}_{\mathbb{T}} = \sum_{j \in U} q_j^{\mathbb{T}}$, where $q_j^{\mathbb{T}}$ is the contribution of item j in $\overline{OPT}_{\mathbb{T}}$. We divide the set of all items U into two sets; the set of *competitive* items, denoted by \mathcal{C} and the set of *non-competitive* items, denoted by $\bar{\mathcal{C}} = \mathcal{M} \setminus \mathcal{C}$. The following lemma upper bounds the contribution of non-competitive items in the optimal solution.

⁶Namely, $\forall i \in [n] : B_i \leq \overline{OPT}/(mc)$, with c a large constant.

⁷See the full version for a discussion.

Algorithm 3 Competitive Market (CM) Algorithm

- 1: Divide the bidders into sets \mathbb{S}, \mathbb{T} uniformly at random, s.t., $|\mathbb{S}| = \frac{n}{2} = |\mathbb{T}|$.
 - 2: Run the greedy algorithm \mathcal{A} for bidders in \mathbb{S} and denote the solution obtained by $\mathcal{A}^{\mathbb{S}}$.
 - 3: **for** $j \in U$ **do**
 - 4: Set $p_j = \frac{1}{2\beta} \bar{v}(\mathcal{A}_j^{\mathbb{S}})$, where $\beta > 1$ is a constant
 - 5: **end for**
 - 6: Fix an internal ordering of bidders in \mathbb{T} , π , and set $U_1 = U$.
 - 7: **for** each bidder $i \in \mathbb{T}$ arriving according to π **do**
 - 8: Let $S_i = \text{BCDQ}(v_i, U_i, \bar{p})$.
 - 9: Set $U_{i+1} = U_i \setminus S_i$.
 - 10: **end for**
-

Lemma 5.3. *Let $\mathcal{C} = \{j | q_j^{\mathbb{T}} > \frac{\bar{v}(\mathcal{A}_j^{\mathbb{S}})}{\beta}\}$ for constant $\beta > 1$. Then, $\sum_{j \in \bar{\mathcal{C}}} q_j^{\mathbb{T}} \leq \frac{\varepsilon}{2(\beta-1)} \overline{OPT}$ and $\sum_{j \in \mathcal{C}} q_j^{\mathbb{T}} \geq \frac{\beta(2-\varepsilon)-2}{2(\beta-1)} \overline{OPT}$.*

Proof. From Definition 5.1, it holds with constant probability (w.c.p) that: $\overline{OPT} \geq \sum_{j \in \mathcal{C}} q_j^{\mathbb{T}} + \sum_{j \in \bar{\mathcal{C}}} q_j^{\mathbb{T}} = \sum_{j \in U} q_j^{\mathbb{T}} \geq \left(1 - \frac{\varepsilon}{2}\right) \cdot \overline{OPT}$. Let $\mathbb{S}_{\bar{\mathcal{C}}} \subseteq \mathbb{S}$ be the set of the bidders that are allocated the non-competitive items from the greedy algorithm \mathcal{A} when running on set \mathbb{S} . Then, in the augmented set $\mathbb{T} \cup \mathbb{S}_{\bar{\mathcal{C}}}$, there exists an allocation \mathcal{Q}^8 with liquid valuation,

$$\bar{v}(\mathcal{Q}) \geq \sum_{j \in \mathcal{C}} q_j^{\mathbb{T}} + \sum_{j \in \bar{\mathcal{C}}} \bar{v}(\mathcal{A}_j^{\mathbb{S}}) \quad (2)$$

and therefore we have w.c.p:

$$\begin{aligned} \overline{OPT} &\geq \bar{v}(\mathcal{Q}) \geq \sum_{j \in \mathcal{C}} q_j^{\mathbb{T}} + \sum_{j \in \bar{\mathcal{C}}} \bar{v}(\mathcal{A}_j^{\mathbb{S}}) \geq \sum_{j \in \mathcal{C}} q_j^{\mathbb{T}} + \beta \sum_{j \in \bar{\mathcal{C}}} q_j^{\mathbb{T}} \\ &\geq \left(1 - \frac{\varepsilon}{2}\right) \overline{OPT} + (\beta - 1) \sum_{j \in \bar{\mathcal{C}}} q_j^{\mathbb{T}} \end{aligned}$$

Re-arranging the latter and using the fact that $\sum_{j \in \mathcal{C}} q_j + \frac{\varepsilon}{2(\beta-1)} \overline{OPT} \geq \sum_{j \in U} q_j^{\mathbb{T}} \geq \left(1 - \frac{\varepsilon}{2}\right) \overline{OPT}$, then, for the items in \mathcal{C} it holds w.c.p that: $\sum_{j \in \mathcal{C}} q_j^{\mathbb{T}} \geq \frac{\beta(2-\varepsilon)-2}{2(\beta-1)} \overline{OPT}$. \square

In the next Lemma, we prove a lower bound on the contribution of competitive items to the solution obtained by the greedy algorithm, with respect to \overline{OPT} .

Lemma 5.4. $\sum_{j \in \mathcal{C}} \bar{v}(\mathcal{A}_j^{\mathbb{S}}) \geq \frac{2(\beta-1)-\varepsilon(3\beta-1)}{4(\beta-1)} \overline{OPT}$.

Proof. Combining Inequality (2) and Lemma 5.3 we get that $\sum_{j \in \bar{\mathcal{C}}} \bar{v}(\mathcal{A}_j^{\mathbb{S}}) \leq \frac{\beta\varepsilon}{2(\beta-1)} \overline{OPT}$. Algorithm \mathcal{A} provides a 2-approximation to the optimal LW of set \mathbb{S} (Lehmann, Lehmann, and Nisan 2006), so w.c.p we have:

⁸Allocation \mathcal{Q} is realized by allocating all items in \mathcal{C} to bidders in \mathbb{T} that also had them in the $\overline{OPT}_{\mathbb{T}}$ allocation and all items in $\bar{\mathcal{C}}$ to the bidders in $\mathbb{S}_{\bar{\mathcal{C}}}$ that had them in the allocation of the greedy \mathcal{A} . The claim is completed by submodularity.

$\sum_{j \in \mathcal{C}} \bar{v}(\mathcal{A}_j^{\mathbb{S}}) + \sum_{j \in \bar{\mathcal{C}}} \bar{v}(\mathcal{A}_j^{\mathbb{S}}) \geq \frac{1}{2} \overline{\text{OPT}}_{\mathbb{S}} \geq \frac{1-\frac{\epsilon}{2}}{2} \overline{\text{OPT}}$. Combining the last two equations, we get the result. \square

Theorem 5.5. *The CM Algorithm is universally truthful and achieves, in expectation, a constant approximation to the optimal LW, i.e., $\mathbb{E}[\bar{v}(\mathcal{S})] \geq (1 - 2\delta) \cdot \frac{2(\beta-1) - \epsilon \cdot (3\beta-1)}{16\beta(\beta-1)} \overline{\text{OPT}}$.*

Proof. Since the bidders that control the prices being posted belong to set \mathbb{S} and they never get any item, it is their (weakly) dominant strategy to report their valuations and their budgets truthfully. Furthermore, the bidders that are buying under the said posted prices belong to set \mathbb{T} and they make BCDQs, which we shown to be truthful. Finally, the bidders are uniformly at random split at sets \mathbb{S} and \mathbb{T} .

For each item $j \in \mathcal{C}$ we have $q_j^{\mathbb{T}} > \bar{v}(\mathcal{A}_j^{\mathbb{S}})/\beta$. Therefore, there exists an allocation for bidders in \mathbb{T} and items in \mathcal{C} that is supported by prices p_1, \dots, p_m , where $p_j = \bar{v}(\mathcal{A}_j^{\mathbb{S}})/\beta$. Thus, a modification of (Dobzinski 2016, Lemma 4.2) implies that if we set $p'_j = p_j/2$, for each $j \in \mathcal{C}$, and run a fixed price auction in \mathbb{T} with prices p'_1, \dots, p'_m , we get that $\bar{v}(\mathcal{S}) \geq \sum_{j \in \mathcal{C}} p_j/4$. Using the latter, along with the prices of the items, we have that $\bar{v}(\mathcal{S}) = \frac{1}{4\beta} \sum_{j \in \mathcal{C}} \bar{v}(\mathcal{A}_j^{\mathbb{S}}) \geq \frac{2(\beta-1) - \epsilon \cdot (3\beta-1)}{16\beta(\beta-1)} \overline{\text{OPT}}$, where the last inequality is due to Lemma 5.4. Thus, we conclude that $\mathbb{E}[\bar{v}(\mathcal{S})] \geq (1 - 2\delta) \frac{2(\beta-1) - \epsilon \cdot (3\beta-1)}{16\beta(\beta-1)} \overline{\text{OPT}}$. \square

6 Bayesian Setting

The Bayesian Setting offers a great middle ground between the unstructured worst-case instances and the very structured Competitive Markets. In this setting, let $\vec{v} = (v_1, \dots, v_n)$ be a profile of bidder valuations and $\mathbf{B} = (B_1, \dots, B_n)$ a profile of bidder budgets. Assume that the bidders' valuations are drawn independently from distributions $\mathcal{V}_1, \dots, \mathcal{V}_n$ and the budgets from distributions $\mathcal{B}_1, \dots, \mathcal{B}_n$. For simplicity, let us assume that their liquid valuations are drawn independently from distributions $\mathcal{D}_1, \dots, \mathcal{D}_n$. We will denote by $\mathcal{D} = \mathcal{D}_1 \times \dots \times \mathcal{D}_n$ the product distribution where liquid valuations profiles, $\vec{v} = (v_1, \dots, v_n)$, are independently drawn from.

We are going to show that the results presented in (Feldman, Gravin, and Lucier 2014) can be extended for budget-constrained bidders. Specifically, we are going to show that, if liquid valuations are fractionally subadditive, then we can create appropriate prices such that, when presented to the bidders in a posted-price mechanism and bidders are making BCDQs, then we can obtain universally truthful constant-factor approximation mechanisms for the LW in Bayesian CAs. Our Lemma 6.2 establishes the existence of such appropriately scaled prices. The key component activating our results is that instead of reasoning about the *utility* achieved from the bundle purchased by bidder i (as received by the BCDQ), we instead have to use Lemma 3.3. We also note that using our techniques one could even achieve the better approximation guarantees presented by (Dütting et al. 2017). Their analysis is significantly more complex, however, and we omit it in the interest of space.

Theorem 6.1. *Let distribution \mathcal{D} over XOS liquid valuation profiles be given via a sample access to \mathcal{D} . Suppose*

that for every $\vec{v} \sim \mathcal{D}$, we have: i) black-box access to a LW maximization algorithm, ALG⁹ ii) an XOS value query oracle (for liquid valuations sampled from \mathcal{D})¹⁰. Then, for any $\epsilon > 0$, we can compute item prices in $\text{POLY}(m, n, 1/\epsilon)$ time such that, for any bidder arrival order, the expected liquid welfare of the posted price mechanism is at least $\frac{1}{4} \mathbb{E}_{\vec{v} \sim \mathcal{D}}[\bar{v}(\text{ALG}(\vec{v}))] - \epsilon$, where by $\text{ALG}(\vec{v})$ we denote the solution produced by algorithm ALG.

Lemma 6.2. *Given a distribution \mathcal{D} over XOS liquid valuations, let \vec{p} be the price vector s.t. $p_j = \frac{1}{2} \mathbb{E}_{\vec{v} \sim \mathcal{D}}[LW_j(\vec{v})]$. Let \vec{p}' be any price vector such that $|p'_j - p_j| < \delta$ for all $j \in [m]$. Then, for any arrival order, π , bidders buying bundles by making BCDQs under prices \vec{p}' results in expected liquid welfare at least $\frac{1}{4} \mathbb{E}_{\vec{v} \sim \mathcal{D}}[\bar{v}(\text{ALG}(\vec{v}))] - \frac{m\delta}{2}$.*

7 Conclusion

In real-life auctions, bidders are *always* constrained by budgets, which we tend to overlook due to the technical difficulties that they add. The role of budgets in welfare/revenue optimization is amplified in CAs, where bidders have richer valuations and hence, studying budgeted CAs is a step towards bridging the gap between the theory on truthful mechanism design for CAs and constraints faced in practice. In this work, we showed how the liquid welfare can be approximated in CAs where bidders are budget-constrained in three settings: worst-case, Competitive Markets and stochastic. The most meaningful question that arises from our work (apart, of course, from the ever existent one of lowering the approximation guarantee in worst-case instances) is related to the competitive markets. We conjecture that the condition that we provide can be made even weaker, and leave it to future research.

Finally, our results can also be used to extend a variety of already known results in CAs without budgets, to CAs with budget-constrained bidders. For example, Lemma 3.3 (with some changes in the constants of (Dütting and Kesselheim 2017)) implies a constant factor approximation for best response dynamics in XOS CAs with budgeted bidders, that apply after a single round of bid updates.

8 Acknowledgments

We are thankful to Yiling Chen for her most valuable feedback in an initial draft of this work, as well as the anonymous reviewers for their helpful comments and suggestions. Part of this work was done while Chara Podimata was visiting Yahoo Research, NY. This work is partially supported by the National Science Foundation under grant CCF-1718549 and the Harvard Data Science Initiative. Any opinions, findings, conclusions, or recommendations expressed here are those of the authors alone.

⁹ALG can be any algorithm that provides a $O(1)$ -approximation to the optimal LW, since we do not care about incentives (access to ALG will only happen for ghost samples).

¹⁰An XOS value oracle takes as input a set T and returns the corresponding additive representative function for the set T , i.e., an additive function $A_i(\cdot)$, such that (i) $\bar{v}_i(S) \geq A_i(\hat{S})$ for any $\hat{S} \subset [m]$ and (ii) $\bar{v}_i(T) = A_i(T)$.

References

- Balkanski, E., and Hartline, J. D. 2016. Bayesian budget feasibility with posted pricing. In *Proceedings of the 25th International Conference on World Wide Web*, 189–203. International World Wide Web Conferences Steering Committee.
- Bei, X.; Chen, N.; Gravin, N.; and Lu, P. 2012. Budget feasible mechanism design: from prior-free to bayesian. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, 449–458. ACM.
- Blumrosen, L., and Nisan, N. 2009. On the computational power of demand queries. *SIAM Journal on Computing* 39(4):1372–1391.
- Borgs, C.; Chayes, J.; Immorlica, N.; Mahdian, M.; and Saberi, A. 2005. Multi-unit auctions with budget-constrained bidders. In *Proceedings of the 6th ACM conference on Electronic commerce*, 44–51. ACM.
- Brero, G., and Lahaie, S. 2018. A bayesian clearing mechanism for combinatorial auctions. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, 941–948.
- Brero, G.; Lubin, B.; and Seuken, S. 2018. Combinatorial auctions via machine learning-based preference elicitation. In *IJCAI*, 128–136.
- Chen, N.; Gravin, N.; and Lu, P. 2011. On the approximability of budget feasible mechanisms. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, 685–699. Society for Industrial and Applied Mathematics.
- Dobzinski, S., and Paes Leme, R. 2014. Efficiency guarantees in auctions with budgets. In *International Colloquium on Automata, Languages, and Programming*, 392–404. Springer.
- Dobzinski, S.; Fu, H.; and Kleinberg, R. 2010. Truthfulness via proxies. *arXiv preprint arXiv:1011.3232*.
- Dobzinski, S.; Nisan, N.; and Schapira, M. 2006. Truthful randomized mechanisms for combinatorial auctions. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, 644–652. ACM.
- Dobzinski, S.; Papadimitriou, C. H.; and Singer, Y. 2011. Mechanisms for complement-free procurement. In *Proceedings of the 12th ACM conference on Electronic commerce*, 273–282. ACM.
- Dobzinski, S. 2007. Two randomized mechanisms for combinatorial auctions. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*. Springer. 89–103.
- Dobzinski, S. 2011. An impossibility result for truthful combinatorial auctions with submodular valuations. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, 139–148. ACM.
- Dobzinski, S. 2016. Breaking the logarithmic barrier for truthful combinatorial auctions with submodular bidders. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, 940–948. ACM.
- Düetting, P.; Feldman, M.; Kesselheim, T.; and Lucier, B. 2017. Prophet inequalities made easy: Stochastic optimization by pricing non-stochastic inputs. In *Foundations of Computer Science (FOCS), 2017 IEEE 58th Annual Symposium on*, 540–551. IEEE.
- Dughmi, S.; Roughgarden, T.; and Yan, Q. 2011. From convex optimization to randomized mechanisms: toward optimal combinatorial auctions. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, 149–158. ACM.
- Dütting, P., and Kesselheim, T. 2017. Best-response dynamics in combinatorial auctions with item bidding. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, 521–533. Society for Industrial and Applied Mathematics.
- Eden, A.; Feldman, M.; and Vardi, A. 2017. Online random sampling for budgeted settings. In *International Symposium on Algorithmic Game Theory*, 29–40. Springer.
- Feige, U. 2009. On maximizing welfare when utility functions are subadditive. *SIAM Journal on Computing* 39(1):122–142.
- Feldman, M.; Gravin, N.; and Lucier, B. 2014. Combinatorial auctions via posted prices. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, 123–135. SIAM.
- Fotakis, D.; Lotidis, K.; and Podimata, C. 2018. A bridge between liquid and social welfare in combinatorial auctions with submodular bidders. <https://arxiv.org/pdf/1809.01803.pdf>.
- Hershberger, J., and Suri, S. 2001. Vickrey prices and shortest paths: What is an edge worth? In *Foundations of Computer Science, 2001. Proceedings. 42nd IEEE Symposium on*, 252–259. IEEE.
- Krysta, P., and Vöcking, B. 2012. Online mechanism design (randomized rounding on the fly). In *International Colloquium on Automata, Languages, and Programming*, 636–647. Springer.
- Lehmann, B.; Lehmann, D.; and Nisan, N. 2006. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior* 55(2):270–296.
- Lu, P., and Xiao, T. 2015. Improved efficiency guarantees in auctions with budgets. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, 397–413. ACM.
- Lu, P., and Xiao, T. 2017. Liquid welfare maximization in auctions with multiple items. In *International Symposium on Algorithmic Game Theory*, 41–52. Springer.
- Rassenti, S. J.; Smith, V. L.; and Bulfin, R. L. 1982. A combinatorial auction mechanism for airport time slot allocation. *The Bell Journal of Economics* 402–417.
- Singer, Y. 2010. Budget feasible mechanisms. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, 765–774. IEEE.
- Wu, W.; Liu, X.; and Li, M. 2018. Budget-feasible procurement mechanisms in two-sided markets. In *IJCAI*, 548–554.