Simple Mechanisms for Subadditive Buyers via Duality

YANG CAI
Yale University
and
MINGFEI ZHAO
McGill University

A central problem in mechanism design is how to design simple and approximately revenue-optimal auctions in multi-item multi-buyer settings. Prior to our work, all results only apply to cases where the buyers' valuations are linear over the items.We unify and improve all previous results, as well as generalize the results to accommodate non-linear valuations [Cai and Zhao 2017]. In particular, we prove that a simple, deterministic and Dominant Strategy Incentive Compatible (DSIC) mechanism, namely, the sequential posted price with entry fee mechanism, achieves a constant fraction of the optimal revenue among all randomized, Bayesian Incentive Compatible (BIC) mechanisms, when buyers' valuations are XOS (a superclass of submodular valuations) over independent items. If the buyers' valuations are subadditive over independent items, the approximation factor degrades to $O(\log m)$, where m is the number of items. We obtain our results by first extending the Cai-Devanur-Weinberg duality framework to derive an effective benchmark of the optimal revenue for subadditive buyers, and then developing new analytic tools that combine concentration inequality of subadditive functions, prophet-inequality type of arguments, and a novel decomposition of the benchmark.

Categories and Subject Descriptors: J.4 [**Social and Behavioral Science**]: Economics

General Terms: Algorithms, Mechanism Design, Economics, Theory,

Additional Key Words and Phrases: Bayesian Mechanism Design, Simple Mechanism, XOS Valuation, Subadditive Valuation, Revenue, Approximation, Duality

1. INTRODUCTION

In mechanism design, we usually focus on obtaining or approximating the optimal mechanism. Clearly, the quality of the mechanism with respect to the designer's objective is crucial. However, perhaps one should also pay equal attention to the simplicity of a mechanism. When facing a complicated mechanism, participants may be confused by the rules and thus unable to optimize their actions and react in unpredictable ways instead. Such behavior can lead to undesirable outcomes and destroy the performance of the mechanism. An ideal mechanism would be optimal and simple. For revenue maximization in multi-item settings, we now know that, even in fairly basic cases, the optimal mechanisms suffer many undesirable properties including randomization, non-monotonicity, and others [Rochet and Chone 1998; Thanassoulis 2004; Pavlov 2011; Hart and Nisan 2013; Hart and Reny 2012; Briest et al. 2010; Daskalakis et al. 2013; 2014]. To move forward, we need to understand the tradeoff between optimality and simplicity.

Authors' email addresses: `yang.cai@yale.edu`, `mingfei.zhao@mcgill.ca`

A central problem on this front is how to design simple and approximately revenue-optimal mechanisms in multi-item settings. We have witnessed a lot of progress in the past few years. For instance, when buyers have unit-demand valuations, we can use a sequential posted price mechanism to approximates the optimal revenue thanks to a line of work initiated by Chawla et al. [Chawla et al. 2007; Chawla et al. 2010; Chawla et al. 2015; Cai et al. 2016]. When buyers are additive, we know that either selling the items separately or running a VCG mechanism with per buyer entry fee approximates the optimal revenue thanks to a series of work initiated by Hart and Nisan [Hart and Nisan 2012; Cai and Huang 2013; Li and Yao 2013; Babaioff et al. 2014; Yao 2015; Cai et al. 2016]. Recently, Chawla and Miller [Chawla and Miller 2016] generalized these two lines of work to matroid rank functions[1]. For subadditive valuations beyond matroid rank functions, we only knew how to handle a single buyer [Rubinstein and Weinberg 2015] prior to our work[2]. It is a major open problem to design a simple and approximately optimal mechanism for multiple subadditive buyers.

1.1 Our Results

We have made major progress on this open problem in [Cai and Zhao 2017]. We unify and strengthen all the results mentioned above via an extension of the duality framework proposed by Cai et al. [Cai et al. 2016]. Moreover, we show that even when there are multiple buyers with XOS valuation functions, a simple, deterministic and Dominant Strategy Incentive Compatible (DSIC) mechanism, namely the sequential posted price with entry fee mechanism (see Mechanism 1) , suffices to extract a constant fraction of the optimal Bayesian Incentive Compatible (BIC) revenue[3]. For subadditive valuations, our approximation ratio degrades to $O(\log m)$. Please see Table I for the approximation ratios we obtained for different valuation classes.

Theorem 1. *There exists a sequential posted price with entry fee mechanism (SPEM) that achieves a constant fraction of the optimal BIC revenue in multi-item settings, when the buyers' valuation distributions are XOS over independent items. When the buyers' valuation distributions are subadditive over independent items, our mechanism achieves at least* $\Omega\left(\frac{1}{\log m}\right)$ *of the optimal BIC revenue, where* m *is the number of items. Note that our mechanism is deterministic and DSIC.*

SPEMs are clearly DSIC, as any buyer can first figure out how much surplus she can obtain by winning her favorite bundle among the available items, and she

[1] Here is a hierarchy of the valuation functions. additive & unit-demand $\subseteq$ matroid rank $\subseteq$ constrained additive & submodular $\subseteq$ XOS $\subseteq$ subadditive. A function is constrained additive if it is additive up to some downward closed feasibility constraints.

[2] All results mentioned above assume that the buyers' valuation distributions are over independent items. For additive and unit-demand valuations, this means a buyer's values for the items are independent. The definition is generalized to subadditive valuations by Rubinstein and Weinberg [Rubinstein and Weinberg 2015].

[3] A mechanism is Bayesian Incentive Compatible (BIC) if it is in every buyer's interest to tell the truth, assuming that all other buyers' reported their values. A mechanism is Dominant Strategy Incentive Compatible (DSIC) if it is in every buyer's interest to tell the truth no matter what reports the other buyers make.

Mechanism 1 Sequential Posted Price with Entry Fee Mechanism (SPEM)

Require: There are n buyers and m items. p_{ij} is the price for buyer i to purchase item j and $\delta_i(\cdot)$ is buyer i's entry fee function.

$S \leftarrow [m]$
for $i \in [n]$ **do**
    Show buyer i the set of available items S, and define entry fee as $\delta_i(S)$.
    if buyer i pays the entry fee $\delta_i(S)$ **then**
        i receives her favorite bundle $S_i^* \subseteq S$ and pays $\sum_{j \in S_i^*} p_{ij}$.
        $S \leftarrow S \setminus S_i^*$.
    else
        i gets nothing and pays 0.
    end if
end for

		Additive or Unit-demand	Matroid Rank	Constrained Additive	XOS	Subadditive
Single Buyer	Previous	6 or 4	31.1*	31.1	338*	338
	This Paper	-	11*	11	40*	40
Multiple Buyers	Previous	8 or 24	133	?	?	?
	This Paper	-	70*	70	268	$O(\log m)$

* The result is implied by another result for a more general setting.

Table I. Comparison of approximation ratios between previous and current work.

only accepts the entry fee when the surplus is larger. Due to the simplicity and strong strategyproofness, SPEMs have been widely adopted in the real world, for example, by Amazon Prime and Costco. We have indeed obtained a stronger result, that is, we only need to use the following two special types of SPEMs to obtain the approximation result in Theorem 1: the *rationed sequential posted price mechanisms* and the *anonymous sequential posted price with entry fee mechanisms.*

Rationed sequential posted price mechanisms (RSPM). The entry fee is 0 for every buyer, and every buyer is allowed to purchase at most one item.

Anonymous sequential posted price with entry fee mechanisms (ASPE). The item prices are anonymous, that is every buyer faces the same collection of item prices $\{p_j\}$. The mechanism may have nonzero entry fee.

THEOREM 2. *Either a* rationed sequential posted price mechanism *or an* sequential anonymous posted price with entry fee mechanism *can achieve a constant fraction (or* $\Omega\left(\frac{1}{\log m}\right)$*) of the optimal BIC revenue when the buyers' valuation distributions are XOS (or subadditive) over independent items.*

In our specification of a SPEM (Mechanism 1), we assume that the buyers arrive in lexicographical order. This assumption is not necessary, and Theorem 2 holds for arbitrary arrival order. In other words, we can construct a set of item prices $\{p_{ij}\}_{i\in[n],j\in[m]}$ and a collection of entry fee functions $\{\delta_i(\cdot)\}_{i\in[n]}$ that are oblivious to the arrival order, such that the corresponding SPEM guarantees the same constant factor approximation as in Theorem 2 for any arrival order.

2. OUR MODEL

We consider revenue maximization in combinatorial auctions with n independent buyers and m heterogenous items. Each buyer has a valuation that is **subadditive over independent items** (see Definition 1). We denote buyer i's type t_i as $\langle t_{ij}\rangle_{j=1}^m$, where t_{ij} is buyer i's private information about item j. Note that this information need not be a scalar. We refer the readers to Example 1 for more concrete examples.

For each i, j, we assume t_{ij} is drawn independently from the distribution D_{ij}. Let $D_i = \times_{j=1}^m D_{ij}$ be the distribution of buyer i's type and $D = \times_{i=1}^n D_i$ be the distribution of the type profile. We use T_{ij} (or T_i, T) and f_{ij} (or f_i, f) to denote the support and density function of D_{ij} (or D_i, D). When buyer i's type is t_i, her valuation for a set of items S is denoted by $v_i(t_i, S)$.

Definition 1. *[Rubinstein and Weinberg 2015] For every buyer i, whose type is drawn from a product distribution $F_i = \prod_j F_{ij}$, her distribution $\mathcal{V}_i$ of valuation function $v_i(t_i, \cdot)$ is* ***subadditive over independent items*** *if:*

—$v_i(\cdot,\cdot)$ ***has no externalities***, *i.e., for each $t_i \in T_i$ and $S \subseteq [m]$, $v_i(t_i, S)$ only depends on $\langle t_{ij}\rangle_{j\in S}$, formally, for any $t_i' \in T_i$ such that $t_{ij}' = t_{ij}$ for all $j \in S$, $v_i(t_i', S) = v_i(t_i, S)$.*

—$v_i(\cdot,\cdot)$ ***is monotone***, *i.e., for all $t_i \in T_i$ and $U \subseteq V \subseteq [m]$, $v_i(t_i, U) \leq v_i(t_i, V)$.*

—$v_i(\cdot,\cdot)$ ***is subadditive***, *i.e., for all $t_i \in T_i$ and $U, V \subseteq [m]$, $v_i(t_i, U \cup V) \leq v_i(t_i, U) + v_i(t_i, V)$.*

We use $V_i(t_{ij})$ to denote $v_i(t_i, \{j\})$, as it only depends on t_{ij}. When $v_i(t_i, \cdot)$ is XOS (or constrained additive) for all $t_i \in T_i$, we say $\mathcal{V}_i$ is XOS (or constrained additive) over independent items.

We focus on the following valuation classes.

Definition 2. *We define several classes of valuations formally. Let t be the type and $v(t, S)$ be the value for bundle $S \in [m]$.*

—***Constrained Additive:*** *$v(t, S) = \max_{R\subseteq S, R\in\mathcal{I}} \sum_{j\in R} v(t, \{j\})$, where $\mathcal{I} \subseteq 2^{[m]}$ is a downward closed set system over the items specifying the feasible bundles. In particular, when $\mathcal{I} = 2^{[m]}$, the valuation is an* ***additive function****; when $\mathcal{I} = \{\{j\} \mid j \in [m]\}$, the valuation is a* ***unit-demand function****; when $\mathcal{I}$ is a matroid, the valuation is a* ***matroid rank function****. An equivalent way to represent any constrained additive valuations is to view the function as additive but the buyer is only allowed to receive bundles that are feasible, i.e., bundles in $\mathcal{I}$. To ease notations, we interpret t as an m-dimensional vector $(t_1, t_2, \cdots, t_m)$ such that $t_j = v(t, \{j\})$.*

—***XOS/Fractionally Subadditive:*** *$v(t, S) = \max_{i\in[K]} v^{(i)}(t, S)$, where K is some finite number and $v^{(i)}(t, \cdot)$ is an additive function for any $i \in [K]$.*

—***Subadditive:*** *$v(t, S_1 \cup S_2) \leq v(t, S_1) + v(t, S_2)$ for any $S_1, S_2 \subseteq [m]$.*

Note that when a buyer has constrained additive valuation, her value for a bundle is still linear over the items as long as the bundle is feasible. Such linearity no longer holds when we consider more general valuations such as XOS or subadditive functions. Next, we provide a few examples of various valuation distributions which are over independent items (Definition 1):

EXAMPLE 1. *$t = \{t_j\}_{j\in[m]}$ where t is drawn from $\prod_j D_j$,*

—*Additive: t_j is the value of item j. $v(t,S) = \sum_{j\in S} t_j$.*

—*Unit-demand: t_j is the value of item j. $v(t,S) = \max_{j\in S} t_j$.*

—*Constrained Additive: t_j is the value of item j. $v(t,S) = \max_{R\subseteq S, R\in\mathcal{I}} \sum_{j\in R} t_j$.*

—*XOS/Fractionally Subadditive: $t_j = \left\{t_j^{(k)}\right\}_{k\in[K]}$ encodes all the possible values associated with item j, and $v(t,S) = \max_{k\in[K]} \sum_{j\in S} t_j^{(k)}$.*

3. BACKGROUND ON DUALITY

In this section, we provide some basic background on the Cai-Devanur-Weinberg duality framework [Cai et al. 2016], which will be helpful for the readers to understand the major contributions of [Cai and Zhao 2017]. Readers who are familiar with the duality framework can consider skipping this section.

We first formulate the revenue maximization problem as an LP (see Figure 1). For all buyers i and types $t_i \in T_i$, we use $p_i(t_i)$ as the interim price paid by buyer i and $\sigma_{iS}(t_i)$ as the interim probability of receiving the exact bundle S. To ease the notation, we use a special type $\varnothing$ to represent the choice of not participating in the mechanism. More specifically, $\sigma_{iS}(\varnothing) = 0$ for any S and $p_i(\varnothing) = 0$. Now a Bayesian IR (BIR) constraint is simply another BIC constraint: for any type t_i, buyer i will not want to lie to type $\varnothing$. We let $T_i^+ = T_i \cup \{\varnothing\}$.

Next, we take the partial Lagrangian dual of the LP by lagrangifying the BIC constraints. Let $\lambda_i(t_i, t_i')$ be the Lagrange multiplier associated with the BIC constraint that when buyer i's true type is t_i she will not prefer to lie to another type t_i'. As shown in [Cai et al. 2016], the dual solution has finite value if and only if the dual variables λ_i form a valid flow for every buyer i. The reason is that the payments $p_i(t_i)$ are unconstrained variables, therefore the corresponding coefficients in the dual problem must be 0 in order for the dual to have finite value. It turns out that when all these coefficients are 0, the dual variables λ satisfy $f_i(t_i) + \sum_{t_i'\in T_i} \lambda_i(t_i', t_i) - \sum_{t_i'\in T_i^+} \lambda_i(t_i, t_i') = 0$ for every buyer i and every type t_i of hers. If we consider a graph with

—super source s and a super sink $\varnothing$, along with a node t_i for every type $t_i \in T_i$,

—an edge from s to t_i with flow $f_i(t_i)$ for all $t_i \in T_i$,

—an edge from t_i to t_i' with flow $\lambda_i(t_i, t_i')$ for all $t_i \in T_i$ and $t_i' \in T_i^+$ (including the super sink),

then the above equation makes sure that the inflow equals to the outflow at every node t_i. Therefore, we only consider dual variables λ that correspond to a flow. For every flow λ, we can define a virtual valuation function.

Definition 3. *(Virtual Value Function) For each flow λ, we define a corresponding virtual value function $\Phi^{\lambda}(\cdot)$, such that for every buyer i, every type $t_i \in T_i$ and every set $S \subseteq [m]$,*

$$\Phi_i^{\lambda}(t_i, S) = v_i(t_i, S) - \frac{1}{f_i(t_i)} \sum_{t_i' \in T_i} \lambda_i(t_i', t_i) \left(v_i(t_i', S) - v_i(t_i, S) \right).$$

Variables:

— $p_i(t_i)$, for all buyers i and types $t_i \in T_i$, denoting the expected price paid by buyer i when reporting type t_i over the randomness of the mechanism and the other buyers' types.

— $\sigma_{iS}(t_i)$, for all buyers i, all bundles of items $S \subseteq [m]$, and types $t_i \in T_i$, denoting the probability that buyer i receives **exactly** the bundle S when reporting type t_i over the randomness of the mechanism and the other buyers' types.

Constraints:

— $\sum_{S \subseteq [m]} \sigma_{iS}(t_i) \cdot v_i(t_i, S) - p_i(t_i) \geq \sum_{S \subseteq [m]} \sigma_{iS}(t_i') \cdot v_i(t_i, S) - p_i(t_i')$, for all buyers i, and types $t_i \in T_i, t_i' \in T_i^+$, guaranteeing that the reduced form mechanism (σ, p) is BIC and Bayesian IR.

— $\sigma \in P(D)$, guaranteeing σ is feasible.

Objective:

— $\max \sum_{i=1}^{n} \sum_{t_i \in T_i} f_i(t_i) \cdot p_i(t_i)$, the expected revenue.

Fig. 1. A Linear Program (LP) for Revenue Optimization.

Duality implies that for any virtual valuation function Φ^{λ} and any BIC mechanism M, the revenue of M is upper bounded by the virtual welfare of M w.r.t. Φ^{λ}. The two quantities are equal if M is the optimal mechanism and λ is the optimal dual.

Theorem 3. *(Virtual Welfare $\geq$ Revenue [Cai et al. 2016]) For any flow λ and any BIC mechanism $M = (\sigma, p)$, the revenue of M is $\leq$ the virtual welfare of σ w.r.t. the virtual valuation $\Phi^{\lambda}(\cdot)$.*

$$\sum_{i=1}^{n} \sum_{t_i \in T_i} f_i(t_i) \cdot p_i(t_i) \leq \sum_{i=1}^{n} \sum_{t_i \in T_i} f_i(t_i) \sum_{S \subseteq [m]} \sigma_{iS}(t_i) \cdot \Phi_i^{\lambda}(t_i, S)$$

Let λ^* be the optimal dual variables and $M^* = (\sigma^*, p^*)$ be the revenue optimal BIC mechanism, then the expected virtual welfare with respect to Φ^{λ^*} under σ^* equals to the expected revenue of M^*.

3.1 Canonical Flow and Decomposition

Theorem 3 shows that every flow induces an upper bound for the optimal revenue, but which one should we use to generate our benchmark? If we choose the optimal flow, then the benchmark becomes too complex for us to analyze. If we choose a flow that is easy to analyze, then the induced benchmark may be too high and not within a constant factor of the optimal revenue. The challenge is how to strike the right balance between optimality and simplicity. The main contribution of [Cai et al. 2016] is to provide a canonical way of setting the flow that give rise to a close to optimal and mathematically analyzable benchmark for additive and unit-demand valuations. In this section, we use the setting of a single additive buyer as a running example to present this canonical flow/dual. For simplicity, we will assume the type space is a m-dimensional bounded integer lattice $\times_{j=1}^{m}[H]$. The readers will soon realize that the same flow can be easily extended to the general case.

As the buyer is additive, we will use a m-dimensional vector t to denote her type, and t_j is her value for item j. We partition the type set into m regions, where region R_j contains all types t such that $j = \text{argmax}_k\{t_k\}$ (break ties lexicographically). Consider the following implicitly defined flow (see Figure 2): for any type t in R_j, if $t' = (t_j - 1; t_{-j})$, i.e. the type that takes t's favorite item and decreases the value by 1 is still in R_j, then t sends all of its incoming flow into t' and it is the only type that sends flow into t'. If t' is not in R_j, then t sends all its incoming flow to the super sink.

Key property of the induced virtual valuation: so what Φ^λ does this flow induce? It is not immediate to see, but some inductive calculation yields that for every item j and every type t in region R_j: $\Phi^\lambda(t, S) = \sum_{k \in S \setminus \{j\}} t_k + \varphi_j(t_j) \cdot \mathbb{1}[j \in S]$, where $\varphi_j(t_j) = t_j - \frac{1-F_j(t_j)}{f_j(t_j)}$, Myerson's single-dimensional virtual value for D_j [Myerson 1981]. In other words, the virtual valuation is also additive over the items, where the virtual value of the buyer's favorite item j is just the Myerson's virtual value w.r.t. distribution D_j and the virtual value for any non-favorite item k is just its true value t_k. The corresponding virtual welfare is a strong enough benchmark for us to recover the result in [Babaioff et al. 2014]. To provide more intuition about why this benchmark is useful, consider a simpler flow where every type sends all of its inflow to the super sink directly. The induced virtual valuation is the same as the real valuation, and therefore the optimal social welfare is the benchmark, which can be infinitely times larger than the optimal revenue. The key property of the canonical flow is that it turns the value of the favorite item into the Myerson's virtual value, which substantially lowers the benchmark. As we will see soon, this property no longer holds for XOS/subadditive valuations. Restoring this key property is one of the major contributions of [Cai and Zhao 2017].

4. PROOF SKETCH AND MAIN CONTRIBUTIONS

In this section, we sketch the proof of our result. Along the way, we also discuss three major challenges that we faced and how we addressed them.

4.1 Extending the canonical flow to subadditive valuations

The first step of the proof is to find a suitable benchmark for the optimal revenue. Let's start with the canonical flow that we defined in Section 3.1 and see what goes wrong when the valuation is no longer linear over the items. We use a single subadditive bidder to illustrate the difficulty. We first divide the type space into m regions in a similar way, where region R_j contains all types t with $V(t_j) \geq V(t_k)$ for all $k \neq j$, i.e. winning item j is better than winning any other item. If the valuation is additive, the contribution to the virtual value function $\Phi^\lambda(t, S)$ from any type $t' \in R_j$ is $\lambda(t', t)(v(t', S) - v(t, S))$ which is either 0 when $j \notin S$ or

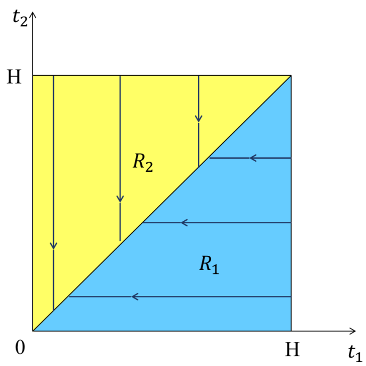

Fig. 2. An example of the canonical flow for an additive buyer with two items.

$\lambda(t', t)(t'_j - t_j)$. Importantly, the contribution does not depend on t_k for any $k \neq j$. This is the crucial property that allows us to replace the value of the favorite item by the corresponding Myerson's virtual value in the virtual valuation induced by the canonical flow $\Phi^{\lambda}(t, \cdot)$ for additive valuations. This property continues to hold for any valuation that is constrained additive. However, for XOS/subadditive valuations, $\lambda(t', t)\left(v(t', S) - v(t, S)\right)$ heavily depends on t_{-j}, even if $t'_{-j} = t_{-j}$. For example, when the valuation is XOS, the additive function that has attained the maximum value may be different under t and t'. As a result, the difference $v(t', S) - v(t, S)$ depends on t_{-j}. Unfortunately, this dependency completely destroys the structure of the induced benchmark and makes it impossible to analyze.

Valuation Relaxation. To overcome this difficulty, we take a different approach. Instead of directly studying the dual of the original problem, we first relax the valuation to $\tilde{v}(t, S) = v(t, S \backslash \{j\}) + V(t_j)$ if $t \in R_j$. The relaxation is inspired by [Rubinstein and Weinberg 2015] and can be viewed as "forcing" the valuation to be additive across the favorite item and all the other items. Next, we apply the canonical flow to the relaxed valuation. With some easy calculation, it is not hard to see that $\lambda(t', t)(\tilde{v}(t', S) - \tilde{v}(t, S))$ is either 0 or $\lambda(t', t)(V(t'_j) - V(t_j))$, and the induced benchmark resembles the appealing format in Section 3.1. However, this benchmark is only an upper bound of the optimal revenue for $\tilde{v}$. Since $\tilde{v}(t, S)$ dominates $v(t, S)$ for any t and S, it seems that the optimal revenue for $\tilde{v}$ should be at least as high as the optimal revenue for v. However, we have examples from [Hart and Reny 2012] showing that this intuitive revenue monotonicity does not always hold. To relate the optimal revenue under $\tilde{v}$ and v, we apply a different technique known as the ϵ-BIC to BIC reduction [Hartline and Lucier 2010; Hartline et al. 2011; Bei and Huang 2011; Daskalakis and Weinberg 2012] to show that the optimal revenue under $\tilde{v}$ is within a constant factor of the optimal revenue under v. Combining these two steps, we obtain a benchmark for subadditive valuations. Using our benchmark, we improve the approximation ratio for a single subadditive buyer from 338 [Rubinstein and Weinberg 2015] to 40.

4.2 An adaptive dual via ex-ante relaxation

So far, our discussion of the dual has focused on a single buyer setting. How do we handle multiple buyers? A natural idea is to set the dual variables λ_i for each buyer i equal to the canonical flow for a single buyer. Unfortunately, we can construct an example where the induced upper bound is unboundedly larger than the optimal revenue even when all buyers have additive valuations. The main reason that the canonical dual works in single buyer setting is because it converts the value of the favorite item into Myerson's virtual value, which substantially decreases the upper bound. If every buyer simply applies the single buyer canonical flow in the multi-buyer case, each buyer converts the value of her favorite item into Myerson's virtual value. However, this is a bad idea. Remember that the upper bound is the virtual welfare. Since it is possible that all buyers end up having the same favorite item, only the buyer who actually receives the favorite item can use her saving in the virtual valuation to reduce the upper bound. In such a scenario, a more effective way to decrease the upper bound will be to have buyers choose different items as their favorite items. This is exactly the reasoning behind the example and why applying the single buyer canonical dual to every buyer fails in multi-buyer settings. How do we design a canonical dual for multiple buyers? Note that the problem we face here can intuitively be thought of as the supply of the "favorite" items does not meet the demand. To balance the supply and demand, we use a basic economic approach – introducing prices.

We use $\boldsymbol{\beta} = \{\beta_{ij}\}_{i\in[n],j\in[m]} \in \mathbb{R}^{nm}_{\geq 0}$ to denote the prices[4], and we define a multi-buyer canonical flow $\lambda^{(\boldsymbol{\beta})}$ for every set of prices $\boldsymbol{\beta}$. For simplicity, we only describe the flow for additive valuations (see Figure 3 for an example), but it can be generalized to subadditive valuations by combining the valuation relaxation technique described in Section 4.1 [Cai and Zhao 2017][5]. Based on $\boldsymbol{\beta}$, we partition the type set T_i of each buyer i into $m+1$ regions: **(i)** $R_0^{(\boldsymbol{\beta}_i)}$ contains all types t_i such that $t_{ij} < \beta_{ij}$ for all $j \in [m]$. **(ii)** $R_j^{(\boldsymbol{\beta}_i)}$ contains all types t_i such that $t_{ij} - \beta_{ij} \geq 0$ and j is the smallest index in $\text{argmax}_k\{t_{ik} - \beta_{ik}\}$. Intuitively, if we view β_{ij} as the price of item j for buyer i, then $R_0^{(\boldsymbol{\beta}_i)}$ contains all types in T_i that cannot afford any item, and any $R_j^{(\boldsymbol{\beta}_i)}$ with $j > 0$ contains all types in T_i whose "favorite" item is j. For every type t_i in region $R_0^{(\boldsymbol{\beta}_i)}$, the flow goes directly to the super sink. For every region $R_j^{(\boldsymbol{\beta}_i)}$ with $j > 0$, the flow is similar to the single buyer canonical flow in region R_j. In particular, $\lambda_i^{(\boldsymbol{\beta})}(t_i, t_i') > 0$ only if t_i and t_i' are both in $R_j^{(\boldsymbol{\beta}_i)}$ and only differ in the j-th coordinate.

Every $\boldsymbol{\beta}$ induces a benchmark for the optimal revenue, but which one should we use? In [Cai et al. 2016], the choice of $\boldsymbol{\beta}$ was inspired by the prices in the VCG mechanism, and the corresponding optimal virtual welfare was used as the benchmark for the optimal revenue when the valuations are either additive or unit-demand. Unfortunately, the benchmark induced by this VCG inspired $\boldsymbol{\beta}$ becomes

[4]The $\boldsymbol{\beta}$ is only used to define the flow. It should not to be confused with the posted prices $\{p_{ij}\}_{i\in[n],j\in[m]}$ in the sequential posted price with entry fee mechanism.

[5]In [Cai and Zhao 2017], we apply the multi-buyer canonical flow on the relaxed valuation $\tilde{v}$ to obtain the benchmark.

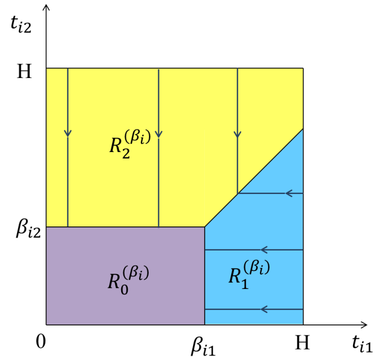

Fig. 3. An example of the canonical flow for multiple additive buyers with two items.

extremely complex and almost impossible to analyze, even when we only move slightly beyond additive or unit-demand valuations. In particular, the benchmark is already too involved to bound when the buyers' valuations are k-demand[6].

In [Cai and Zhao 2017], we propose a more flexible approach. In [Cai et al. 2016], the $\boldsymbol{\beta}$ and the corresponding flow $\lambda^{(\boldsymbol{\beta})}$ are chosen up front, then for every mechanism M, they use M's virtual welfare under the virtual valuation $\Phi^{\lambda^{(\boldsymbol{\beta})}}(\cdot)$ as the upper bound of the revenue of M. In [Cai and Zhao 2017], we take a more flexible approach to obtaining the upper bound. Instead of using a fixed $\boldsymbol{\beta}$, we tailor a special $\boldsymbol{\beta}(M)$ for every mechanism M, such that t_{ij} is above price $\beta(M)_{ij}$ with probability equal to $\frac{1}{2}$ of buyer i's ex-ante probability of winning item j in M[7]. The choice of $\boldsymbol{\beta}(M)$ is inspired by Chawla and Miller's ex-ante relaxation [Chawla and Miller 2016]. According to Theorem 3, M's virtual welfare under virtual valuation $\Phi^{\lambda^{\boldsymbol{\beta}(M)}}(\cdot)$ is still a valid upper bound of M's revenue. Since the virtual valuation is designed specifically for M, the induced virtual welfare is much easier to analyze. Indeed, the adaptiveness of the upper bound is crucial for our analysis when buyers have complex valuations such as XOS functions. We will point out exactly why this is critical in our proof in the next section.

4.3 New mechanism and analysis for the core

With the tools in Section 4.1 and 4.2, we can derive an upper bound of the optimal revenue for multiple subadditive buyers. The third and probably the most important contribution of our paper is a novel analysis of this upper bound and a new mechanism. Similar to prior work, we break the upper bound into three different terms – Single, Tail and Core – and bound them separately.

Both the Single and Tail can be covered by the revenue of a few RSPMs, and the analysis is relatively standard by now, so we will focus on the analysis of the Core. When the buyers are additive or unit-demand, the Core is not too difficult to analyze, but it soon becomes a monster once we go beyond these two

[6]The valuation is k-demand is a constrained additive function, where the buyer can enjoy at most k items.
[7]The particular choice of $\frac{1}{2}$ is not critical.

basic valuation functions. For example, the tools in [Cai et al. 2016] can handle additive or unit-demand buyers, but are insufficient to tackle the CORE even when the buyers have k-demand valuations, a very special case of matroid rank valuations. Rubinstein and Weinberg [Rubinstein and Weinberg 2015] showed how to approximate the CORE for a single subadditive buyer using grand bundling, but their approach is limited to a single buyer. Yao [Yao 2015] showed how to approximate the CORE for multiple additive buyers using a VCG with per buyer entry fee mechanism, but again it is unlikely his approach can be extended to even multiple k-demand buyers. Chawla and Miller [Chawla and Miller 2016] finally broke the barrier. They showed how to bound the CORE for matroid rank valuations using a sequential posted price mechanism by applying the *online contention resolution scheme (OCRS)* developed by Feldman et al. [Feldman et al. 2016]. The connection with OCRS is an elegant observation, and one might hope the same technique applies to more general valuations. Unfortunately, OCRS is only known to exist for special cases of downward closed constraints, and as we will explain, the approach by Chawla and Miller cannot yield any constant factor approximation for general constrained additive valuations.

As all results in the literature [Chawla et al. 2010; Yao 2015; Cai et al. 2016; Chawla and Miller 2016] only study matroid rank valuations, for simplicity, we will restrict our attention to a similar but more general class – the constrained additive valuations – in the comparison, but our approach also applies to XOS and subadditive valuations.

We provide a 100 feet overview of the approach by Chawla and Miller [Chawla and Miller 2016]. Essentially, all analyses prior to [Cai and Zhao 2017] follow the same path. The CORE is the optimal social welfare obtainable from a truncated version of v. In particular, the truncated valuation of a feasible set S is $v_i'(t_i, S) = \sum_{j\in S} t_{ij} \cdot \mathbb{1}[t_{ij} \leq \beta_{ij} + c_i]$, where $\{\beta_{ij}\}_{i\in[n],j\in[m]}$ is the set of prices we use in Section 4.2 to define the canonical dual, and $\{c_i\}_{i\in[n]}$ is a set of thresholds used to separate the CORE from the TAIL. We will not define $\{c_i\}_{i\in[n]}$, as it is not important for our discussion. We use $\boldsymbol{\beta}+\boldsymbol{c}$ to denote the set $\{\beta_{ij}+c_i\}_{i\in[n],j\in[m]}$. The analysis in [Chawla and Miller 2016] first separates the CORE into two parts: (i) the lower core: the welfare obtained from values below $\boldsymbol{\beta}$, and (ii) the upper core: the welfare obtained from values between $\boldsymbol{\beta}$ and $\boldsymbol{\beta}+\boldsymbol{c}$[8]. It is not hard to show that the upper core is upper bounded by the revenue of a two-part tariff mechanism[9]. In particular, only the revenue from the entry fees is used in the analysis to cover the upper core, while the revenue from the posted prices is ignored. Due to the choice of $\boldsymbol{\beta}$ (similar to the way we choose $\boldsymbol{\beta}(M)$ as described in Section 4.2), the lower core is upper bounded by $\sum_{i,j} \beta_{ij} \cdot \Pr_{t_{ij}}[t_{ij} \geq \beta_{ij}]$. When every buyer's feasibility constraint is a matroid, one can use the OCRS from [Feldman et al. 2016] to design a sequential

[8]In particular, if buyer i is awarded a bundle S that is feasible for her, the contribution to the lower core is $\sum_{j\in S} \min\{\beta_{ij}, t_{ij}\} \cdot \mathbb{1}[t_{ij} < \beta_{ij} + c_i]$ and the contribution to the upper core is $\sum_{j\in S} (t_{ij} - \beta_{ij})^+ \cdot \mathbb{1}[t_{ij} < \beta_{ij} + c_i]$

[9]A two-part tariff mechanism also has an entry fee component and a item pricing component, but it is different from SPEM. In particular, it charges all buyers the entry fees up front and then the buyers can enter the mechanism sequentially. In other words, a buyer needs to decide whether to pay the entry fee before she knows what items are still available, as she may not be the first to enter. As a result, the two-part tariff mechanism is only BIC but not DSIC.

posted price mechanism to approximate this quantity. However, we constructed an example where $\sum_{i,j} \beta_{ij} \cdot \Pr_{t_{ij}}[t_{ij} \geq \beta_{ij}]$ is $\Omega\left(\frac{\sqrt{m}}{\log m}\right)$ times larger than even the optimal social welfare when the buyers have general downward closed feasibility constraints. Readers can find the example in the full version of our paper [Cai and Zhao 2016]. Hence, such approach cannot yield any constant factor approximation for general constrained additive valuations.

We take a different path. Instead of decomposing the CORE and bound the two parts separately, we analyze it as a whole and use our new mechanism, the SPEM, to approximate it. Remember that for every mechanism M, we tailor a dual for M and use the corresponding virtual welfare under M's as the upper bound for M's revenue. The CORE part of the upper bound is M's social welfare w.r.t. the truncated valuation v'. Our goal is to show that for any mechanism M, we can design a SPEM to extract a constant fraction of M's social welfare under v' as revenue.

We provide some intuition as to why a SPEM can approximate the CORE well. First, consider a sequential posted price mechanism $\widetilde{\mathcal{M}}$. A key property of posted price mechanisms is that when buyer i's valuation is subadditive over independent items, her utility in the mechanism, which is her surplus from winning her favorite bundle among the unsold items, is also subadditive over independent items. If we can argue that her utility function under v'_i is α-Lipschitz with some small α, the concentration inequality for subadditive functions [Talagrand 1995; Schechtman 2003] allows us to set an entry fee for the buyer, so that we can extract a constant fraction of her utility just through the entry fee. If we modify the $\widetilde{\mathcal{M}}$ by introducing an entry fee for every buyer according to the concentration inequality for subadditive functions, the new mechanism $\widetilde{\mathcal{M}}'$, which is a SPEM, should have revenue that is a constant fraction of $\widetilde{\mathcal{M}}$'s social welfare. The reason is simple. $\widetilde{\mathcal{M}}$'s welfare is simply its revenue plus the sum of the buyers' utilities, and $\widetilde{\mathcal{M}}'$ can extract some extra revenue from the entry fees, which is a constant fraction of the sum of the buyers' utilities. Therefore, if there exists a sequential posted price mechanism that achieves a constant fraction of M's social welfare under the truncated valuation v', the modified mechanism can obtain a constant fraction of the CORE as revenue. Surprisingly, when the buyers have XOS valuations, Feldman et al. [Feldman et al. 2015] showed that there exists an *anonymous sequential posted price mechanism* that always obtains at least $\frac{1}{2}$ of the optimal social welfare with a prophet inequality type argument. Hence, a SPEM, more specifically an ASPE as the posted prices are anonymous, should approximate the CORE well. However, this intuition cannot be directly translated into a mathematical proof. Aside from some small issues, the main technical hurdle is that the Lipschitz constant α of the buyers' utilities are too large for us to bound. In the next paragraph, we will discuss how to use a new technique – the *shifted core* – to handle the Lipschitz constant and turn the intuition into a proof.

Shifted Core. Instead of considering v', we carefully construct a new valuation $\hat{v}$ that is always dominated by the true valuation v. In particular, it is also a truncated version of v, but unlike v' which is truncated at $\boldsymbol{\beta} + \boldsymbol{c}$, $\hat{v}$ is truncated at a different threshold. We consider the social welfare of M under $\hat{v}$ and define it as

$\widehat{\text{CORE}}$. This is called the shifted core as it is similar to the CORE except that the valuations are truncated at different places. Here we make use of the adaptiveness of the dual. As we only care about the social welfare under M, and both v' and $\hat{v}$ are chosen according to M, it turns out that when the two valuations are far away from each other for some buyer i and item j, M does not allocate j to i too often. As a result, we show that the difference between $\widehat{\text{CORE}}$ and CORE can be bounded by the revenue of a RSPM. So it suffices to approximate $\widehat{\text{CORE}}$. But why is $\widehat{\text{CORE}}$ easier to approximate? The reason is two-fold: **(i)** thanks to the way that $\hat{v}$ is constructed, any buyer i's surplus under $\hat{v}_i$ is not only subadditive over independent items, but also has a small Lipschitz constant τ_i. These Lipschitz constants are so small that $\sum_{i=1}^{n} \tau_i$ and can be upper bounded by a constant number of RSPMs; **(ii)** if we construct an ASPE with item prices similar to Feldman et al. [Feldman et al. 2015] and the entry fee function $\delta_i(S)$ to be the median of buyer i's surplus for S over the randomness of t_i[10], using a proof inspired by [Feldman et al. 2015], we can show that the revenue from the posted prices plus the expected surpluses of the buyers approximates the $\widehat{\text{CORE}}$. How can we approximate the buyers' expected surpluses? Again by the concentration inequality for subadditive functions, we can bound buyer i's expected surplus by her entry fee and τ_i. As $\hat{v}$ is dominated by the true valuation v, thus for any type t_i of buyer i and any set of available items, the surplus under the true valuation v_i also dominates the surplus under $\hat{v}_i$. Remember that the entry fee is the median of the surplus under $\hat{v}_i$, so buyer i must accept the entry fee with probability at least 1/2. Hence, the revenue from the entry fees is at least a constant fraction of the sum of buyers' surpluses in the ASPE.

In all previous results, when the mechanism has an entry fee component, the analysis can only makes use of the revenue from the entry fees. We proposed a new mechanism SPEM, and by combining the concentration inequality for subadditive functions with a prophet inequality type argument, our analysis take into account both the revenue from the entry fees and the posted prices. This key improvement allows us to approximate the CORE.

REFERENCES

BABAIOFF, M., IMMORLICA, N., LUCIER, B., AND WEINBERG, S. M. 2014. A Simple and Approximately Optimal Mechanism for an Additive Buyer. In *the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.

BEI, X. AND HUANG, Z. 2011. Bayesian Incentive Compatibility via Fractional Assignments. In *the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.

BRIEST, P., CHAWLA, S., KLEINBERG, R., AND WEINBERG, S. M. 2010. Pricing Randomized Allocations. In *the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.

CAI, Y., DEVANUR, N. R., AND WEINBERG, S. M. 2016. A duality based unified approach to bayesian mechanism design. In *the 48th Annual ACM Symposium on Theory of Computing (STOC)*.

CAI, Y. AND HUANG, Z. 2013. Simple and Nearly Optimal Multi-Item Auctions. In *the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.

[10]The surplus is the buyer's value for her favorite bundle minus the corresponding prices. It is different from the utility of a buyer in an ASPE due to the existence of the entry fee.

Cai, Y. and Zhao, M. 2016. Simple mechanisms for subadditive buyers via duality. *CoRR abs/1611.06910*.

Cai, Y. and Zhao, M. 2017. Simple mechanisms for subadditive buyers via duality. In *the 49th Annual ACM Symposium on Theory of Computing (STOC)*.

Chawla, S., Hartline, J. D., and Kleinberg, R. D. 2007. Algorithmic Pricing via Virtual Valuations. In *the 8th ACM Conference on Electronic Commerce (EC)*.

Chawla, S., Hartline, J. D., Malec, D. L., and Sivan, B. 2010. Multi-Parameter Mechanism Design and Sequential Posted Pricing. In *the 42nd ACM Symposium on Theory of Computing (STOC)*.

Chawla, S., Malec, D. L., and Sivan, B. 2015. The power of randomness in bayesian optimal mechanism design. *Games and Economic Behavior 91*, 297–317.

Chawla, S. and Miller, J. B. 2016. Mechanism design for subadditive agents via an ex-ante relaxation. In Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016. *CoRR*, 579–596.

Daskalakis, C., Deckelbaum, A., and Tzamos, C. 2013. Mechanism design via optimal transport. In *ACM Conference on Electronic Commerce, EC '13, Philadelphia, PA, USA, June 16-20, 2013*. 269–286.

Daskalakis, C., Deckelbaum, A., and Tzamos, C. 2014. The complexity of optimal mechanism design. In *the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.

Daskalakis, C. and Weinberg, S. M. 2012. Symmetries and Optimal Multi-Dimensional Mechanism Design. In *the 13th ACM Conference on Electronic Commerce (EC)*.

Feldman, M., Gravin, N., and Lucier, B. 2015. Combinatorial auctions via posted prices. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*. 123–135.

Feldman, M., Svensson, O., and Zenklusen, R. 2016. Online contention resolution schemes. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*. 1014–1033.

Hart, S. and Nisan, N. 2012. Approximate Revenue Maximization with Multiple Items. In *the 13th ACM Conference on Electronic Commerce (EC)*.

Hart, S. and Nisan, N. 2013. The menu-size complexity of auctions. In *the 14th ACM Conference on Electronic Commerce (EC)*.

Hart, S. and Reny, P. J. 2012. Maximal revenue with multiple goods: Nonmonotonicity and other observations. *Discussion Paper Series dp630, The Center for the Study of Rationality, Hebrew University, Jerusalem*.

Hartline, J., Kleinberg, R., and Malekian, A. 2011. Bayesian incentive compatibility via matchings. In *the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.

Hartline, J. D. and Lucier, B. 2010. Bayesian Algorithmic Mechanism Design. In *the 42nd ACM Symposium on Theory of Computing (STOC)*.

Li, X. and Yao, A. C.-C. 2013. On revenue maximization for selling multiple independently distributed items. *Proceedings of the National Academy of Sciences 110,* 28, 11232–11237.

Myerson, R. B. 1981. Optimal Auction Design. *Mathematics of Operations Research 6,* 1, 58–73.

Pavlov, G. 2011. Optimal mechanism for selling two goods. *The B.E. Journal of Theoretical Economics 11,* 3.

Rochet, J.-C. and Chone, P. 1998. Ironing, sweeping, and multidimensional screening. *Econometrica 66,* 4 (July), 783–826.

Rubinstein, A. and Weinberg, S. M. 2015. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*. 377–394.

Schechtman, G. 2003. Concentration, results and applications. *Handbook of the geometry of Banach spaces 2*, 1603–1634.

Talagrand, M. 1995. Concentration of measure and isoperimetric inequalities in product spaces. *Publications Mathématiques de l'Institut des Hautes Etudes Scientifiques 81,* 1, 73–205.

THANASSOULIS, J. 2004. Haggling over substitutes. *Journal of Economic Theory 117,* 2 (August), 217–245.

YAO, A. C. 2015. An n-to-1 bidder reduction for multi-item auctions and its applications. In SODA. *CoRR*.