A Prior-Independent Revenue-Maximizing Auction for Multiple Additive Bidders*

Kira Goldner and Anna R. Karlin

Department of Computer Science and Engineering, University of Washington,
Box 352350, Seattle, WA 98195, USA
{kgoldner,karlin}@cs.washington.edu

Abstract. Recent work by Babaioff et al. [1], Yao [30], and Cai et al. [7] shows how to construct an approximately optimal auction for additive bidders, given access to the priors from which the bidders' values are drawn. In this paper, building on the single sample approach of Dhangwatnotai et al. [15], we show how the auctioneer can obtain approximately optimal expected revenue in this setting *without* knowing the priors, as long as the item distributions are regular.

Keywords: Mechanism design; approximately optimal auctions; prior-independence; additive bidders

1 Introduction

In a *multiple additive bidders setting*, there are n agents and a seller selling a set of m distinct items. Each agent i has a private value v_{ij} for item j, and value $v_i(S) = \sum_{j \in S} v_{ij}$ for the set of items S. The seller runs an auction to determine who (if anyone) to sell each item to and at what price. The auction (or mechanism) takes as input the collection of bids, and determines a feasible allocation and a price to charge each agent. The seller knows ahead of time the distribution from which each v_{ij} is drawn.[1] A key question is how to design a truthful and optimal[2] (or approximately optimal) auction.

This is a notoriously difficult problem, but in the past decade, several breakthrough results have been obtained. There are three main lines of work related to optimal auctions for additive bidders. For the case of finite type spaces, [3–6, 29] are able to use linear and convex programming techniques to formulate and solve the optimal auction problem. This gives a black-box reduction from mechanism design to algorithm design that yields a polynomial time algorithm for revenue maximization in additive settings. A second strand of work [20, 22, 1, 30, 7] handles arbitrary distributions and develops approximately optimal auctions.

* This research was done in part while the authors were visiting the Simons Institute for Theoretical Computer Science. The authors are funded by the National Science Foundation under CCF grant 1420381.

[1] Thus, from the seller's perspective this value is a random variable V_{ij}.

[2] i.e., revenue-maximizing, in expectation.

Finally, [11, 12, 10, 18, 19, 16] use duality frameworks to optimally solve the problem for certain settings with a small number of items, and to provide necessary and sufficient conditions under which grand bundle selling is optimal.

In this paper, we consider the question of *prior-independent* optimal mechanism design in the multiple additive bidders setting. By prior-independent we mean two things: first, that there exist prior distributions from which the agents' values are drawn (as in all the work discussed above), and, second, that the mechanism designer has no knowledge of these priors. Thus, without any knowledge of the priors, we seek to construct a mechanism that guarantees a constant fraction of the expected profit achieved by the optimal mechanism tailored to the particular prior distributions. This guarantee should hold no matter what the distributions happen to be, as long as they satisfy the fairly standard condition of regularity. A growing body of work obtains prior-independent mechanisms in a number of settings [15, 13, 26, 25, 17].

The main result of this paper is an auction that achieves this goal for the additive bidder setting when the V_{ij}'s are all independent and drawn from regular distributions. We give a mechanism that requires only a single sample from the distribution of each V_{ij}, and when there are at least two bidders from any prior distribution, we can implement a sample mechanism as a prior-independent mechanism. Thus, we add to the short list of prior-independent results in multi-parameter settings [13, 26].

Our work builds on the breakthrough results of Babaioff, Imorlica, Lucier, and Weinberg [1] and Yao [30], on the one hand, and Dhangwatnotai, Roughgarden, and Yan [15] on prior-independent mechanism design on the other hand. A crucial lemma in [15][3] is that, for a single-item single-bidder problem, access to a *single sample* from a regular distribution is sufficient to approximate the optimal revenue, which in this case is the revenue that results from pricing at the reserve price for the distribution.

Amazingly, the approximately optimal auctions of Babaioff et al., Yao, and Cai et al. essentially only use second-price auctions and reserve pricing for either single items or bundles of items, and therefore, the single sample paradigm nearly suffices to construct a prior-independent version of these auctions. There is only one detail to resolve and that relates to the issue of pricing bundles: the sum of regular random variables is not necessarily regular. However, delving into the proof from Babaioff et al., we find that the solution to this problem essentially writes itself: in the "bad" case, when bundle pricing is necessary for approximating the optimal revenue, it happens to be that the relevant random variable concentrates so that, in fact, a sample bundle price is sufficient.

1.1 Other Related Work

An important line of recent research [9, 21, 24, 14, 23] has explored the sample complexity of auctions. For example, [15] shows that with a single sample, one can design an auction that gets a constant factor approximation to the optimal

[3] This is a reinterpretation of the Bulow-Klemperer Theorem [2].

single item auction. How much better can you do with more samples? This question has been explored in a number of auction settings; e.g. by Morgenstern and Roughgarden [23] in the additive bidder setting we study in this paper.

1.2 Organization

After preliminaries, we show in Section 2 how to approximate the optimal single additive bidder revenue when given access to a sample from each item distribution. Then, in Section 3, we use the latter result to give two approximately optimal auctions for the multiple additive bidders setting: one that is given access to a sample from every item distribution from each bidder and one is that is fully prior-independent. In Section 4 we discuss an improved analysis for bidders with finite support distributions. We conclude with open problems in Section 5.

1.3 Preliminaries

In this paper, we consider the setting of a revenue-maximizing monopolist seller with m items to sell to n additive bidders. Each bidder i has his value V_{ij} for item j drawn from an unknown prior distribution F_{ij}. All bidders are additive: for any set of items S, bidder i's value for the set is

$$V_i(S) = \sum_{j \in S} V_{ij}.$$

We will assume that each of the distributions F_{ij} is regular. That is, $\varphi_{ij}(v) = v - \frac{1-F_{ij}(v)}{f_{ij}(v)}$ is non-decreasing.

We also use the following notation:

- The *revenue curve* $R(\cdot)$ gives the expected revenue for selling an item at a price x to a bidder with value V drawn from distribution F. That is, $R(x) := x \cdot Pr[V \geq x]$.
- The *monopoly price* r^* is the price that maximizes revenue: $r^* := \text{argmax}_x R(x)$.
- Consider a single additive bidder with value $V_j \sim G_j$ for item j. Then $\text{SRev}(V_1, \dots, V_m)$ denotes the optimal expected revenue that can be obtained by posting a price for each item j individually. That is,

$$\text{SRev}(V_1, \dots, V_m) := \sum_j \max_{x_j} R_j(x_j),$$

  where R_j is the revenue curve associated with the distribution G_j.
- $\text{BRev}(V_1, \dots, V_m)$ denotes the optimal expected revenue for posting a price on the "grand bundle" of all of the items to this same additive bidder with $V_j \sim G_j$. That is,

$$\text{BRev}(V_1, \dots, V_m) := \max_x x \cdot \Pr[\sum_j V_j \geq x].$$

- For any number x, let $(x)^+$ denote $\max\{x, 0\}$.

2 A prior independent mechanism for a single additive bidder, given samples

Our mechanism draws heavily on two prior results. The first demonstrates that access to a single sample from a bidder's distribution can be used to obtain a $\frac{1}{2}$-approximation of optimal revenue in the single-item setting.

Theorem 1 (Dhangwatnotai, Roughgarden, and Yan 2010). *Consider a bidder whose value for a particular item is drawn from F, a regular distribution with monopoly price r^* and revenue function $R(\cdot)$. Let $S \sim F$ be a random sample from the distribution F. Then, for every nonnegative number t,*

$$\mathbb{E}(R(\max\{t, S\})) \geq \frac{1}{2} R(\max\{t, r^*\}).$$

Therefore, in particular, for $t = 0$, the expected revenue from posting a price of S yields at least half of the optimal posted price revenue, which is $R(r^*)$.

The second result we use demonstrates that a combination of two very simple mechanisms can be used to obtain a constant factor of the optimal revenue in the single additive bidder setting.

Theorem 2 (Babaioff, Immorlica, Lucier, and Weinberg 2014). *Consider a single additive bidder with value V_j for item j drawn independently from distribution G_j. Denote by $\textsc{opt}(V_1, \ldots, V_m)$ the revenue of the optimal mechanism. Let $t = \mathrm{SRev}(V_1, \ldots, V_m)$ denote the optimal expected revenue from selling the items separately, and define $V := \sum_{j=1}^{m} V_j$, the bidder's value for the grand bundle. Then*

- *If $\mathbb{E}[V | \; V_j \leq t \quad \forall j] \leq 4\mathrm{SRev}(V_1, \ldots, V_m)$, then*

$$\mathbb{E}[\textsc{opt}(V_1, \ldots, V_m)] \leq 6\, \mathrm{SRev}(V_1, \ldots, V_m).$$

- *Otherwise, if $\mathbb{E}[V | \; V_j \leq t \quad \forall j] > 4\mathrm{SRev}(V_1, \ldots, V_m)$, then*

$$\Pr\left[V \geq \frac{2}{5} \cdot \mathbb{E}[V | \; V_j \leq t \quad \forall j]\right] \geq \frac{47}{72}$$

and

$$\mathbb{E}[\textsc{opt}(V_1, \ldots, V_m)] \leq 2\, \mathrm{SRev}(V_1, \ldots, V_m) + \mathbb{E}[V \mid V_j \leq t \;\; \forall j].$$

From this, Babaioff et al. obtain the following corollary:

Corollary 1 (Babaioff, Immorlica, Lucier, and Weinberg 2014). *Consider a single additive bidder with value V_j for item j drawn independently from distribution G_j. Let $\mathrm{SRev}(V_1, \ldots, V_m)$ denote the optimal expected revenue from selling the items separately and let $\mathrm{BRev}(V_1, \ldots, V_m)$ denote the optimal expected revenue from selling the grand bundle. Then*

$$\mathbb{E}[\textsc{opt}(V_1, \ldots, V_m)] \leq 6 \cdot \max\{\mathrm{SRev}(V_1, \ldots, V_m) + \mathrm{BRev}(V_1, \ldots, V_m)\}.$$

We now combine the single additive bidder analysis with samples from the distributions to give an approximately optimal mechanism for a single additive bidder that does not rely on knowledge of the priors, but rather uses a single sample from each distribution.

The multi-bidder analogue of the better of selling separately or selling the grand bundle is a two-part tariff mechanism as used in [30, 7, 8]. Here, each bidder is offered a list of item prices and an entry fee. Typically in these mechanisms, some item prices are determined first. Then, the buyer's surplus values above each item's price are analyzed to understand either (a) whether to increase the item prices or to use an entry-fee or (b) how to compute the entry fee. This is equivalent to the buyer's prior distribution for each item shifted down by the item's price. We define the mechanism in a slightly more general way than is necessary here with a parameter $\boldsymbol{\Delta}$ in order to easily extend to the case where we want to analyze the shifted distributions. In the single bidder setting, we do not shift the distributions, so we will set the shift $\Delta_j = 0$ for all j. However, this parameter will allow us to use this mechanism as a black box in the multiple bidders setting.

Definition 1. *Define the* Sample Mechanism *as follows. Given a set* $(S_1, \ldots, S_m)$ *of samples from an additive bidder's distribution, and a set of nonnegative values* $\Delta_1, \ldots, \Delta_m$*,*

(a) with probability $\frac{1}{2}$*: Offer a price of* $\max\{\Delta_j, S_j\}$ *for each item* j *separately.*
(b) with probability $\frac{1}{2}$*: Offer the bidder a price of* $S^+ = \sum_{j=1}^m (S_j - \Delta_j)^+$ *to enter the auction. If he pays the entrance fee, he can take any item* j *he wants at price* Δ_j*. (When* $\boldsymbol{\Delta} = \mathbf{0}$*, this is simply pricing the grand bundle.)*

Denote the revenue from this mechanism as $\textsc{Samp}(V_1, \ldots, V_m; \boldsymbol{\Delta})$*.*

Theorem 3. *Consider a single additive bidder with value* V_j *for item* j *drawn independently from regular distribution* G_j*. Let* $\Delta_1, \ldots, \Delta_j \geq 0$*, and define* $V_j^+ = (V_j - \Delta_j)^+$*. The Sample Mechanism has expected revenue which is a constant fraction of the optimal expected revenue for* $(V_1^+, \ldots, V_m^+)$*.*

Proof. The first step of the Sample Mechanism obtains expected revenue which is a constant fraction of $\textsc{SRev}(V_1^+, \ldots, V_m^+)$. To see this, note that

$$\textsc{SRev}(V_1^+, \ldots, V_m^+) = \sum_j \max_{x \geq \Delta_j} \left[(x - \Delta_j)^+ (1 - G_j(x))\right]$$

$$\leq \sum_j \max_{x \geq \Delta_j} R_j(x) \leq \sum_j R_j(\max\{\Delta_j, r_j^*\}),$$

where $R_j(\cdot)$ is the revenue curve for G_j and r_j^* is the optimal reserve price for G_j. The last inequality follows from regularity of G_j, and thus concavity of $R_j(\cdot)$. An application of Lemma 1 shows that offering the bidder a price of $\max\{\Delta_j, S_j\}$ for item j yields expected revenue at least half of $R_j(\max\{\Delta_j, r_j^*\})$, hence step (a) is a $\frac{1}{4}$-approximation to $\mathrm{SRev}(V_1^+, \dots, V_m^+)$.

To complete the proof, we show how the Sample Mechanism approximates $\mathbb{E}[\textsc{opt}(V_1^+, \dots, V_m^+)]$ within a constant factor using the two cases of Theorem 2. To this end, let $t = \mathrm{SRev}(V_1^+, \dots, V_m^+)$ and define $C := \mathbb{E}[V^+ |\ V_j^+ \le t \quad \forall j]$, where $V^+ = \sum_j V_j^+$. If $C \le 4t$, we are in the first case of Theorem 2 applied to the random variables $(V_1^+, \dots, V_m^+)$, hence

$$\mathbb{E}[\textsc{opt}(V_1^+, \dots, V_m^+)] \le 6\,\mathrm{SRev}(V_1^+, \dots, V_m^+) \le 24\,\textsc{Samp}(V_1, \dots, V_m; \boldsymbol{\Delta}).$$

Otherwise, $C > 4t$, and we also need to consider the revenue from the second step of the Sample Mechanism. In this case, from Theorem 2, we have

$$\mathbb{E}[\textsc{opt}(V_1^+, \dots, V_m^+)] \le 2\,\mathrm{SRev}(V_1^+, \dots, V_m^+) + C \quad \text{and} \quad \Pr\left[V^+ \ge \frac{2}{5}C\right] \ge \frac{47}{72}.$$

Next, recall that $S^+ = \sum_{j=1}^m (S_j - \Delta_j)^+$ is entry fee that is offered. Observe that if $V^+ > S^+$, then the bidder will enter the auction, since his utility will then be $\sum_{j|V_j \ge \Delta_j}(V_j - \Delta_j) - S^+ = V^+ - S^+$. Therefore, the expected revenue from step (b), the entry fee portion of the auction, can be bounded as follows:

$$\begin{aligned}
&\mathbb{E}\left[S^+ \mid V^+ \ge S^+\right]\Pr[V^+ \ge S^+] \\
&\quad \ge \mathbb{E}\left[S^+ \,\middle|\, V^+ \ge S^+,\ V^+, S^+ \ge \frac{2}{5}C\right] \cdot \Pr\left[V^+, S^+ \ge \frac{2}{5}C,\ V^+ \ge S^+\right] \\
&\quad \ge \frac{2}{5}C \cdot \Pr\left[V^+ \ge S^+ \,\middle|\, V^+, S^+ \ge \frac{2}{5}C\right] \cdot \left(\frac{47}{72}\right)^2 \\
&\quad > \frac{1}{2} \cdot \frac{1}{6}C \\
&\quad = \frac{1}{12}C.
\end{aligned} \tag{1}$$

The third line follows from the second part of Theorem 2, and the independence of V^+ and S^+. The fourth line follows from the fact that V^+ and S^+ are identically distributed.

In this case, it is now clear that the Sample Mechanism obtains a constant factor approximation:

$$\mathbb{E}[\textsc{opt}(V_1^+, \dots, V_m^+)] \le 2\,\mathrm{SRev}(V_1^+, \dots, V_m^+) + C \le 24\,\textsc{Samp}(V_1, \dots, V_m; \boldsymbol{\Delta}).$$

Thus in either case, $\mathbb{E}[\textsc{opt}(V_1^+, \dots, V_m^+)] \le 24\,\textsc{Samp}(V_1, \dots, V_m; \boldsymbol{\Delta})$. This mechanism loses a factor of 4 compared to the prior-dependent $\max\{\mathrm{SRev}, \mathrm{BRev}\}$ mechanism in [1].

Corollary 2. *Consider a single additive bidder with value V_j for item j drawn independently from regular distribution G_j. The Sample Mechanism with $\Delta_j = 0$ for all j has expected revenue which is a constant fraction of the optimal expected revenue $\mathbb{E}[\textsc{opt}(V_1, \dots, V_m)]$.*

3 Multiple additive bidders

Our mechanism builds on the following breakthrough result from which Yao constructs a simple, approximately optimal mechanism for the multiple additive bidders setting.

Theorem 4 (Yao 2015). *Consider n additive bidders, where V_{ij} is the value bidder i has for item j, where $1 \leq i \leq n$ and $1 \leq j \leq m$. Assume the set of random variables $\{V_{ij}\}_{j=1}^{m}$ are independent for each i. Define the following auxiliary random variables:*

$$X_{ij} := \max_{k \neq i} V_{kj} \quad \text{and} \quad A_{ij} := (V_{ij} - X_{ij})^{+}.$$

Then for $\boldsymbol{V}_j = (V_{1j}, \ldots, V_{nj})$, the expected revenue of the optimal mechanism for the multiple additive bidders setting satisfies

$$\mathbb{E}[\textsc{opt}(\boldsymbol{V}_1, \ldots, \boldsymbol{V}_m)] \leq 8 \sum_i \mathbb{E}[\textsc{opt}(A_{i1}, \ldots, A_{im})] + 9\, \mathbb{E}[\mathrm{SPA}(\boldsymbol{V}_1, \ldots, \boldsymbol{V}_m)],$$

where $\mathrm{SPA}(\cdot)$ *is the revenue from running a separate second-price auction for each item and* $\textsc{opt}(A_{i1}, \ldots, A_{im})$ *denotes the revenue obtained by the optimal single additive bidder auction, when that bidder's value for item j is A_{ij}.*

3.1 A Sample Auction for Multiple Bidders

By randomly choosing to either (a) run a second-price auction separately on each item or (b) run sample mechanisms on each bidder with $\Delta_{ij} = X_{ij}$, we can achieve a constant-fraction of the optimal revenue from only samples in the multiple-additive-bidder setting as well.

Definition 2. *Define the* Multiple-Additive-Bidders Sample Mechanism *parameterized by p as follows, given a sample S_{ij} from each bidder i's distribution for item j:*

(a) with probability p: Run a Second-Price Auction on each item j. That is, offer each bidder i the option to take item j at a price equal to X_{ij}.
(b) with probability $1-p$: Offer each bidder i an entry fee of $\sum_j (S_{ij} - X_{ij})^{+}$. Any bidder willing to pay the entry fee can then take[4] item j at price X_{ij}.

Let $\text{MAB-}\textsc{Samp}(\boldsymbol{V}_1, \ldots, \boldsymbol{V}_m; p)$ *denote the revenue from the Multiple-Additive-Bidders Sample Mechanism with parameter p.*

Theorem 5. *In the setting of Theorem 4, when the random variables V_{ij} are all independent and, for each j, the random variables V_{ij} is drawn from regular distribution F_{ij} for each bidder i and item j, with access to a sample S_{ij} from each F_{ij}, the Multiple-Additive-Bidders Sample Mechanism with parameter $p = \frac{105}{201}$ obtains at least a constant fraction of the optimal expected revenue.*

[4] This guarantees that each item is taken by at most one bidder.

Proof. Recall that $X_{ij} := \max_{k\neq i} V_{kj}$ is the highest bid for item j excluding bidder i's bid, and $A_{ij} := (V_{ij} - X_{ij})^+$ is the surplus from buyer i's value for item j over this price.

Given sample S_{ij} for each item j, an application of Theorem 3 on each bidder i where $\Delta_{ij} = X_{ij}$ gives that $\textsc{opt}(A_{i1}, \ldots, A_{im}) \leq 24\,\textsc{Samp}(V_{i1}, \ldots, V_{im}; \boldsymbol{X}_i)$, where $\boldsymbol{X}_i = (X_{i1}, \ldots, X_{im})$. Using Theorem 4, this gives that

$$\begin{aligned}\mathbb{E}[\textsc{opt}(\boldsymbol{V}_1, \ldots, \boldsymbol{V}_m)] &\leq 8\sum_i \mathbb{E}[\textsc{opt}(A_{i1}, \ldots, A_{im})] + 9\,\mathbb{E}[\mathrm{SPA}(\boldsymbol{V}_1, \ldots, \boldsymbol{V}_m)] \\ &\leq 8 \cdot 24 \sum_i \mathbb{E}_{\boldsymbol{X}_i}[\textsc{Samp}(V_{i1}, \ldots, V_{im}; \boldsymbol{X}_i)] \\ &\quad + 9\,\mathbb{E}[\mathrm{SPA}(\boldsymbol{V}_1, \ldots, \boldsymbol{V}_m)] \\ &\leq 201\,\text{MAB-}\textsc{Samp}(\boldsymbol{V}_1, \ldots, \boldsymbol{V}_m; \frac{105}{201})\end{aligned}$$

Running the sample mechanisms with probability $\frac{192}{201}$ and a second-price auction separately on each item with probability $\frac{9}{201}$ gives a 201-approximation, which is equivalent to running a second-price auction with probability $\frac{105}{201}$ and the entry fee mechanism with the remaining probability.. Similarly to the single bidder case, this loses less than a factor of 4 compared to the Bundling Mechanism in [30] which requires full knowledge of all of the prior distributions and achieves a 57-approximation.

3.2 A Prior-Independent Auction

We can also use sample mechanisms to sell to multiple additive bidders without extra samples. Analogously to [15], if the seller can identify which bidders come from the same distribution, she can take a sample bidder from each group a of identically distributed bidders and use it to set the prices for the rest of the group. This requires at least two bidders from each distribution group a. The mechanism is the same as the Multiple-Additive-Bidders Sample Mechanism, but with randomly excluded bidders used as samples.

Definition 3. *Define the* Multiple-Additive-Bidders Prior-Independent Mechanism *parameterized by p as follows:*

(a) with probability p: Run a Second-Price Auction on each item j. That is, offer each bidder i the option to take item j at a price equal to X_{ij}.
(b) with probability $1-p$: Remove a random bidder i_a from each group of bidders a and let S_{aj} be his bid for item j (i.e., $S_{aj} := V_{i_a j}$). Let $\mathcal{S}$ be the set of bidders sampled from each group a. Also, let $\beta_{ij} = \max_{k\notin\mathcal{S},\ k\neq i} V_{kj}$. Offer each remaining bidder i from group a an entry fee of $\sum_j (S_{aj} - \beta_{ij})^+$. Any bidder willing to pay the entry fee can then take item j at price β_{ij}.

Let MAB-PI$(\boldsymbol{V}_1, \ldots, \boldsymbol{V}_m; p)$ *denote the revenue from the Multiple-Additive-Bidders Prior-Independent Mechanism with parameter p.*

Theorem 6. *In the setting of Theorem 4, when the random variables V_{ij} are all independent and, for each j, the random variables V_{ij} is drawn from regular distribution F_{ij} for each bidder i and item j, with at least 2 bidders from every distribution group a, the Multiple-Additive-Bidders Prior-Independent Mechanism with parameter $p = \frac{585}{1161}$ obtains at least a constant fraction of the optimal expected revenue.*

Proof. If n_a is the number of bidders from distribution group a and $n_a^{\min}$ is the number of bidders in the smallest such group, then

$$\begin{aligned}\sum_i \mathbb{E}[\textsc{opt}(A_{i1}, \ldots, A_{im})] &\leq \sum_a \frac{n_a}{n_a - 1} \sum_{i \in a, i \notin \mathcal{S}} \mathbb{E}[\textsc{opt}(A_{i1}, \ldots, A_{im})] \\ &\leq \frac{n_a^{\min}}{n_a^{\min} - 1} \sum_{i \notin \mathcal{S}} \mathbb{E}[\textsc{opt}(A_{i1}, \ldots, A_{im})].\end{aligned}$$

Also notice that

$$A_{ij} := \Big(V_{ij} - \max\{\max_a\{S_{aj}\}, \beta_{ij}\}\Big)^+ \qquad \text{and define} \qquad V_{ij}^+ := (V_{ij} - \beta_{ij})^+.$$

Clearly, the random variable V_{ij}^+ dominates the random variable A_{ij} (i.e., $\Pr(V_{ij}^+ \geq x) \geq \Pr(A_{ij} \geq x)$ for all x). Therefore,

$$\textsc{SRev}(A_{i1}, \ldots, A_{im}) \leq \textsc{SRev}(V_{i1}^+, \ldots, V_{im}^+)$$

and

$$\textsc{BRev}(A_{i1}, \ldots, A_{im}) \leq \textsc{BRev}(V_{i1}^+, \ldots, V_{im}^+)$$

Thus, by Corollary 1, it suffices to obtain a constant fraction of $\mathbb{E}[\textsc{opt}(V_{i1}^+, \ldots, V_{im}^+)]$ for each i.

Using the analysis from Theorem 5, we put it all together to see that

$$\begin{aligned}\mathbb{E}[\textsc{opt}(\boldsymbol{V}_1, \ldots, \boldsymbol{V}_m)] &\leq 8 \cdot \frac{n_a^{\min}}{n_a^{\min} - 1} \sum_{i \notin \mathcal{S}} \mathbb{E}[\textsc{opt}(A_{i1}, \ldots, A_{im})] \\ &\quad + 9\,\mathbb{E}[\mathrm{SPA}(\boldsymbol{V}_1, \ldots, \boldsymbol{V}_m)] \\ &\leq 8 \cdot \frac{n_a^{\min}}{n_a^{\min} - 1} \cdot 6 \sum_{i \notin \mathcal{S}} \mathbb{E}[\textsc{opt}(V_{i1}^+, \ldots, V_{im}^+)] \\ &\quad + 9\,\mathbb{E}[\mathrm{SPA}(\boldsymbol{V}_1, \ldots, \boldsymbol{V}_m)] \\ &\leq 8 \cdot \frac{n_a^{\min}}{n_a^{\min} - 1} \cdot 6 \cdot 24 \sum_{i \notin \mathcal{S}} \mathbb{E}_{\boldsymbol{\beta}_i}[\textsc{Samp}(V_{i1}, \ldots, V_{im}; \boldsymbol{\beta}_i)] \\ &\quad + 9\,\mathbb{E}[\mathrm{SPA}(\boldsymbol{V}_1, \ldots, \boldsymbol{V}_m)] \\ &\leq 1161 \cdot \frac{n_a^{\min}}{n_a^{\min} - 1} \,\text{MAB-PI}(\boldsymbol{V}_1, \ldots, \boldsymbol{V}_m; \frac{585}{1161})\end{aligned}$$

where of course, since $n_a^{\min} \geq 2$, then $\frac{n_a^{\min}}{n_a^{\min}-1} \leq 2$.

Note that the loss due to excluding bidders to use as samples is a factor of $6\frac{n_a^{\min}}{n_a^{\min}-1}$.

4 Finite Support Distributions

Cai, Devanur, and Weinberg [7] present a new framework that analyzes revenue from multiple additive bidders with finite support distributions (over discrete value spaces) via a similar core-tail decomposition. These results also hold for discretizing a continuous value space and losing at most a factor of $1+\varepsilon$ in the revenue due to the discretization. Utilizing this analysis improves the constant of our approximation.

Precisely, they show that

$$E[\textsc{opt}(\boldsymbol{V}_1, \ldots, \boldsymbol{V}_m)] \leq 4\,\textsc{SRev}(\boldsymbol{V}_1, \ldots, \boldsymbol{V}_m) + \textsc{Core}$$

where, for the highest other bid $X_{ij} := \max_{k \neq i} V_{kj}$,

$$\textsc{Core} = \mathbb{E}[\sum_{i=1}^{n}\sum_{j=1}^{m}(V_{ij} - X_{ij})^+ \mathbb{1}_{V_{ij} \in [X_{ij}, X_{ij}+t_i]}]$$

and

$$t_i = \textsc{SRev}((V_{i1} - X_{i1})^+, \ldots, (V_{im} - X_{im})^+).$$

Then if $A_{ij} = (V_{ij} - X_{ij})^+ \cdot \mathbb{1}_{V_{ij} \in [X_{ij}, X_{ij}+t_i]}]$, we have that

$$\textsc{Core} = \sum_{i=1}^{n} \mathbb{E}_{\boldsymbol{V}_{-i}}[\mathbb{E}_{V_i}[\sum_{j=1}^{m} A_{ij}]].$$

Then in a proof nearly identical to that of Theorem 3, we can show that $\textsc{Core} \leq 24\sum_i \mathbb{E}_{\boldsymbol{V}_{-i}}[\textsc{Samp}(V_{i1}, \ldots, V_{im}; \boldsymbol{X}_i)]$. For each bidder, we set $\Delta_j = X_{ij}$ and bound $\mathbb{E}[\sum_{j=1}^{m} A_{ij}] \leq 24\textsc{Samp}(V_{i1}, \ldots, V_{im}; \boldsymbol{\Delta})$.

In one case, $\mathbb{E}[\sum_{j=1}^{m} A_{ij}] \leq 4t = 4\textsc{SRev}((V_{i1} - X_{i1})^+, \ldots, (V_{im} - X_{im})^+) \leq 16\,\textsc{Samp}(V_{i1}, \ldots, V_{im}; \boldsymbol{X}_i)$. In the other case, $\mathbb{E}[\sum_{j=1}^{m} A_{ij}] > 4t$. In this case, similarly to the proof of Theorem 2, we get that

$$\begin{aligned}
\Pr[|\sum_j A_{ij} - \mathbb{E}[\sum_{j=1}^{m} A_{ij}]| \geq \frac{3}{5}\mathbb{E}[\sum_{j=1}^{m} A_{ij}]] \\
&\leq \frac{var(\sum_{j=1}^{m} A_{ij})}{\frac{3}{5}^2 \mathbb{E}[\sum_{j=1}^{m} A_{ij}]^2} && \text{by Chebyshev's inequality} \\
&< \frac{var(\sum_{j=1}^{m} A_{ij})}{\frac{9}{25}^2 \cdot 16t^2} && \text{since } \mathbb{E}[\sum_j A_{ij}] > 4t \\
&\leq \frac{2t^2}{\frac{9}{25}^2 \cdot 16t^2} = \frac{25}{72}.
\end{aligned}$$

The final inequality follows from the fact that $var(\sum_{j=1}^m A_{ij}) \leq t^2$ by Lemma 9 of [7].

Identically to the proof of Theorem 3, we get that if we offer the bidder a price of $S^+ = \sum_{j=1}^m (S_{ij} - X_{ij})^+$ that

$$\mathbb{E}[\sum_{j=1}^m A_{ij}] \leq 12\,\mathbb{E}[S^+ \cdot \Pr[\sum_{j=1}^m A_{ij} \geq S^+]] \leq 24\textsc{Samp}(V_{i1}, \ldots, V_{im}; \boldsymbol{X}_i).$$

Hence $\textsc{Core} \leq 24 \sum_i \mathbb{E}_{\boldsymbol{V}_{-i}}[\textsc{Samp}(V_{i1}, \ldots, V_{im}; \boldsymbol{X}_i)]$. Moreover, since $(V_{ij} - \beta_{ij})^+$ stochastically dominates $(V_{ij} - X_{ij})^+$, then we only lose a factor of $\frac{n_a^{\min}}{n_a^{\min}-1}$ for excluding bidders to use as samples to make a prior-independent auction.

As in the proof of Theorem 3, an application of Theorem 1 gives that a second-price auction on each item is a constant factor of the revenue from selling each item separately. Then again, since $\beta_{ij} = \max_{k \notin \mathcal{S},\ k \neq i} V_{kj}$,

$$\begin{aligned}
E[\textsc{opt}(\boldsymbol{V}_1, \ldots, \boldsymbol{V}_m)] &\leq 4\,\mathrm{S}\textsc{Rev}(\boldsymbol{V}_1, \ldots, \boldsymbol{V}_m) + \textsc{Core} \\
&\leq 8\,\mathrm{SPA}(\boldsymbol{V}_1, \ldots, \boldsymbol{V}_m) + \\
&\qquad 24 \cdot \frac{n_a^{\min}}{n_a^{\min} - 1} \sum_{i \notin \mathcal{S}} \mathbb{E}_{\boldsymbol{\beta}_i}[\textsc{Samp}(V_{i1}, \ldots, V_{im}; \boldsymbol{\beta}_i)] \\
&\leq 32 \cdot \frac{n_a^{\min}}{n_a^{\min} - 1} \text{MAB-PI}(\boldsymbol{V}_1, \ldots, \boldsymbol{V}_m; \frac{5}{8})
\end{aligned}$$

We lose a factor of 4 compared to the mechanism of [7] when given samples, and a factor of $4 \cdot \frac{n_a^{\min}}{n_a^{\min}-1}$ without samples.

5 Open Problems

Beyond Additive Bidders One interesting problem for future work is to design prior-independent mechanisms for more general valuations. Recent work in revenue maximization for more general multi-item settings gives mechanisms that have constant-factor approximation guarantees for a single subadditive buyer [28] and for multiple matroid-constrained buyers [8]. Both of these results rely on an analysis that chooses prices in the bidders' distributions that would sell with a constrained ex-ante probability. As these probabilities are aimed at segmenting off the tails of the distributions and samples are unlikely to come from the tail, it is unclear how to design a prior independent mechanism for these settings.

Lower Bounds Another interesting open problem is to obtain a lower bound on the gap in revenue between the optimal mechanism and the Sample Mechanism, and for the Multiple-Additive-Bidders Sample Mechanism as well. However, stronger lower bounds are still open problems for the mechanisms from [1, 30, 7] as well.

A lower bound from [15] shows that the factor of 2 in Theorem 1 is tight when a bidder's distribution for a single item is the distribution where the revenue curve is a triangle, that is, where $F(v) = \frac{v}{v+1}$ on $[0, H)$ as $H \to \infty$.

The best known lower bound on the approximation of the $\max\{\text{SRev}, \text{BRev}\}$ is a factor of 2. The example, given by Rubinstein [27], has n items from the equal revenue distribution and n rare but expensive items. The optimal revenue gets an equal fraction of revenue from each group; however, selling the grand bundle does well for the first set and poorly for the second while selling separately captures the revenue of the second set but not the first. Of course, this gap gives the Sample Mechanism a lower bound of 2 as well.

References

1. Babaioff, M., Immorlica, N., Lucier, B., Weinberg, S.M.: A simple and approximately optimal mechanism for an additive buyer. In: Proceedings of the 55th Annual Symposium on Foundations of Computer Science. FOCS '14 (2014), `http://arxiv.org/abs/1405.6146`
2. Bulow, J., Klemperer, P.: Auctions vs. negotiations. Tech. rep., National Bureau of Economic Research (1994)
3. Cai, Y., Daskalakis, C., Weinberg, S.M.: An algorithmic characterization of multidimensional mechanisms. In: Proceedings of the forty-fourth annual ACM symposium on Theory of computing. pp. 459–478. ACM (2012)
4. Cai, Y., Daskalakis, C., Weinberg, S.M.: Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In: Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on. pp. 130–139. IEEE (2012)
5. Cai, Y., Daskalakis, C., Weinberg, S.M.: Reducing revenue to welfare maximization: Approximation algorithms and other generalizations. In: Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms. pp. 578–595. SIAM (2013)
6. Cai, Y., Daskalakis, C., Weinberg, S.M.: Understanding incentives: Mechanism design becomes algorithm design. In: Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on. pp. 618–627. IEEE (2013)
7. Cai, Y., Devanur, N.R., Weinberg, S.M.: A duality based unified approach to bayesian mechanism design. In: Proceedings of the forty-eighth annual ACM symposium on Theory of computing. ACM (2016)
8. Chawla, S., Miller, J.B.: Mechanism Design for Subadditive Agents via an Ex Ante Relaxation. In: Proceedings of the 2016 ACM Conference on Economics and Computation. pp. 579–596. EC '16, ACM, New York, NY, USA (2016), `http://doi.acm.org/10.1145/2940716.2940756`
9. Cole, R., Roughgarden, T.: The sample complexity of revenue maximization. In: Proceedings of the 46th Annual ACM Symposium on Theory of Computing. pp. 243–252. ACM (2014)
10. Daskalakis, C.: Multi-item auctions defying intuition? ACM SIGecom Exchanges 14(1), 41–75 (2015)
11. Daskalakis, C., Deckelbaum, A., Tzamos, C.: Mechanism design via optimal transport. In: Proceedings of the 14th ACM Conference on Electronic Commerce. pp. 269–286. EC '13 (2013), `http://doi.acm.org/10.1145/2482540.2482593`

12. Daskalakis, C., Deckelbaum, A., Tzamos, C.: Strong duality for a multiple-good monopolist. In: Proceedings of the Sixteenth ACM Conference on Economics and Computation. pp. 449–450. ACM (2015)
13. Devanur, N., Hartline, J., Karlin, A., Nguyen, T.: Prior-independent multi-parameter mechanism design. In: Internet and Network Economics, pp. 122–133. Springer (2011), http://users.eecs.northwestern.edu/~hartline/papers/auctions-WINE-11.pdf
14. Devanur, N.R., Huang, Z., Psomas, C.A.: The sample complexity of auctions with side information. arXiv preprint arXiv:1511.02296 (2015)
15. Dhangwatnotai, P., Roughgarden, T., Yan, Q.: Revenue maximization with a single sample. In: Proceedings of the 11th ACM Conference on Electronic Commerce. pp. 129–138. EC '10, ACM, New York, NY, USA (2010), http://doi.acm.org/10.1145/1807342.1807364
16. Fiat, A., Goldner, K., Karlin, A.R., Koutsoupias, E.: The FedEx Problem. In: Proceedings of the 2016 ACM Conference on Economics and Computation. pp. 21–22. EC '16, ACM, New York, NY, USA (2016), http://doi.acm.org/10.1145/2940716.2940752
17. Fu, H., Hartline, J., Hoy, D.: Prior-independent auctions for risk-averse agents. In: Proceedings of the fourteenth ACM conference on Electronic commerce. pp. 471–488. ACM (2013)
18. Giannakopoulos, Y., Koutsoupias, E.: Duality and optimality of auctions for uniform distributions. In: Proceedings of the 15th ACM Conference on Economics and Computation. pp. 259–276. EC '14 (2014), http://doi.acm.org/10.1145/2600057.2602883
19. Giannakopoulos, Y., Koutsoupias, E.: Selling two goods optimally. In: ICALP (2). Lecture Notes in Computer Science, vol. 9135, pp. 650–662. Springer (2015)
20. Hart, S., Nisan, N.: Approximate revenue maximization with multiple items. arXiv preprint arXiv:1204.1846 (2012)
21. Huang, Z., Mansour, Y., Roughgarden, T.: Making the most of your samples. In: Proceedings of the Sixteenth ACM Conference on Economics and Computation. pp. 45–60. ACM (2015)
22. Li, X., Yao, A.C.C.: On revenue maximization for selling multiple independently distributed items. Proceedings of the National Academy of Sciences 110(28), 11232–11237 (2013), http://dx.doi.org/10.1073/pnas.1309533110
23. Morgenstern, J., Roughgarden, T.: Learning simple auctions. CoRR abs/1604.03171 (2016), http://arxiv.org/abs/1604.03171
24. Morgenstern, J.H., Roughgarden, T.: On the pseudo-dimension of nearly optimal auctions. In: Advances in Neural Information Processing Systems. pp. 136–144 (2015)
25. Roughgarden, T., Talgam-Cohen, I.: Optimal and near-optimal mechanism design with interdependent values. In: Proceedings of the Fourteenth ACM Conference on Electronic Commerce. pp. 767–784. EC '13, ACM, New York, NY, USA (2013), http://doi.acm.org/10.1145/2482540.2482606
26. Roughgarden, T., Talgam-Cohen, I., Yan, Q.: Supply-limiting mechanisms. In: Proceedings of the 13th ACM Conference on Electronic Commerce. pp. 844–861. ACM (2012), http://theory.stanford.edu/~tim/papers/limit_conf.pdf
27. Rubinstein, A.: On the computational complexity of optimal simple mechanisms. In: Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science. pp. 21–28. ITCS '16, ACM, New York, NY, USA (2016), http://doi.acm.org/10.1145/2840728.2840736

28. Rubinstein, A., Weinberg, S.M.: Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. In: Proceedings of the Sixteenth ACM Conference on Economics and Computation. pp. 377–394. ACM (2015), http://arxiv.org/abs/1501.07637
29. Weinberg, S.M.: Algorithms for Strategic Agents. Ph.D. thesis, Massachusetts Institute of Technology (2014), http://people.csail.mit.edu/smweinberg/thesis.pdf
30. Yao, A.C.C.: An n-to-1 bidder reduction for multi-item auctions and its applications. In: Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms. pp. 92–109 (2015), http://arxiv.org/abs/1406.3278