Auction Design with a Revenue Target

Paul W. Goldberg[1] and Bo Tang[2](✉)

[1] Oxford University, Oxford, UK
[2] University of Liverpool, Liverpool, UK
bo.tang@liv.ac.uk

Abstract. In many fund-raising situations, a revenue target is specified. This suggests that the fund-raiser is interested in maximizing the probability to achieve this revenue target, rather than in maximizing the expected revenue. We study this topic from the perspective of Bayesian mechanism design, in a setting where a seller has a certain good that he can supply at no cost, and there are buyers whose joint valuation for the good comes from some given prior distribution. We present an algorithm to find the optimal truthful auction for two buyers with independent valuations via a direct characterization of the optimal auction. In contrast, we show the problem is NP-hard when the number of buyers is arbitrary or the distributions are correlated. Both negative results can be modified to show NP-hardness of designing auctions for risk-averse sellers.

Our main results address the design of *simple* auctions for many buyers, again in the context of a revenue target. For *Sequential Posted Price Auctions*, we provide a FPTAS to compute the optimal posted prices for a given sequence of buyers. For *Monopoly Price Auctions*, we apply the results of [8] on sparse covers of distributions to obtain a PTAS in a setting where the seller has a constraint on discriminatory pricing, consisting of a fixed set of prices he may use.

1 Introduction

There is a considerable literature on the algorithmic challenge of designing auctions that maximise the expected revenue obtained from a set of buyers. In this paper we consider a related objective where instead of maximising the expected revenue, the auctioneer has been given some revenue target T, and wishes to maximise the probability of raising at least T. This objective gives rise to new and interesting algorithmic challenges, and has some plausible real-world motivations, discussed below.

We work in the classical Bayesian setting of a collection of buyers whose valuations (prices they are willing to pay) for items being sold, are assumed to be drawn from some known prior distribution D. We are interested in designing mechanisms that are incentive compatible and individually rational. D in combination with a mechanism M results in a distribution over the revenue R obtained. A standard objective is to choose M to maximise the expected

P.W. Goldberg — Supported by EPSRC under grant EP/K01000X/1.

M. Hoefer (Ed.): SAGT 2015, LNCS 9347, pp. 137–149, 2015.
DOI: 10.1007/978-3-662-48433-3_11

value of R. A more general setting assumes a non-decreasing "utility of money" function u, and aims to maximize the expectation of $u(R)$. In this revenue-target setting, u is a shifted Heaviside function, equal to 0 for $R < T$ and 1 for $R \geq T$. Certain concave functions u have been used to model risk aversion, however the functions u considered here are not concave.

In this paper we focus on the "digital goods" setting, where the seller can supply unlimited copies of some good, at no cost. We also assume that the buyers have unit demand, so that a buyer's type is represented by a probability distribution over his valuation for a copy of the item. This special case is a simplified model of the fund-raising situations mentioned below. In the context of digital goods and unit demands, maximisation of the *expected* revenue can be decomposed into revenue-maximisation from each buyer independently. In contrast, when we switch to a revenue target, we find that the deal offered to a buyer should depend on the outcomes of the deals offered to other buyers.

This revenue-target setting is motivated by various real-world scenarios. Charitable fund-raising typically identify a target revenue to be raised. Similarly, in Internet crowd-sourcing platforms that support fund-raising for business start-ups (Kickstarter, Indiegogo, RocketHub etc.), it is typical to aim for some amount of money, and if that target is not reached, the would-be investors get their money back. (Our model doesn't properly capture this situation; we mention it to emphasise the importance of revenue targets in practice.) While a fund-raising effort is not the same thing as an auction, to some extent it can be modelled as one: an approach to a donor (or investor) corresponds to an attempt to sell an item to a would-be buyer. In cases where goods are sold at auction, it may be more desirable to raise a particular amount of money than to maximise the expected revenue. For example, in a bankruptcy situation, the administrator may wish to sell a collection of items so as to prioritise repaying the top-tier creditors. And while the FCC spectrum auction wants to raise as much money as possible, it is also required to cover its costs.

1.1 Our Results

We consider the problem parametrised by the number of buyers n, and the support size m of their value distributions. With multiple buyers, it is #P-complete to compute the exact success probability (probability to achieve revenue target T) for a given auction (Proposition 2). Given this obstacle, in Sect. 3 we consider a basic case of two buyers having uncorrelated valuations. We exhibit a polynomial-time algorithm to exactly compute the optimal truthful auction that maximises the probability to achieve T, given as input any discrete prior distributions. We do this via a structural characterisation of auctions that optimise the probability of achieving a given revenue target. This characterisation totally differs from the one maximising expected revenue and allows us to restrict to auctions with a geometric property that makes the problem tractable.

We show contrasting hardness results for correlated valuations or n buyers with independent distributions. Specifically, it is shown to be NP-complete to compute the optimal auction for three buyers having correlated valuations and

NP-hard for n buyers with independent distributions. Note that, in the latter case, a truthful auction may not necessarily be succinctly representable. We overcome the obstacle via proving the hardness for a class of succinct auctions and showing there exists a truthful auction with good performance if and only if there exists a good succinct auction in the constructed instance.

Our main algorithmic results are in Sect. 4, for two prevalent auctions following the trend of designing *simple* auctions. The first one is the *Sequential Posted Price Auction* introduced by Chawla et al. [4] to approximate the expected revenue in multi-dimensional Bayesian mechanism design. In this auction, the seller offers a take-it-or-leave-it price to each buyer sequentially. Given a sequence of buyers, we are able to provide a *fully polynomial-time approximation scheme* (FPTAS) to compute an approximately optimal sequential posted price auction that maximizes the success probability with an additive error. Second, we consider the *Monopoly Price Auction* where the seller offers take-it-or-leave-it prices to buyers simultaneously. This type of auctions was studied in [13] for selling goods with limited supply. We apply results of [8] on sparse covers of Poisson binomial distributions to obtain a PTAS when the seller has a limitation on discriminatory pricing, i.e., is only allowed to use few distinct prices.

1.2 Related Work

There has been a long line of research on maximizing expected revenue in Bayesian mechanism design starting from the seminal work by Myerson [15]. Recently, Cai et al [3] developed a general framework reducing revenue maximization to social welfare maximization. They also applied the framework to optimize certain non-linear functions [2]. However, the mechanisms they derived are randomized and Bayesian truthful, not deterministic truthful mechanisms studied in this paper.

Another line of research studied auction design for risk-averse sellers that can be regarded as maximizing a concave function of the revenue (cf. [17]). Sundararajan and Yan [18] studied the auction design problem for a risk averse seller and gave robust mechanisms (without knowledge of the concave function) which achieve constant approximations when buyers' distributions are independent. The approximation ratio has been improved to $e/(e-1)$ by Bhalgat et al. [1] by using the knowledge of concave functions. Our work complements their results by providing some corresponding intractability results.

We mention several negative results on revenue maximization in deterministic mechanism design. Diakonikolas et al. [10] showed that it is NP-hard to maximize revenue given a welfare constraint. Chen et al. [6] proved that it is NP-hard to maximize revenue in a multi-dimensional setting with a single unit-demand buyer when the valuations of items are independently distributed. For correlated buyers, Papadimitriou and Pierrakos [16] proved that it is NP-hard to approximate the optimal expected revenue for a single-item auction. However, in digital goods setting, the revenue maximizing auction can be constructed easily by computing the optimal price for each bidder separately based on their distributions conditioned on others' bids.

The study of digital goods auctions was initiated by Goldberg et al. [11]. Recently, Chen et al. [5] derived the optimal competitive auction with the benchmark defined to measure worst-case over all buyer profiles. In contrast, our benchmark measure is the average cases based on the prior distribution. Another related concept is "profit extractor" (see Sect. 6.2.4 in [12]) which is a decision problem the profit maximization in the prior-free setting.

Threshold probability maximization is a classical objective in stochastic optimization and has been studied for several combinatorial optimization problems (cf. [14] and references therein). However no incentive issues were considered before when optimizing this objective. The technique we apply to approximate the optimal monopoly price auction is based on [8]. These results have been shown helpful in computing Nash Equilibria [9] and learning sums of random variables [7]. But to our knowledge, this paper is their first application in auction design.

2 Preliminaries

Auction Setting. We study an auction environment where a seller wants to sell copies of an item to n bidders. Each bidder/buyer i is interested in a single copy of the item and values it at a privately known value v_i. A valuation profile $\mathbf{v}$ is the vector of all bidders' valuations, i.e. $\mathbf{v} = (v_1, \ldots, v_n)$. We consider a deterministic single-round sealed-bid auction where each bidder submits a bid b_i to express how much he is willing to pay for the item. After soliciting submitted bids $\mathbf{b} = (b_1, \ldots, b_n)$, the seller must decide whether each bidder i wins an item and how much he needs to pay. Bidder i's utility is the difference between his value v_i and his payment if he wins a item; otherwise he pays 0 and gets utility 0 to guarantee *individual rationality*, that is, no bidders will get a negative utility in the auction.

We assume every bidder in the auction is rational and aims to maximize his own utility by choosing the best bidding strategy. An auction is said to be truthful if for each bidder i, bidding his true valuation (i.e. $b_i = v_i$) is a dominant strategy no matter what the other bidders bid. It is known that truthful auctions can be characterized by *bid-independent* auctions where for each bidder i, the auction computes a threshold price p_i that does not depend on b_i but may depend on the bids of the other bidders $\mathbf{b}_{-i} = (b_1, \ldots, b_{i-1}, b_{i+1}, \ldots, b_n)$. In other words, there exists a pricing function for bidder i such that $p_i = f_i(\mathbf{b}_{-i})$ and i wins the item iff $b_i \geq p_i$ and his payment is p_i if he wins. So it suffices to consider bid-independent auctions when designing truthful auctions.

Thus any truthful or bid-independent auction A can be represented by n pricing functions $(f_1, \ldots, f_n)$ where f_i is the pricing function for bidder i which maps other bidders' valuations $\mathbf{v}_{-i}$ to the threshold price p_i. For convenience, we use $x_i(\mathbf{v})$ to denote the allocation rule of the auction, i.e. $x_i(\mathbf{v}) = 1$ if i wins an item when the valuation profile is $\mathbf{v}$; otherwise $x_i(\mathbf{v}) = 0$. Hence, the revenue of A on profile $\mathbf{v}$ is $R^A(\mathbf{v}) = \sum_{i\in[n]} x_i(\mathbf{v}) f_i(\mathbf{v}_{-i})$ where $[n]$ denotes the set $\{1, \ldots, n\}$. We also use $R_i^A(\mathbf{v})$ to denote the revenue of the auction A from bidder i, i.e., $R_i^A(\mathbf{v}) = x_i(\mathbf{v}) f_i(\mathbf{v}_{-i})$. We will omit A from the notation if the auction is clear from the context.

Representation of Prior Distribution. We assume the seller has prior knowledge of the bidders' valuations, which is represented by a distribution on the valuation profile $\mathbf{v}$. In particular, we use D to denote the distribution on the valuation profile and V to denote the support of D. We denote the probability that the valuation profile is $\mathbf{v}$ by $\Pr[\mathbf{v}]$ for all $\mathbf{v} \in V$. Obviously, the distribution D can be represented in the size of V (denoted by $|V|$ or $|D|$) by explicitly describing $\Pr[\mathbf{v}]$ for all $\mathbf{v} \in V$. We also use $V_i = \{v_i^1, \ldots, v_i^{m_i}\}$ to denote the set of all possible value of v_i in D, where m_i is $|V_i|$ and $v_i^1 < v_i^2 < \cdots < v_i^{m_i}$. For convenience, we define $v_i^0 = 0$ and assume $0 \in V_i$.

We say the bidders' valuations are independently distributed if D is a product distribution, i.e. $D = \times_{i \in [n]} D_i$ where D_i is the distribution on buyer i's valuations; otherwise they are correlated. For convenience, we say the bidders are independent (or correlated) according to whether their valuations are independently distributed. For independent bidders, D can be represented using space $O(n \cdot m)$ where $m = \max_i m_i$.

We consider a seller with revenue target T and his utility is 1 if the revenue raised in the auction is at least T; otherwise his utility is 0. Given an instance $\mathcal{I} = (D, T)$ with the profile distribution D and revenue target T, the seller's utility in an auction A is $\Pr_{\mathbf{v} \sim D}[R^A(\mathbf{v}) \geq T]$. We also call this value the *performance* of auction A on instance $\mathcal{I}$. So an auction is an optimal truthful auction for an instance $\mathcal{I}$ if no truthful auction can outperform A on the instance $\mathcal{I}$. Similarly, we say A is c-additive approximately optimal if no truthful auction can perform better than the performance of A plus a parameter c. It is without loss of generality to assume the range of pricing function for bidder i is V_i as shown in the following proposition. The intuition is that rounding prices up to the next valuation of the agent will not decrease the revenue of the auction.

Proposition 1. *For any distribution profile D and truthful auction A, there exists another truthful auction A' such that the range of pricing functions for bidder i in A' is V_i for all $i \in [n]$ and $R^{A'}(\mathbf{v}) \geq R^A(\mathbf{v})$ for all profiles $\mathbf{v}$.*

Simple Auctions. We consider two types of simple auctions called monopoly price auctions and sequential posted price auctions. A *monopoly price auction* is a truthful auction with pricing functions $(f_1, \ldots, f_n)$ where each function f_i depends only on the prior distribution D and not on the other bids $\mathbf{b}_{-i}$. We say an auction is a *sequential posted price auction* with respect to an order σ if f_i may depend on D together with the bids of buyers who precede i in σ, i.e. $(b_1, \ldots, b_{i-1})$ if buyers are indexed according to σ. The following proposition shows the hardness of evaluating the performance of a given monopoly price auction. This is proved via a reduction from counting the solutions of KNAPSACK.

Proposition 2. *Given a monopoly price auction for independent bidders, it is #P-complete to compute the probability of achieving a revenue target.*

3 Optimal Truthful Auction for Two Independent Bidders

Recall that any truthful auction for two bidders can be represented by two pricing functions f_1 and f_2. By Proposition 1, we only need to consider $f_1 : V_2 \to V_1$ which maps bidder 2's valuations to bidder 1's threshold prices and $f_2 : V_1 \to V_2$. First of all, we show that the general problem reduces to a restricted version where bidders' distributions have support $\{0,\ldots,m\}\times\{0,\ldots,m\}$ and the target revenue is m, for some positive integer m. The intuition is mapping values of one agent to indices and mapping values of the second agent to intervals of $T - v_1$.

Lemma 3. *Given any instance $\mathcal{I} = (D, T)$ with an independent profile distribution $D = D_1 \times D_2$ (D_i having support V_i) and a target revenue T, there exists an integer $m \le \min\{|V_1|, |V_2|\} + 1$ and another instance $\mathcal{I}' = (D', T')$ such that*

(a) $D' = D'_1 \times D'_2$ has the support $\{0,\ldots,m\} \times \{0,\ldots,m\}$ and $T' = m$
(b) Given an instance $\mathcal{I}$, the instance $\mathcal{I}'$ can be found in time linear in m
(c) Given any optimal truthful auction for $\mathcal{I}'$, it is possible to construct an optimal truthful auction for $\mathcal{I}$ in time linear in m.

For the case with two independent bidders, we assume $V_1 = V_2 = \{0,\ldots,m\}$ and $T = m$. We also use q_1^i and q_2^j to denote probabilities $\Pr[v_1 = i]$ and $\Pr[v_2 = j]$ respectively and $R(i,j)$ to be the revenue from the profile (i,j). Regarding pricing functions, we can assume $f_1(0) = m$ and $f_2(0) = m$, since otherwise we can increase $f_1(0)$ or $f_2(0)$ to m without loss of the objective. In the following lemmas, we show that there exists an optimal auction with several nice properties. The first one is monotonicity of f_1 and f_2. Intuitively, the lemma says once one bidder's valuation increases, the seller will get more revenue from this bidder and set a lower price for the other bidder as a consequence.

Lemma 4. *There exists an optimal truthful auction for two independent bidders such that the pricing functions are monotonically non-increasing.*

By Lemma 3 we assume the valuations of both bidders are in $\{0,\ldots,m\}$ and the target revenue is m. So for any profile $\mathbf{v}$ such that $v_1 < m$ and $v_2 < m$, the seller must sell items to both bidders to achieve the target revenue. Based on this observation, we are able to show another property of f_1 and f_2.

Lemma 5. *There exists an optimal truthful auction $A = (f_1, f_2)$ for two independent bidders such that f_1 is non-increasing and for any $i \in \{0,\ldots,m\}$,*

$$f_2(i) = \begin{cases} m & \text{if } \forall j \in \{0,\ldots,m\},\ i < f_1(j) \\ j & \text{if } \exists j \in \{0,\ldots,m\},\ f_1(j) \le i < f_1(j-1) \\ f_2(m-1) & \text{if } \forall j \in \{0,\ldots,m\},\ i \ge f_1(j),\ \text{i.e.i} = \text{m since } f_1(0) = \text{m} \end{cases}$$

Intuitively, the optimal auction described in the above lemma divides all profiles into four areas. In area one, the auction allocates nothing and in area two it sells both items. In area three (or four), the auction only sells a single copy with a price m to bidder 1 (or bidder 2). In addition, as shown in Fig. 1, the values of f_2 in this auction only depend on f_1. Thus, in order to design the optimal auction, we only need to find the optimal f_1, then a suitable f_2 follows by Lemma 5. Before characterizing the optimal f_1, we introduce some new notations. Given a non-increasing function f_1, let $J \subseteq [m]$ be the set of indices such that $f_2(j) < f_2(j-1)$. We denote the set J by $\{j_1, j_2, \ldots, j_{|J|}\}$ with an increasing order, i.e. $j_\ell < j_{\ell+1}$. Let $i_\ell = f_1(j_\ell)$ as illustrated in Fig. 2. We also define $i_0 = j_{|J|+1} = m + 1$ for simplicity. Then for all $\ell = 1, \ldots, |J|$ and $j_\ell \le j < j_{\ell+1}$, $f_1(j) = i_\ell$ by the definition of j_ℓ. In addition, for all $\ell = 1, \ldots, |J|$ and $i_\ell \le i < i_{\ell-1}$, $f_2(i) = j_\ell$ by Lemma 5. This is because j_ℓ is the j such that $f_1(j) \le i < f_1(j-1)$. Then we can prove the following lemma.

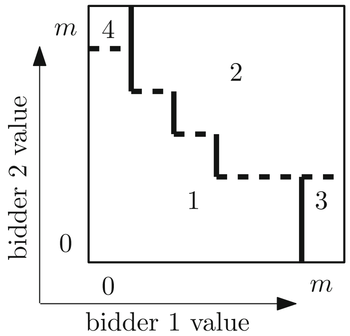

Fig. 1. Illustration of the computation of f_2 for a given f_1 based on Lemma 5. Again, the vertical bold lines are f_1 and the horizontal dashed lines are the resulting f_2. We also mark the four areas mentioned in the text.

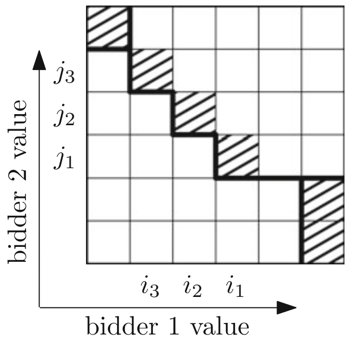

Fig. 2. Illustration of the definition of the set J, the values j_ℓ and i_ℓ when the pricing functions f_1 and f_2 are given as vertical and horizontal bold lines respectively. The shawed squares illustrate the profiles with revenue at least m.

Lemma 6. *There exists an optimal auction $A = (f_1, f_2)$ such that $i_\ell + j_\ell = m$ for all $\ell = 1, \ldots, |J|$ where i_ℓ and j_ℓ are defined by f_1 as above.*

By the above lemma, we can characterize the optimal auction by only using the set J, i.e. the values of $\{j_1, \ldots, j_{|J|}\}$. Given the set J, we can compute f_1 and f_2 by Lemmas 6 and 5 respectively. Based on this characterization, we are able to show the main theorem in this section.

Theorem 7. *Given a distribution $D = D_1 \times D_2$ for two independent bidders and a target revenue for the seller, an optimal truthful auction can be found in time $O(m^3)$ where $m = \min\{|D_1|, |D_2|\}$.*

We have contrasting NP-hardness for more general cases. Both results can be modified for the cases with risk-averse sellers.

Theorem 8. *It is NP-complete to compute an optimal auction for three correlated bidders, having a joint prior distribution presented as a set of probabilities on a finite set of support points.*

Theorem 9. *It is NP-hard to compute the optimal auction for n independent bidders even when each bidder has only two possible valuations, i.e. $|V_i| = 2$.*

4 Near-Optimal Simple Auctions for Independent Bidders

In this section, we study the following simple auctions for sellers with a target revenue when the bidders are independent. In Sect. 4.1, we present an additive FPTAS for computing approximately optimal *sequential posted price* auctions with respect to a fixed order σ. Then in Sect. 4.2 we show an additive PTAS for optimal *monopoly price auctions*, in a setting where the seller is restricted to using a constant number of distinct prices.

4.1 Approximately Optimal Sequential Posted Price Auction

We first present a pseudo-polynomial time algorithm to compute optimal sequential posted prices via dynamic programming. Then we show that this algorithm can be modified to be a FPTAS with respect to additive error. We order the bidders with respect to the fixed order σ.

Recall that in a sequential posted price mechanism, the seller offers take-it-or-leave-it prices to the buyers sequentially with respect to a given order σ and the computation of the price for buyer i is based on the results of all buyers preceding i, together with the valuation distributions. Note that the optimal sequential posted price for any sequence of buyers, performs at least as well as the optimal monopoly price auction. In contrast with the objective of expected revenue maximization, our objective of a target revenue means that the price offered to bidder i may depend on the revenue gained from the first $i-1$ bidders. This allows us to solve the problem by the following dynamic programming. Let $Q[i, r]$ be the maximal probability to achieve revenue r by selling items to buyers from i to n. By Proposition 1, it is sufficient to consider the case that $p_i \in V_i$ where V_i is the support of buyer i's valuation distribution. It is easy to see $Q[i, r] = 1$ if $r \leq 0$ and $Q[i, r] = 0$ if $i > n$ and $r > 0$. For the other cases when $i \leq n$ and $r > 0$ we have

$$Q[i, r] = \max_{p_i \in V_i} \{Q[i + 1, r - p_i] \cdot \Pr[v_i \geq p_i] + Q[i + 1, r] \cdot (1 - \Pr[v_i \geq p_i])\}.$$

Thus the maximal probability to achieve target revenue T from all buyers is $Q[1, T]$. Note that solving the above dynamic programming gives a pseudo-polynomial time algorithm for the problem. Actually, we can get an additive FPTAS by rounding the dynamic programming properly.

Theorem 10. *There exists an additive FPTAS for computing approximately optimal sequential posted price auctions with respect to a fixed order of the buyers. In particular, given* $\epsilon \in (0,1)$*, an instance* $\mathcal{I} = (D, T)$ *with* n *independent buyers and a buyer sequence* σ*, an* ϵ*-additive approximately optimal sequential posted price auction with respect to* σ *can be computed in time* $O(m^2 n^2 \log n \cdot 1/\epsilon \log(1/\epsilon))$ *where* m *is the maximal support size, i.e.* $\max_{i\in[n]}\{|D_i|\}$.

4.2 Approximately Optimal Monopoly Price Auction

In this section, we present a PTAS for computing the optimal monopoly price auction when the seller is restricted to a given constant-sized set of distinct prices, and for each buyer has to select one of those prices for that buyer. Recall that in a monopoly price auction, the seller offers those take-it-or-leave-it prices to the buyers simultaneously, and the prices are only based on the valuation distributions. Our PTAS uses results of [8] on Poisson Binomial Distributions. First of all, we review the definitions and results. For any two random variables X and Y supported on a finite set A, their total variation distance is defined as

$$d_{\mathrm{TV}}(X, Y) = \frac{1}{2} \sum_{a \in A} |\Pr[X = a] - \Pr[Y = a]|.$$

We use the following result in the proof of Theorems 14 and 15.

Lemma 11 (Lemma 2 in [8]). *Let* $X_1, \ldots, X_n$ *be mutually independent random variables, and let* $Y_1, \ldots, Y_n$ *be mutually independent random variables. Then*

$$d_{\mathrm{TV}}(\sum_{i=1}^{n} X_i, \sum_{i=1}^{n} Y_i) \le \sum_{i=1}^{n} d_{\mathrm{TV}}(X_i, Y_i).$$

A distribution is said to be a *Poisson Binomial Distribution* (PBD) of order n if it is a discrete probability distribution consisting of the sum of n independent indicator random variables. The distribution is parameterized by a vector $(r_i)_{i=1}^n \in [0,1]^n$ of probabilities and is denoted by $\mathrm{PBD}(r_1, \ldots, r_n)$. Let $\mathcal{S}_n$ be the set of all PBDs of order n. We review a construction of an efficient and proper ϵ-cover for $\mathcal{S}_n$.

Theorem 12 (Theorem 1 in [8]). *For all* $n, \epsilon > 0$*, there exists a set* $\mathcal{S}_{n,\epsilon} \subset \mathcal{S}_n$ *such that*

1. $\mathcal{S}_{n,\epsilon}$ *is an* ϵ*-cover of* $\mathcal{S}_n$ *in total variation distance; that is, for all* $D \in \mathcal{S}_n$*, there exists some* $D' \in \mathcal{S}_{n,\epsilon}$ *such that* $d_{\mathrm{TV}}(D, D') \le \epsilon$*,*
2. $|S_{n,\epsilon}| \le n^2 + n \cdot (\frac{1}{\epsilon})^{O(\log^2 1/\epsilon)}$*,*
3. $S_{n,\epsilon}$ *can be computed in time* $O(n^2 \log n) + O(n \log n) \cdot (\frac{1}{\epsilon})^{O(\log^2 1/\epsilon)}$.

Moreover, all distributions $\mathrm{PBD}(r_1, \ldots, r_n) \in \mathcal{S}_{n,\epsilon}$ *in the cover satisfy at least one of the following properties, for some positive integer* $t = t(\epsilon) = O(1/\epsilon)$.

- *(t-sparse form) there is some* $\ell \le t^3$ *such that, for all* $i \le \ell$*,* $r_i \in \{\frac{1}{t^2}, \frac{2}{t^2}, \ldots, \frac{t^2-1}{t^2}\}$ *and for all* $i > \ell$*,* $r_i \in \{0, 1\}$*; or*

– *((n,t)-Binomial form) there is some $\ell \in [n]$ and $q \in \{\frac{1}{n}, \frac{2}{n}, \ldots, \frac{n}{n}\}$ such that, for all $i \le \ell$, $r_i = q$ and for all $i > \ell$, $r_i = 0$; moreover ℓ and q satisfy $\ell q \ge t^2$ and $\ell q(1-q) \ge t^2 - t - 1$.*

In words, every PBD can be approximated by either a sparse PBD or a binomial distribution. Moreover, the following theorem tells us that if the first $O(\log 1/\epsilon)$ moments of two PBDs are the same, then the total variation distance between them is at most ϵ.

Theorem 13 (Theorem 3 in [8]). *Let $\mathcal{P} := (p_i)_{i=1}^n \in [0,1/2]^n$ and $\mathcal{Q} := (q_i)_{i=1}^n \in [0,1/2]^n$ be two collections of probability values. Let also $\mathcal{X} := (X_i)_{i=1}^n$ and $\mathcal{Y} := (Y_i)_{i=1}^n$ be two collections of mutually independent indicators with $\mathrm{E}[X_i] = p_i$ and $\mathrm{E}[Y_i] = q_i$, for all $i \in [n]$. If for some $d \in [n]$ the following condition is satisfied: $\sum_{i=1}^n p_i^\ell = \sum_{i=1}^n q_i^\ell$ for all $\ell = 1, \ldots, d$, then $d_{\mathrm{TV}}(\sum_i X_i, \sum_i Y_i) \le 13(d+1)^{1/4} 2^{-(d+1)/2}$.*

It is easy to see that Theorem 13 holds if we replace $[0,1/2]$ with $[1/2,1]$. Moreover, by setting $d = O(\log 1/\epsilon)$, this bound becomes at most ϵ. Theorem 12 shows that there exists an efficient cover for the set of all PBDs. However, we cannot directly apply this theorem to our problem, since (given prices and prior distributions of a problem instance) the set of associated PBDs (call it S) is a proper subset of $\mathcal{S}_n$, and we need to find a cover that consists of a subset of S. Theorem 14 is intended to overcome this obstacle. Given n finite sets $W_1, \ldots, W_n$ where $W_i \subset [0,1]$ for all $i \in [n]$, let $W = \times_{i=1}^n W_i$, and let $\mathcal{S}_n(W)$ denote the set of all PBDs such that the probability of the indicator i is in W_i for all $i \in [n]$. That is $\mathcal{S}_n(W) = \{\mathrm{PBD}(r_1, \ldots, r_n) | (r_i)_{i=1}^n \in W\}$.

Theorem 14. *For all $n, \epsilon > 0$ and any n finite subsets of $[0,1]$, $W_1, \ldots, W_n$ let $W = \times_{i=1}^n W_i$. Then there exists a set $\mathcal{S}_{n,\epsilon}(W) \subset \mathcal{S}_n(W)$ such that*

1. *$\mathcal{S}_{n,\epsilon}(W)$ is an ϵ-cover of $\mathcal{S}_n(W)$ in total variation distance; that is, for all $D \in \mathcal{S}_n(W)$, there exists some $D' \in \mathcal{S}_{n,\epsilon}(W)$ such that $d_{\mathrm{TV}}(D, D') \le \epsilon$,*
2. *$\mathcal{S}_{n,\epsilon}(W)$ can be computed in time $(\frac{n}{\epsilon})^{O(\log^2 1/\epsilon)}$ and has size at most $(\frac{n}{\epsilon})^{O(\log^2 1/\epsilon)}$.*

Given the above theorem, we can obtain an additive PTAS for computing approximately optimal monopoly price auctions, given a fixed set of allowed prices.

Theorem 15. *There exists an additive PTAS for computing approximately optimal monopoly price auctions when the seller is restricted to a fixed number of distinct prices. In particular, given $\epsilon \in (0,1)$, an instance with n independent bidders and k distinct prices the seller may use, an ϵ-additive approximately optimal monopoly price auction can be computed in time $(\frac{nk}{\epsilon})^{O(k \log^2 1/\epsilon)}$.*

Proof. We use $a_1, \ldots, a_k$ to denote the k distinct prices the seller may use. Given a monopoly price auction with price vector $(p_1, p_2, \ldots, p_n)$, we use an indicator random variable H_{ij} to indicate that the seller gets revenue a_j from buyer i, that is $H_{ij} = 1$ iff $p_i = a_j$ and $v_i \ge a_j$. Let $H_j = \sum_{i \in [n]} H_{ij}$ and

$H = \sum_{j \in [k]} a_j H_j$. Note that H is the random variable for the total revenue raised in this auction. Since the H_{ij} are indicator random variables, the H_j are Poisson Binomial random variables due to the independence among bidders. So H can be viewed as a weighted sum of k Poisson Binomial random variables. Let r_{ij} denote the probability of getting revenue exactly a_j from buyer i. Then the distribution of H_j is $\mathrm{PBD}(r_{1j}, \cdots, r_{nj})$. The distribution of H can be represented by the vector $\mathbf{r} = (r_{ij})_{i \in [n], j \in [k]}$. Let W_i be the set of all possible $(r_{i1}, \ldots, r_{in})$ such that $r_{ij} = \Pr[v_i \geq a_j]$ if the seller use price a_j for bidder i and $r_{ij} = 0$ otherwise . It is clear that the set $W = \times_{i \in [n]} W_i$ is the set of all probability vector $\mathbf{r}$ corresponding to a feasible pricing vector $\mathbf{p}$.

Note that for any two random variables X, Y and any value T, $|\Pr[X \geq T] - \Pr[Y \geq T]| \leq d_{\mathrm{TV}}(X, Y)$. So if there exists an ϵ-cover for the set of all possible distribution of H parameterized by $\mathbf{r} \in W$, we can explore the pricing rules in the cover instead of all possible pricing rules to find a sequence of monopoly prices which approximately maximize $\Pr[H \geq T]$. In order to get such a cover, we need to modify the dynamic programming used in the proof of Theorem 14 to be k-dimensional. The moment profile $(\mu^1, \ldots, \mu^k, \nu^1, \ldots, \nu^k)$ is defined as $\mu^j = (\mu^j_1, \ldots, \mu^j_d), \nu^j = (\nu^j_1, \ldots, \nu^j_d)$ and $\mu^j_\ell, \nu^j_\ell \in \{0, (\frac{\epsilon}{nk})^\ell, 2(\frac{\epsilon}{nk})^\ell, \ldots, n\}$ for all $\ell \in [d]$ and $j \in [k]$. By a similar argument to Theorem 14 and Lemma 11, all the possible moment profiles is already an ϵ-cover. Define $A[i, \mu^1, \ldots, \mu^k, \nu^1, \ldots, \nu^k]$ to be the indicator such that it is equal to 1 iff there exists $\mathbf{r}_1 \in W_1, \ldots, \mathbf{r}_i \in W_i$ such that for all $j \in [k]$ and $\ell \in [d]$, $\sum_{i' \leq i : r'_{i'j} \in [0,1/2]} (r'_{i'j})^\ell = \mu^j_\ell$ and $\sum_{i' \leq i : r'_{i'j} \in (1/2,1]} (r'_{i'j})^\ell = \nu^j_\ell$ where $\mathbf{r'}$ is a $\frac{\epsilon}{nk}$-rounding of $\mathbf{r}$ such that r'_{ij} is a multiple of $\frac{\epsilon}{nk}$ and $r_{ij} - \frac{\epsilon}{nk} < r'_{ij} \leq r_{ij}$ for all $i \in [n]$ and $j \in [k]$.

Similarly to the proof of Theorem 14, A can be computed by the following dynamic programming. Inductively, to compute layer $i+1$, we consider all the non-zero entries of layer i and for every such non-zero entry and every possible prices a_j, we find which entry of layer $i+1$ we would transition to if we choose $p_i = a_j$, i.e. $r_{ij} = \Pr[v_i \geq a_j]$ and $r_{ij'} = 0$ for all $j' \neq j$. It is easy to see the overall running time to compute A is $(\frac{nk}{\epsilon})^{O(k \log^2 1/\epsilon)}$. In addition, we can find the corresponding monopoly prices for any distribution in this cover by tracing the pointers in the computation of A. Therefore, we can enumerate all possible pricing rules in this cover with size at most $(\frac{nk}{\epsilon})^{O(k \log^2 1/\epsilon)}$ to find the optimal pricing which maximize $\Pr[H \geq T]$.

The final step is to compute $\Pr[H \geq T]$ given a price vector $\mathbf{p}$. By Theorem 12, we know any PBD can be approximated by a sparse PBD or a binomial distribution. For the given price vector, we can get the corresponding H_j for all $j \in [k]$. We use Theorem 12 to compute H'_j from H_j such that H'_j is either a k/ϵ-sparse PBD or a binomial distribution and $d_{\mathrm{TV}}(H'_j, H_j) \leq \epsilon/k$ for all $j \in [k]$. Then we compute $\Pr[H'_j = T_j]$ for any value $T_j \in [0, \ldots, n]$ and $j \in [k]$. This computation can be done efficiently since H'_j is either a k/ϵ-sparse PBD or a binomial distribution. By Lemma 11, we have $d_{\mathrm{TV}}(H', H) \leq \epsilon$ where $H' = \sum_j a_j H_j$. Finally we compute $\Pr[H' \geq T] = \sum_{(T_j)_j : \sum_j a_j T_j \geq T} \prod_j \Pr[H'_j = T_j]$ by enumerating all possible $T_1, \ldots, T_k$. Since the distance between H and H' is at most ϵ, we have $\Pr[H \geq T] \geq \Pr[H' \geq T] - \epsilon$. Combine all these together, we get the additive PTAS with running time $(\frac{nk}{\epsilon})^{O(k \log^2 1/\epsilon)}$. $\square$

5 Conclusion

We see several promising directions for future work. For independent buyers, a direct open problem is to generalize our characterization to three or more buyers. That may be achievable via an induction on the number of buyers, characterizing the optimal auction for three buyers by using the case with two buyers as a substructure. Another direction is to approximate the optimal auction via designing simple auctions. We find several examples to show the lower bounds (see full version for more details) but the upper bound is still open. Finally, we point out an interesting problem of computing optimal monopoly prices without the limitation on distinct prices.

References

1. Bhalgat, A., Chakraborty, T., Khanna, S.: Mechanism Design for a Risk Averse Seller. In: Goldberg, P.W. (ed.) WINE 2012. LNCS, vol. 7695, pp. 198–211. Springer, Heidelberg (2012)
2. Cai, Y., Daskalakis, C., Weinberg, S.: Understanding incentives: Mechanism design becomes algorithm design, In: FOCS 2013, pp. 618–627. IEEE, October 2013
3. Cai, Y., Daskalakis, C., Weinberg, S.M.: Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In: FOCS 2012, pp. 130–139. IEEE Computer Society, Washington (2012)
4. Chawla, S., Hartline, J.D., Malec, D.L., Sivan, B.: Multi-parameter mechanism design and sequential posted pricing. In: STOC 2010, New York, NY, pp. 311–320 (2010)
5. Chen, N., Gravin, N., Lu, P.: Optimal competitive auctions. In: STOC 2014, pp. 253–262. ACM, New York (2014)
6. Chen, X., Diakonikolas, I., Paparas, D., Sun, X., Yannakakis, M.: The complexity of optimal multidimensional pricing. In: SODA 2014, pp. 1319–1328. SIAM (2014)
7. Daskalakis, C., Diakonikolas, I., Servedio, R.A.: Learning poisson binomial distributions. In: STOC 2012, pp. 709–728. ACM, New York (2012)
8. Daskalakis, C., Papadimitriou, C.: Sparse covers for sums of indicators. Probab. Theory Relat. Fields **162**, 679–705 (2014)
9. Daskalakis, C., Papadimitriou, C.H.: Approximate Nash equilibria in anonymous games. J. Econ. Theory **156**, 207–245 (2015)
10. Diakonikolas, I., Papadimitriou, C., Pierrakos, G., Singer, Y.: Efficiency-revenue trade-offs in auctions. In: Czumaj, A., Mehlhorn, K., Pitts, A., Wattenhofer, R. (eds.) ICALP 2012, Part II. LNCS, vol. 7392, pp. 488–499. Springer, Heidelberg (2012)
11. Goldberg, A.V., Hartline, J.D., Wright, A.: Competitive auctions and digital goods. In: SODA 2001, pp. 735–744. SIAM, Philadelphia (2001)
12. Hartline, J.D.: Mechanism design and approximation. Book draft, October 2013

13. Hartline, J.D., Roughgarden, T.: Simple versus optimal mechanisms. SIGecom Exch. **5:8**(1), 1–5:3 (2009)
14. Li, J., Yuan, W.: Stochastic combinatorial optimization via poisson approximation, In: STOC 2013, pp. 971–980. ACM, New York (2013)
15. Myerson, R.B.: Optimal auction design. Math. Oper. Res. **6**(1), 58–73 (1981)
16. Papadimitriou, C.H., Pierrakos, G.: On optimal single-item auctions. In: STOC 2011, pp. 119–128. ACM, New York (2011)
17. Rothschild, M., Stiglitz, J.E.: Increasing risk: I. A definition. J. Econ. Theory **2**(3), 225–243 (1970)
18. Sundararajan, M., Yan, Q.: Robust mechanisms for risk-averse sellers. In: EC 2010, pp. 139–148. ACM, New York (2010)