

Optimal Pricing Is Hard

Constantinos Daskalakis^{1,*}, Alan Deckelbaum^{2,**}, and Christos Tzamos^{1,***}

¹ MIT, EECS, CSAIL
costis@csail.mit.edu, tzamos@mit.edu

² MIT, Math, CSAIL
deckel@mit.edu

Abstract. We show that computing the revenue-optimal deterministic auction in unit-demand single-buyer Bayesian settings, i.e. the optimal item-pricing, is computationally hard even in single-item settings where the buyer’s value distribution is a sum of independently distributed attributes, or multi-item settings where the buyer’s values for the items are independent. We also show that it is intractable to optimally price the grand bundle of multiple items for an additive bidder whose values for the items are independent. These difficulties stem from implicit definitions of a value distribution. We provide three instances of how different properties of implicit distributions can lead to intractability: the first is a $\#P$ -hardness proof, while the remaining two are reductions from the SQRT-SUM problem of Garey, Graham, and Johnson [14]. While simple pricing schemes can oftentimes approximate the best scheme in revenue, they can have drastically different underlying structure. We argue therefore that either the specification of the input distribution must be highly restricted in format, or it is necessary for the goal to be mere approximation to the optimal scheme’s revenue instead of computing properties of the scheme itself.

1 Introduction

Designing auctions to maximize revenue in a Bayesian setting is a problem of high importance in both theoretical and applied economics [19–21]. While substantial progress has been made on designing mechanisms with revenue guarantees that are approximately optimal [4, 6, 9, 10], the question of determining the optimal mechanism exactly has been much more intricate [1, 2, 7, 8, 11, 16, 17, 22].

In this paper, we study the complexity of designing optimal deterministic auctions for single-bidder problems, i.e. optimal pricing mechanisms. Prior to our work, Briest showed that finding the optimal pricing mechanism for a unit-demand bidder is highly inapproximable when the bidder’s values for different

* Supported by a Microsoft Research Fellowship, a Sloan Foundation Fellowship and NSF Award CCF-0953960 (CAREER) and CCF-1101491.

** Supported by Fannie and John Hertz Foundation Daniel Stroock Fellowship and NSF Award CCF-1101491.

*** Supported by NSF Award CCF-1101491.

items are correlated [5]. Our work complements his by either considering single-item settings, or multi-item settings with product value distributions. We also investigate the complexity of optimally pricing the grand bundle of multiple items for an additive buyer whose values for the items are independent. For these problems we demonstrate that even when the optimal mechanism can only be one of two possibilities, it can be computationally difficult to determine which one achieves the highest expected revenue.

We note that all hard instances presented in this paper have fully polynomial-time approximation schemes, and thus our results preclude exact algorithms but not computationally efficient approximation schemes. From a practical perspective, a nearly optimal mechanism may be almost as desirable as an exact one. From a theoretical perspective, however, it is important to understand the structure of the exactly optimal mechanism [20], which may be drastically different than that of approximate ones. Computational barriers to determining the best mechanism, such as the ones presented here, reflect barriers to understanding its structure.

Our results suggest in particular that great care must be taken in how a bidder's value distributions are specified. Intricate distributions can be described succinctly, providing a simple outlet to encode computationally hard problems. We present three concrete scenarios where succinctly-represented distributions lead to computational hardness: Easy-to-describe discrete distributions may have exponential size support, may have mild irrationality in their support, or have mild irrationality in the probabilities they assign. Indeed, many (or all) of these features of discrete distributions can be present in simple continuous distributions. Thus, to obtain a robust theory of optimal Bayesian mechanism design, we must either aim for only approximate revenue guarantees or severely limit the types and specification format of allowable value distributions.

2 Preliminaries

In our model, there is a seller with n items and a buyer whose values for the items v_1, \dots, v_n are random variables drawn from known distributions F_1, \dots, F_n . We will consider both *unit-demand* and *additive* buyer types:

- A (quasi-linear) unit-demand buyer is interested in buying at most one item; if the item prices are p_1, \dots, p_n , the buyer buys the item maximizing his utility, $v_i - p_i$, as long as it is positive, breaking ties among the maximizers in some pre-determined way, e.g. lexicographic or in favor of the cheapest/most expensive item.
- A (quasi-linear) additive buyer values a subset S of items $\sum_{i \in S} v_i$. If subset S is priced P_S , his utility for buying that subset is $\sum_{i \in S} v_i - P_S$. The buyer buys the subset of items that maximizes his utility, as long as it is positive, breaking ties among subsets in some pre-determined way.

In the case of a unit-demand bidder, the seller's goal is to price the items to optimize the expected price paid by the buyer. Finding the optimal such prices is called the *unit-demand pricing problem*. In the case of an additive bidder,

the seller’s goal is to price all subsets of items to optimize the expected price paid by the buyer. Of course, the seller may not want to explicitly list the price of every subset but describe their prices in some succinct manner, or may want to offer only some subsets at a finite price. We are particularly interested in the *grand bundle pricing problem* where the seller wants to optimally price the set of all items (the grand bundle) and the buyer must take all items or nothing. As shown in [18], pricing just the grand bundle is optimal in several natural settings. Furthermore, it oftentimes achieves revenue close to the optimal mechanism [3, 15]. Optimally pricing the grand bundle is furthermore interesting in its own right [13].

Finally, the distributions F_1, \dots, F_n may be provided explicitly, by listing their support and the probabilities placed on each point in the support, or implicitly giving a closed-form formula for them. In this paper, we study how various ways to describe the distributions affect the complexity of the pricing problem.

3 Complexity of Sum-of-Attributes Distributions

We first consider the problem of optimally pricing a single item for a single buyer whose value for the item is a sum of independent random variables. The probability distribution of the item’s value has an exponentially sized support, but has a succinct description in terms of each component variable’s distribution. The seller must choose a price P for the item. The buyer will accept the offer (and pay P) if his value for it is at least P , and will reject the offer (giving the seller zero revenue) if his value is strictly less than P . The seller’s goal is to choose P to maximize his expected revenue. In fact it follows from Myerson [20] that pricing the item at the optimal price is the optimal mechanism in this setting, even among randomized mechanisms.

This problem occurs fairly naturally. When selling a complex product (for example, a car), there are a number of attributes (color, size, etc) that a buyer may or may not value highly, and his value for the product may be the sum of his values for the individual attributes. If his values for the attributes are independent, the buyer’s value for the product can be modeled as a sum of independent random variables.

Formally, the problem we study in this section is the following.

Definition 1 (The Sum-of-Attributes Pricing (SoAP) Problem). *Given n pairs of nonnegative integers $(u_1, v_1), (u_2, v_2), \dots, (u_n, v_n)$ and rational probabilities p_1, p_2, \dots, p_n , determine the price P^* which maximizes $P^* \cdot \Pr[\sum_{i=1}^n X_i \geq P^*]$, where the X_i are independent random variables taking value u_i with probability p_i and v_i with probability $1 - p_i$.*

Notice that we can always view an instance of the sum-of-attributes pricing problem as an instance of the grand bundle pricing problem where we seek the optimal price to sell the “grand bundle” of a collection of n items that are independently distributed.

Theorem 1. *The Sum-of-Attributes Pricing problem and the Grand Bundle Pricing problem are $\#P$ -hard.*

Proof. We show how to use oracle access to the SoAP problem to solve the counting analog of the SUBSET-SUM problem, defined next, which is #P-complete.¹

#-SUBSET-SUM: Given as input a set of positive integers $\{a_1, a_2, \dots, a_n\}$ and a positive integer $T \leq \sum_i a_i$, the goal is to determine the number of subsets of the a_i 's which sum to at least T .

The idea of our reduction is to design an instance of the SoAP problem with $n+1$ attributes for which the optimal price is one of two possible prices. A single parameter (in particular, the probability p_{n+1} of the last attribute) determines which of these two prices is optimal. By repeatedly querying a SoAP oracle with varying values of p_{n+1} , we can determine the exact threshold value of p_{n+1} , which provides sufficient information to deduce the answer to the #-subset sum instance.

We proceed to provide the details of our reduction. Given an instance of the #-subset sum problem, we create an instance of SoAP with $n + 1$ attributes, where for all $i \in \{1, \dots, n\}$ we take $u_i = a_i$ and $v_i = 0$, while for the last attribute we take $u_{n+1} = T + 1$ and $v_{n+1} = 1$. Moreover, for all $i \in \{1, \dots, n\}$, we set

$$p_i \triangleq \frac{1}{2^n n(n + 1 + \sum_{j=1}^n a_j)^2}.$$

Notice in particular that the first n attributes have the same probability of taking their highest value. Moreover, the probability that all the first n attributes have value 0 is:

$$\left(1 - \frac{1}{2^n n(n + 1 + \sum_{j=1}^n a_j)^2}\right)^n > 1 - \frac{1}{2^n(n + 1 + \sum_{j=1}^n a_j)^2}$$

i.e. very close to 1. We leave the probability p_{n+1} that the last attribute takes its highest value a free parameter, which we denote by p for convenience.

Now, suppose that we use price B for the SoAP instance. We claim the following:

1. If $B = 1$, the expected revenue is 1.
2. If $1 < B < T + 1$, then the expected revenue is at most

$$B \left(p + \frac{1 - p}{2^n(n + 1 + \sum_{j=1}^n a_j)^2} \right).$$

3. If $B = T + 1$, then the expected revenue is at least $p(T + 1)$.
4. If $T + 1 < B \leq T + 1 + \sum_{j=1}^n a_j$, then the expected revenue is at most

$$\left(T + 1 + \sum_{i=1}^n a_i \right) \left(\frac{1}{2^n(n + 1 + \sum_{j=1}^n a_j)^2} \right) \leq \frac{1 + \sum_{j=1}^n a_j}{2^{n-1}(n + 1 + \sum_{j=1}^n a_j)^2} < 1.$$

5. If $B > T + 1 + \sum_{j=1}^n a_j$, then the expected revenue is 0.

¹ Indeed, the reduction from SAT to SUBSET-SUM as presented in [23] is parsimonious.

The fourth and fifth cases are never optimal, since they are both dominated by using $B = 1$. We claim that the second case is also never optimal. Suppose for the sake of contradiction that some integral price B strictly between 1 and $T + 1$ were optimal. Then we would have the following two constraints:

$$\begin{aligned} - B \left(p + \frac{1-p}{2^n(n+1+\sum_{j=1}^n a_j)^2} \right) &\geq 1 \\ - B \left(p + \frac{1-p}{2^n(n+1+\sum_{j=1}^n a_j)^2} \right) &\geq (T + 1)p. \end{aligned}$$

To show a contradiction, define for convenience

$$\epsilon \triangleq \frac{1}{2^n(n + 1 + \sum_{j=1}^n a_j)^2}.$$

We will show that no value of p exists for which both of the above constraints are simultaneously satisfied. From the first constraint, we deduce $p + \epsilon(1 - p) \geq 1/B$ and thus

$$p \geq \frac{1/B - \epsilon}{1 - \epsilon} \geq \frac{1/T - \epsilon}{1 - \epsilon} > 1/T - \epsilon,$$

where for the last inequality we used that $T \leq \sum_{j=1}^n a_j$. Moreover,

$$1/T - \epsilon \geq \frac{1}{\sum_{j=1}^n a_j} - \frac{1}{2^n(n + 1 + \sum_{j=1}^n a_j)^2} \geq \frac{1}{\sum_{j=1}^n a_j} - \frac{1}{2^n \sum_{j=1}^n a_j} \geq \frac{1}{2 \sum_{j=1}^n a_j}.$$

Therefore, the first constraint implies that $p > \frac{1}{2 \sum_{j=1}^n a_j}$. From the second constraint, we deduce $B(p + \epsilon(1 - p)) \geq (T + 1)p$ and thus

$$p \leq \frac{B\epsilon}{T + 1 - B + B\epsilon},$$

where we used that $B \leq T$ so $T + 1 - B + B\epsilon > 1$. We further have

$$p < B\epsilon \leq T\epsilon \leq \sum_{j=1}^n a_j \epsilon = \frac{\sum_{j=1}^n a_j}{2^n(n + 1 + \sum_{j=1}^n a_j)^2} < \frac{1}{2 \sum_{j=1}^n a_j}.$$

We get a contradiction as both constraints on p cannot be satisfied simultaneously. In summary, we have shown the following:

“For any p , the optimal price is either 1 or $T + 1$.”

We also note the following monotonicity property. If, for some p , the optimal price is $T + 1$, then the optimal price is $T + 1$ for any $p' > p$.² Therefore, there exists a unique p^* for which the expected revenue of selling at price $T + 1$ is exactly the same as the expected revenue of selling at price 1.

² This follows from the fact that the expected revenue from selling at $T + 1$ will only increase as p increases.

Suppose that we knew some p^* such that the expected revenue of selling at $T+1$ is exactly 1. Then, if we denote by V_n the total value of the first n attributes, p^* should satisfy:

$$1 = (T + 1) (p^* + (1 - p^*)P[V_n \geq T]);$$

so

$$P[V_n \geq T] = \frac{1/(T + 1) - p^*}{1 - p^*}.$$

Therefore, it is simple arithmetic to compute $P[V_n \geq T]$ from p^* . We also note that

$$P[V_n \geq T] = \sum_{k=0}^n p_1^k (1 - p_1)^{n-k} \cdot S(k, T) = p_1^n \cdot \sum_{k=0}^n \left(\frac{1 - p_1}{p_1}\right)^{n-k} \cdot S(k, t),$$

where $S(k, T)$ is the number of size k subsets of the a_i 's which sum to at least T . By our choice of p_1 being sufficiently small, we know that $\frac{1-p_1}{p_1} = \frac{1}{p_1} - 1$ is an integer greater than 2^n . Therefore, the $S(k, t)$ are the unique integers in the base- $(\frac{1}{p_1} - 1)$ representation of $P[V_n \geq T]/p_1^n$, and can be found efficiently. So given p^* we can compute the total number of subsets of the a_i 's that sum up to at least T , thereby solving the given instance of #-SUBSET SUM.

It remains to argue that we can compute p^* using oracle access to SoAP. We do binary search on p while maintaining all other parameters of the SoAP instance fixed, as described above. In every step of the binary search, we solve the corresponding SoAP instance, determining if the optimal price is 1 or $T + 1$ and respectively increasing or decreasing the value of p for the next step, until we have pinned down p^* exactly. To argue that this takes polynomial time we notice that:

$$p^* = \frac{1/(T + 1) - P[V_n \geq T]}{1 - P[V_n \geq T]}.$$

We also notice that $P[V_n \geq T]$ is a rational number that can be specified with a polynomial number of bits.³ So p^* has polynomial accuracy and we need polynomially many calls to SoAP to determine it exactly. \square

4 Complexity of Mildly Irrational Valuations

Issues of numerical precision may arise when analyzing value distributions which are implicitly described. Even very mild irrationality, such as the support of the distribution containing square roots of integers, can cause the resulting pricing problem to be computationally intricate. In particular, optimization may require deciding between two mechanisms whose expected utility differs only by an exponentially small amount. In this section, we present an example of how we can reduce a numerical problem whose status even in NP remains unknown to the pricing problem for a unit-demand buyer with mildly irrational valuations.

³ In particular, each number of the form $p_1^i(1 - p_1)^{n-i}$ has polynomial bit-length.

Definition 2 (The Square Root Sum Problem). *Given positive integers $\alpha_1 \leq \alpha_2 \leq \dots \leq \alpha_n$ and K , the SQRT-SUM problem is to determine whether or not $\sum_{i=1}^n \sqrt{\alpha_i} > K$.*

While known to be in PSPACE, it remains an important open problem whether the square root sum problem is solvable in NP, let alone whether it is in P. [12, 14]

Remark 1. Checking whether $\sum_i \sqrt{a_i} = K$ for positive integers a_i , $i = 1, \dots, n$, and K can be done in polynomial time [14]. So the square root sum problem draws its computational difficulty from instances where equality between $\sum_i \sqrt{a_i}$ and K does not hold and we need to decide whether $\sum_i \sqrt{a_i}$ is $>$ or $<$ than K . In the hardness proofs of Theorems 2 and 3 we will implicitly assume that the given instance of the square root sum problem satisfies $\sum_i \sqrt{a_i} \neq K$. Given such instance we will construct a unit-demand pricing instance whose solution answers the question of whether $\sum_i \sqrt{a_i}$ is $>$ or $<$ than K .

Remark 2. The important computational difference between the square root of an integer and the sum of square roots of multiple integers is that the i -th bit of the former can be computed in time polynomial in i and the number's description complexity, while the same is not known to be true for the latter.

Theorem 2. *The unit-demand pricing problem is SQRT-SUM-hard when the item values are independent of support two with rational probabilities and each possible item value is the square root of an integer.⁴*

Proof. We will reduce SQRT-SUM to the pricing problem for a single unit-demand buyer whose values for the items are distributed independently, take one of two possible values with rational probabilities, and each of these possible values is the square root of an integer.

Given an input $\alpha_1 \leq \alpha_2 \leq \dots \leq \alpha_n$ and K to the SQRT-SUM problem, we construct an input to the unit-demand pricing problem with $n + 1$ items. For $i = 1, \dots, n$, item i has value $\sqrt{\alpha_i}$ with probability $1/i$, and value 0 with probability $1 - 1/i$. Finally, item $n + 1$ has value $T/2$ with probability $1/2 + \epsilon$ and value T with probability $1/2 - \epsilon$, where:

$$\epsilon \triangleq \frac{K}{4n \max(K, \alpha_n)} \leq \frac{1}{2}; \quad T \triangleq \frac{(1/2 + \epsilon)K}{n\epsilon}.$$

Notice that $T/2 > \frac{K}{4n\epsilon} = \max(K, \alpha_n) \geq \alpha_n \geq \sqrt{\alpha_n}$.

We now claim that the optimal expected revenue for the unit-demand pricing instance we defined is the maximum of $T/2$ and

$$(1/2 - \epsilon)T + \frac{1/2 + \epsilon}{n} (\sqrt{\alpha_1} + \dots + \sqrt{\alpha_n}).$$

⁴ The item values are mildly irrational since the i -th bit of the square root of an integer can be computed exactly in time polynomial in i and the description complexity of the integer.

Indeed, it is clearly possible to achieve revenue $T/2$ by pricing item $n + 1$ at $T/2$ and all other items at a price greater than $\sqrt{\alpha_n}$. Since $T/2 > \sqrt{\alpha_n}$, if item $n + 1$ is priced less than or equal to $T/2$, the revenue cannot be higher than $T/2$.

Now what if item $n + 1$ were priced at a price higher than $T/2$? Suppose, e.g., that we price item $n + 1$ at T and all other items i at $\sqrt{\alpha_i}$. Then the expected revenue we would get is⁵

$$(1/2 - \epsilon)T + (1/2 + \epsilon) \left(\frac{1}{n}\sqrt{\alpha_n} + \frac{n-1}{n} \cdot \frac{1}{n-1}\sqrt{\alpha_{n-1}} + \dots + \frac{1}{n}\sqrt{\alpha_1} \right) \quad (1)$$

We claim that this is the best revenue we could possibly achieve if item $n + 1$ is priced at a price higher than $T/2$. Indeed, it is easy to see that the maximum of the values of items $1, \dots, n$ is independent of the value of item $n + 1$, it has expectation $\frac{1}{n} \sum_i \sqrt{\alpha_i}$ and, because $T/2 > \sqrt{\alpha_n}$, it is smaller than T with probability 1. So consider any pricing where the price of item $n + 1$ is larger than $T/2$. In the event that the value of item $n + 1$ is T (which happens with probability exactly $1/2 - \epsilon$) the best revenue that the pricing could possibly get is at most T , while in the event that the value of item $n + 1$ is $T/2$ (which happens with probability exactly $1/2 + \epsilon$) the revenue cannot exceed the maximum of the values of items $1, \dots, n$ which has expectation $\frac{1}{n} \sum_i \sqrt{\alpha_i}$ even after conditioning on the value of item $n + 1$ as it is independent from the value of item $n + 1$.

Observe that (1) is higher than $T/2$ if and only if

$$\epsilon T < \frac{(1/2 + \epsilon)}{n} (\sqrt{\alpha_1} + \dots + \sqrt{\alpha_n}),$$

which occurs precisely when $K < \sqrt{\alpha_1} + \dots + \sqrt{\alpha_n}$. □

5 Complexity of Mildly Irrational Probabilities

The reduction of the previous section used distributions that were supported on irrational values. A possible critique of this in a discrete setting is that it may be unnatural for an individual to hold irrational values for an item. Contrastingly, it seems more natural to allow for a person’s values to be rational but to depend on certain mildly irrational probabilities.

Perhaps the simplest form of an irrational probability is one for which we can efficiently compute arbitrary bits of its binary expansion correctly.⁶ Notice that using a fair coin to sample exactly such probability, e.g. $\sqrt{1/3}$, is no more work than sampling exactly a rational probability, e.g. $1/3$: Imagine an infinite sequence of coin tosses. We reveal a prefix of that sequence until, viewed as a binary number, we can certify that the sequence lies above or below the target probability written in binary; if above, we output 1, otherwise we output 0.

⁵ Suppose that ties are broken in favor of the most expensive item.

⁶ This property is satisfied, for example, by a probability of the form \sqrt{r} , where r is a rational number; but, as remarked in section 4, it is unknown whether it is satisfied by a probability of the form $\sum_i \sqrt{r_i}$, for rational r_i ’s.

We now consider unit-demand pricing instances as in the previous section, except where the values are integral and the probabilities are irrational. As in the previous section, we will give a SQRT-SUM-hardness reduction.

Theorem 3. *The unit-demand pricing problem is SQRT-SUM-hard when the item values are independent of support two, have probabilities for which the i^{th} bit of their binary expansions can be computed in time polynomial in i , and each possible item value is integral.*

Proof. Let $a_1 \leq \dots \leq a_n$ and K be an instance of the SQRT-SUM problem. Also let X be a large integer with $X > \max\{3K/n, a_n\}$. We define $a_{n+1} = X^2$ maintaining the monotonicity of the sequence a_i since $X > a_n$.

We reduce the given SQRT-SUM instance to an instance of the unit-demand pricing problem with $n + 1$ items. For $i = 1, \dots, n$, item i has value i with probability $p_i = 1 - \sqrt{a_i/a_{i+1}}$, and value 0 with probability $\sqrt{a_i/a_{i+1}}$. Finally, item $n + 1$ has value $T/2$ with probability $3/4$ and value T with probability $1/4$, where:

$$T \triangleq 3 \left(n - \frac{K}{X} \right).$$

Notice that by the choice of $X > 3K/n$ we have that $T/2 > n$, the highest possible value of any other item. Also, since the sequence of a_i 's is non-decreasing, all probabilities p_i are well defined.

As in the proof of Theorem 2, we can argue that the optimal pricing either prices item $n + 1$ at $T/2$ and the other items at infinity (call this ‘‘Scheme 1’’), or prices all items at their high value (call this ‘‘Scheme 2’’). In the former case the revenue is $T/2$. In the latter case the bidder will choose to buy the largest item he values high, i.e. will choose item $n + 1$ if he values it high, otherwise item n if he values it high, and so on.⁷ Therefore, Scheme 1 beats Scheme 2 if and only if:

$$\frac{T}{2} > \frac{T}{4} + \frac{3}{4} (p_n n + p_{n-1} (1 - p_n) (n - 1) + \dots + p_1 \prod_{i=2}^n (1 - p_i)),$$

which becomes, after substituting for the p_i 's:

$$\frac{T}{2} > \frac{T}{4} + \frac{3}{4} \sum_{i=1}^n \left(i \left(\sqrt{\frac{a_{i+1}}{a_{n+1}}} - \sqrt{\frac{a_i}{a_{n+1}}} \right) \right).$$

Simplifying and using the fact that $\sqrt{a_{n+1}} = X$, our condition becomes

$$\frac{T}{2} > \frac{T}{4} + \frac{3}{4} \left(n - \frac{\sum_{i=1}^n \sqrt{a_i}}{X} \right).$$

⁷ As in the proof of Theorem 2 we assume that ties are broken in favor of the most expensive item.

This occurs precisely when:

$$\sum_{i=1}^n \sqrt{a_i} > X(n - T/3) = K.$$

Therefore, Scheme 1 is strictly better than Scheme 2 precisely when $\sum_{i=1}^n \sqrt{a_i} > K$, concluding our reduction from the SQRT-SUM problem. \square

6 Future Work

Studying the complexity of optimal pricing in a Bayesian context is an important question, both theoretically and practically. However, to have a robust complexity model, great care must be taken in specifying the input distributions. Indeed, as shown in this paper, implicit distributions can easily embed hard problems into the distribution's parameters, and therefore any complexity theoretic model of pricing must take into account the complexity of the distributions themselves, and not just the length of a minimal specification.

A setting that avoids the computational barriers raised in this paper is that of several items, each distributed independently on some finite size support, with all values and probabilities rational and explicitly given. This problem is not yet resolved for either unit-demand or additive bidders. Moreover, while our paper has focused only on discrete distributions, issues of distributional specification are perhaps even more vital if one wishes to model the complexity of pricing with continuous distributions. It is of interest to propose a robust computational framework for studying the pricing problem with continuous distributions.

Finally, our results apply to computing the optimal deterministic mechanism, which in the case of a single buyer is tantamount to finding an optimal pricing scheme. It is an important open question to determine the complexity of the optimal mechanism design problem when randomized mechanisms are also allowed.

References

- [1] Alaei, S.: Bayesian Combinatorial Auctions: Expanding Single Buyer Mechanisms to Many Buyers. In: The 52nd Annual IEEE Symposium on Foundations of Computer Science, FOCS (2011)
- [2] Alaei, S., Fu, H., Haghpanah, N., Hartline, J., Malekian, A.: Bayesian Optimal Auctions via Multi- to Single-agent Reduction. In: The 13th ACM Conference on Electronic Commerce (EC) (2012)
- [3] Armstrong, M.: Price Discrimination by a Many-Product Firm. *Review of Economic Studies* 66(1), 151–168 (1999)
- [4] Bhattacharya, S., Goel, G., Gollapudi, S., Munagala, K.: Budget Constrained Auctions with Heterogeneous Items. In: The 42nd ACM Symposium on Theory of Computing, STOC (2010)

- [5] Briest, P.: Uniform Budgets and the Envy-Free Pricing Problem. In: Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfssdóttir, A., Walukiewicz, I. (eds.) ICALP 2008, Part I. LNCS, vol. 5125, pp. 808–819. Springer, Heidelberg (2008)
- [6] Cai, Y., Daskalakis, C.: Extreme-Value Theorems for Optimal Multidimensional Pricing. In: The 52nd Annual IEEE Symposium on Foundations of Computer Science, FOCS (2011)
- [7] Cai, Y., Daskalakis, C., Matthew Weinberg, S.: An Algorithmic Characterization of Multi-Dimensional Mechanisms. In: The 44th Annual ACM Symposium on Theory of Computing, STOC (2012)
- [8] Cai, Y., Daskalakis, C., Matthew Weinberg, S.: Optimal Multi-Dimensional Mechanism Design: Reducing Revenue to Welfare Maximization. In: The 53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS (2012)
- [9] Chawla, S., Hartline, J.D., Kleinberg, R.D.: Algorithmic Pricing via Virtual Valuations. In: The 8th ACM Conference on Electronic Commerce, EC (2007)
- [10] Chawla, S., Hartline, J.D., Malec, D.L., Sivan, B.: Multi-Parameter Mechanism Design and Sequential Posted Pricing. In: The 42nd ACM Symposium on Theory of Computing, STOC (2010)
- [11] Daskalakis, C., Matthew Weinberg, S.: Symmetries and Optimal Multi-Dimensional Mechanism Design. In: The 13th ACM Conference on Electronic Commerce, EC (2012)
- [12] Etessami, K., Yannakakis, M.: On the Complexity of Nash Equilibria and Other Fixed Points. *SIAM J. Comput.* 39(6), 2531–2597 (2010)
- [13] Fang, H., Norman, P.: To bundle or not to bundle. *RAND Journal of Economics* 37(4), 946–963 (2006)
- [14] Garey, M.R., Graham, R.L., Johnson, D.S.: Some NP-Complete Geometric Problems. In: The 8th Annual ACM Symposium on Theory of Computing, STOC (1976)
- [15] Hart, S., Nisan, N.: Approximate revenue maximization with multiple items. In: The 13th ACM Conference on Electronic Commerce, EC (2012)
- [16] Lev, O.: A Two-Dimensional Problem of Revenue Maximization. Discussion Paper Series dp542, The Center for the Study of Rationality, Hebrew University, Jerusalem (April 2010)
- [17] Manelli, A.M., Vincent, D.R.: Multidimensional Mechanism Design: Revenue Maximization and the Multiple-Good Monopoly. *Journal of Economic Theory* 137(1), 153–185 (2007)
- [18] Manelli, A.M., Vincent, D.R.: Bundling as an optimal selling mechanism for a multiple-good monopolist. *Journal of Economic Theory* 127(1), 1–35 (2006)
- [19] Milgrom, P.: *Putting Auction Theory to Work*. Cambridge University Press (2004)
- [20] Myerson, R.B.: Optimal Auction Design. *Mathematics of Operations Research* 6(1), 58–73 (1981)
- [21] Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V.: *Algorithmic Game Theory*. Cambridge University Press, New York (2007)
- [22] Papadimitriou, C.H., Pierrakos, G.: On optimal single-item auctions. In: The 43rd Annual ACM Symposium on Theory of Computing, STOC (2011)
- [23] Sipser, M.: *Introduction to the theory of computation*, 2nd edn. PWS Pub., Boston (2006)