

Algorithmic Pricing via Virtual Valuations

Shuchi Chawla
Computer Science Dept.
University of Wisconsin
Madison, WI, 53706
shuchi@cs.wisc.edu

Jason D. Hartline
Microsoft Research
Silicon Valley Campus
Mountain View, CA, 94043
hartline@microsoft.com

Robert Kleinberg*
Dept. of Computer Science
Cornell University
Ithaca, NY 14853
rdk@cs.cornell.edu

ABSTRACT

Algorithmic pricing is the computational problem that sellers (e.g., in supermarkets) face when trying to set prices for their items to maximize their profit in the presence of a known demand. Guruswami et al. [9] propose this problem and give logarithmic approximations (in the number of consumers) for the unit-demand and single-parameter cases where there is a specific set of consumers and their valuations for bundles are known precisely. Subsequently several versions of the problem have been shown to have poly-logarithmic inapproximability. This problem has direct ties to the important open question of better understanding the Bayesian optimal mechanism in multi-parameter agent settings; however, for this purpose approximation factors logarithmic in the number of agents are inadequate. It is therefore of vital interest to consider special cases where constant approximations are possible.

We consider the unit-demand variant of this pricing problem. Here a consumer has a valuation for each different item and their value for a set of items is simply the maximum value they have for any item in the set. Instead of considering a set of consumers with precisely known preferences, like the prior algorithmic pricing literature, we assume that the preferences of the consumers are drawn from a distribution. This is the standard assumption in economics; furthermore, the setting of a specific set of customers with specific preferences, which is employed in all of the prior work in algorithmic pricing, is a special case of this general Bayesian pricing problem, where there is a discrete Bayesian distribution for preferences specified by picking one consumer uniformly from the given set of consumers. Notice that the distribution over the valuations for the individual items that this generates is obviously correlated. Our work complements these existing works by considering the case where the consumer's valuations for the different items are independent random variables. Our main

*Research supported by an NSF Mathematical Sciences Postdoctoral Research Fellowship. Portions of this work were completed while the author was an NSF postdoctoral fellow at U.C. Berkeley.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'07, June 13–16, 2007, San Diego, California, USA.

Copyright 2007 ACM 978-1-59593-653-0/07/0006 ...\$5.00.

result is a constant approximation algorithm for this problem that makes use of an interesting connection between this problem and the concept of *virtual valuations* from the single-parameter Bayesian optimal mechanism design literature.

Categories and Subject Descriptors

F.0 [Theory of Computation]: General

General Terms

Algorithms, Economics, Theory

Keywords

Pricing, Approximation algorithms, Virtual valuations

1. INTRODUCTION

It is vital to the study of resource allocation in distributed settings such as the Internet that inherent economic issues be addressed. Recently the area of *algorithmic pricing* [1, 4, 5, 9, 10] has emerged as a setting for studying optimization in resource allocation in the presence of a natural fairness constraint: there is a uniform pricing rule under which the consumers are allowed to choose the allocation they most desire. This area has important connections to *algorithmic mechanism design* [2, 3, 9] in addition to obvious applications in traditional market settings such as pricing items in a supermarket.

A *pricing* can be thought of as a menu listing the prices for all possible allocations to a consumer. Given a pricing, a consumer's *preference* indicates a most desired allocation. The algorithmic pricing problem, then, is to take an instance given by a class of allowable pricings and a set of consumers, and compute the pricing maximizing (or approximately maximizing) a specific objective. An *item-pricing* is one where each individual item is assigned a price, and the price of any bundle is the sum of the prices of the items in the bundle. In this arena, the objective of maximizing the profit of the seller presents significant challenges, even when there are no supply constraints. Indeed when the items are pure complements, i.e., consumers are single-minded and combinatorial; and when the items are pure substitutes, i.e., consumers have unit demands, recent works show hardness results for item-pricing that essentially match the poor performance of trivial algorithms [7, 5]. This motivates the search for relevant special cases where algorithmic theory gives an improved understanding of pricing.

Algorithmic pricing and algorithmic mechanism design have important connections. Indeed, for unlimited supply profit

maximization, Balcan et al. [3] give a general reduction from truthful mechanism design to algorithmic pricing. These results are important, in particular, as they address the challenging problem of optimal mechanism design in multi-parameter settings (e.g., general combinatorial auctions and multi-item unit-demand auctions). On the other hand, in single parameter Bayesian settings (e.g., a single-item auction) the optimal auction is well-understood—the well known result of Myerson [12] gives a closed form characterization for the optimal auction. In this paper we bring these results full-circle by showing that techniques from Myerson’s optimal auction are useful for the unit-demand pricing problem. In fact, by designing a pricing to mimic the reserve prices of an optimal auction, we are able to approximate the profit of the optimal auction (and thus, also, the profit of the optimal pricing). Of course, a key difference between single-item auctions and unit-demand pricing is that in an auction the different bids compete against each other, while in a pricing problem there is no competition. Accordingly, our pricing algorithm “simulates competition” by setting higher reserve values.

Formally, the algorithmic mechanism design and algorithmic pricing problems are defined as follows:

DEFINITION 1. (Bayesian Single-item Auction Problem (BSAP))

Given,

- a single item for sale,
- n consumers, and
- distribution \mathbf{F} from which consumer valuations are drawn.

Goal: design seller optimal auction for \mathbf{F} .

DEFINITION 2. (Bayesian Unit-demand Pricing Problem (BUPP))

Given,

- a single unit-demand consumer,
- n items for sale, and
- distribution \mathbf{F} from which the consumer’s valuations for each item are drawn.

Goal: compute seller optimal item-pricing for \mathbf{F} .

In the special case where \mathbf{F} is the product distribution $F_1 \times \dots \times F_n$, the *Bayesian single-item auction problem* was solved by Myerson in his seminal paper on mechanism design [12]. His solution is based on determining the allocation of the item for sale using the consumers’ *virtual valuations* (see Section 2), instead of their actual valuations. We consider product distributions for the *Bayesian unit-demand pricing problem* and show that:

- The optimal revenue of a single-item auction is an upper bound on the revenue of the optimal unit-demand pricing.
- The optimal unit-demand pricing that uses a single *virtual price*¹ for all items obtains a constant fraction of the revenue of the optimal auction.
- If all the distributions satisfy the monotone hazard rate condition (defined in Section 2), a nearly-optimal virtual price can be computed in polynomial time.

We first demonstrate the connection between BSAP and BUPP in the context of identically distributed valuations,

¹The virtual price is a price in virtual valuation space instead of valuation space.

i.e. when $F_i = F_j$ for all $i \neq j$. In this i.i.d. case, our algorithm outputs the same price (not just the same virtual price) for all the items. Note that it is easy to optimize revenue over the space of all pricings that price each item at the same value—just consider the distribution of the maximum valuation ($\max_i v_i$) and solve this problem as a single-consumer single-item revenue maximization problem. One might expect that in the i.i.d. case this optimal single price is in fact the overall optimal pricing. However, a simple example shows that this is not true—consider two items, each with a value independently equal to 1 with probability 2/3 and 2 with probability 1/3; then a simple calculation shows that the pricings (1, 2) and (2, 1) are optimal with respect to revenue² and the pricings (1, 1) and (2, 2) are strictly sub-optimal. In Section 3.1 we prove that in the i.i.d. case, the revenue of the optimal single-price solution is a 3.47-approximation to the optimal revenue of the single-item auction.

We extend this result to the case of general product distributions in Section 3.2, proving that the revenue of the optimal single virtual-price solution is a 4-approximation to the optimal revenue of the corresponding BSAP.

In Section 4 we consider distributions that do not satisfy the monotone hazard rate condition (see Definition 3 in Section 2), also called the *non-regular case* in the economics literature. Myerson’s solution to BSAP in this case uses a smoothed or “ironed” version of virtual valuations. We show that the same fix can be applied to the pricing problem and again the revenue of the optimal single *ironed-virtual-price* solution is a 4-approximation to the optimal revenue of the single-item auction.

We consider the question of computing the optimal virtual price in Section 5. For general discrete distributions that satisfy the MHR condition, we obtain a polynomial time constant-factor approximation algorithm. For general continuous distributions, we consider a computational model in which we have oracle access to the cumulative distribution function and probability density functions the distributions (see Section 2 for more detail), and again obtain a polynomial time approximation algorithm. We leave open the problem of designing a polynomial time approximation algorithm for the non-regular case. The challenge in this case is to come up with a polynomial time algorithm for computing ironed-virtual-valuations.

2. NOTATION AND PRELIMINARIES

The Bayesian Unit-demand Pricing Problem (BUPP)

The input to BUPP is a distribution over n -tuples of valuations. We use $\mathbf{v} = (v_1, \dots, v_n)$ to denote the valuation vector. The value v_i is drawn independently from the distribution F_i over the range $[\ell_i, h_i]$. Following standard notation, we use \mathbf{v}_{-i} to denote all the valuations except the i th one. $\mathbf{F} = F_1 \times \dots \times F_n$ denotes the product distribution from which \mathbf{v} is drawn, and $f_i(v_i)$ denotes the probability density of valuation v_i . Our goal is to determine a price vector $\mathbf{p} = (p_1, \dots, p_n)$ such that the expected revenue $\mathcal{R}^{\mathbf{p}}$, as

²We assume that whenever the consumer faces a tie, i.e. two or more items bring equal utility to her, the seller has the ability to break the tie in favor of any of the items (in particular, the most expensive item). The seller could enforce this by giving a negligibly small discount to the consumer for the most expensive item.

defined below, is maximized.

$$\mathcal{R}^{\mathbf{P}} = \sum_i p_i \Pr_{\mathbf{v} \sim \mathbf{F}} \left[(v_i - p_i) \geq \max_{j \leq n} (v_j - p_j) \right]$$

The Monotone Hazard Rate condition

In much of the paper we will assume that the distributions F_i satisfy the *monotone hazard rate* (MHR) condition defined below. This is a standard assumption used in economics. In the single-consumer, single-item case, this condition essentially implies that the revenue as a function of price has a unique maximum.

DEFINITION 3 (MONOTONE HAZARD RATE). *A distribution F with density f is said to satisfy the monotone hazard rate (MHR) condition if $\frac{1-F(v)}{f(v)}$ is monotonically non-increasing for all v .*

In Section 4 we extend our results to distributions that do not satisfy the MHR condition. This is also called the non-regular case in the literature.

The computational model

We consider two different computational models for BUPP:

- **(Discrete explicit)** In this model, each of the distributions F_i are discrete distributions with small support. These distributions are specified explicitly, and our algorithm is required to run in time polynomial in the number of items n , and the size of the largest support.
- **(Continuous with oracles)** In this model, the distributions F_i are continuous with known supports $[\ell_i, h_i]$. The algorithm is provided the following oracles: an oracle to determine $F_i(v)$ given a value v and an index i , an oracle to determine the density $f_i(v)$ given a value v and an index i , and an oracle to sample from the product distribution \mathbf{F} . The algorithm is required to run in time polynomial in the number of items n and the range $(\max_i h_i)/(\min_i \ell_i)$.

Myerson’s optimal mechanism

The Bayesian Single-item Auction Problem (BSAP) is described as follows: there is a single item for sale and n bidders with values given by the vector \mathbf{v} ; each bidder’s value v_i is drawn independently from a distribution F_i ; the goal of the mechanism designer is to design a *truthful* auction so as to maximize the revenue obtained by the seller from the sale of the item.

In one of the seminal works of Bayesian mechanism design, Myerson developed a mechanism for this problem that obtains the maximum revenue for the seller over the class of all truthful mechanisms [12]. Myerson’s mechanism (denoted \mathcal{M} hereafter) first computes a function of the valuation of each bidder, known as the virtual valuation³:

$$\phi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)} \quad (1)$$

³The virtual valuations computed by Myerson’s mechanism are actually functions of the bids reported by the participants, and not their true values. But given that the mechanism is truthful and these bids are equal to the true values, for simplicity we express the virtual valuations as functions of the true values.

The virtual valuation of a bidder essentially denotes the marginal revenue obtained by allocating the item to this bidder. Myerson’s mechanism offers the item to the bidder with the highest virtual valuation, at a price equal to the virtual-valuation-inverse of the second highest virtual valuation (i.e. the value at which the bidder’s virtual valuation equals the second highest one, or the minimum bid the bidder needs to make to win the item).

In our subsequent discussions it will sometimes be useful to consider the following alternate description of Myerson’s mechanism. The mechanism offers each bidder a take-it-or-leave-it price. The price offered to bidder i is equal to $\phi_i^{-1}(\nu_i)$, where $\nu_i = \max_{j \neq i} \phi_j(v_j)$. Only one bidder (the one with the highest virtual valuation) accepts.

We use $\widehat{\mathcal{R}}^{\mathcal{A}}$ to denote the expected revenue of a truthful mechanism \mathcal{A} for BSAP. $\widehat{\mathcal{R}}^{\mathcal{M}}$ denotes the expected revenue of Myerson’s mechanism \mathcal{M} . In the following theorem and lemma, we assume that the distributions F_i satisfy the monotone hazard rate condition.

THEOREM 1 (MYERSON [12]). $\widehat{\mathcal{R}}^{\mathcal{M}} \geq \widehat{\mathcal{R}}^{\mathcal{A}}$ for all truthful mechanisms \mathcal{A} .

We first note that $\widehat{\mathcal{R}}^{\mathcal{M}}$ is at least as large as the expected revenue of any price vector $\mathcal{R}^{\mathbf{P}}$ for the BUPP.

LEMMA 2. For any price vector \mathbf{p} , $\widehat{\mathcal{R}}^{\mathcal{M}} \geq \mathcal{R}^{\mathbf{P}}$.

PROOF. For a given pricing \mathbf{p} , consider the following mechanism $\mathcal{A}^{\mathbf{P}}$: given a valuation vector \mathbf{v} , we allocate the item to the bidder i with $v_i \geq p_i$ that maximizes $v_i - p_i$. Prices are determined by the standard “threshold payment” rule: the winning bidder, i , pays the minimum bid value which would still make i the winner. $\mathcal{A}^{\mathbf{P}}$ is truthful because it gives a monotone allocation procedure: if a winning bidder unilaterally increases her bid, she still wins. Therefore, $\widehat{\mathcal{R}}^{\mathcal{A}^{\mathbf{P}}} \leq \widehat{\mathcal{R}}^{\mathcal{M}}$.

Now consider any valuation vector \mathbf{v} and suppose that $\mathcal{A}^{\mathbf{P}}$ allocates the item to bidder i . Then the minimum bid at which this bidder is allocated the item is $p_i + \max_{j \neq i} (v_j - p_j, 0)$, which is at least p_i . Therefore, the revenue of $\mathcal{A}^{\mathbf{P}}$ when the valuation vector is \mathbf{v} is at least p_i . However, the revenue of the pricing \mathbf{p} when the valuation vector is \mathbf{v} is exactly p_i . Therefore, $\widehat{\mathcal{R}}^{\mathcal{A}^{\mathbf{P}}} \geq \mathcal{R}^{\mathbf{P}}$. Combining the two inequalities proves the lemma. \square

Two observations lead us to connect Myerson’s setting to BUPP. First, as the number of bidders gets large (especially in the case of identically distributed valuations), the price offered to a bidder in Myerson’s mechanism becomes tightly concentrated around a single value (the expectation of the virtual-value-inverse of the maximum over virtual valuations of other bidders). This value is a reasonable candidate for the price of item i in the pricing problem, and is indeed roughly what we use (with some modifications to allow for an easier analysis).

This approach does not immediately work however. The problem is that in Myerson’s mechanism, by allowing the price to be an appropriate function of other bidders’ values, we ensure that only one bidder accepts the offered price. In the BUPP, with some probability, more than one of the values is above the corresponding offered price, and the consumer gets to pick which item to buy (i.e. there is a lack of competition). In this case, the price earned by our solution may be much worse than the price earned by Myerson’s mechanism

for the same valuation vector. Our second observation is that this situation happens with a low probability. Furthermore, we are able to use the monotone hazard rate condition to charge the revenue earned by Myerson's mechanism against the revenue earned by our pricing in the case of such an event.

3. APPROXIMATING PRICING IN THE REGULAR CASE

In this section we demonstrate that the unit-demand optimal pricing that uses a single virtual price for all items obtains a constant fraction of the revenue of the optimal single-item auction. We begin with the i.i.d. case and then extend our results to general product distributions. Throughout this section we assume that all distributions satisfy the MHR condition (Definition 3).

3.1 The i.i.d. case

Let F denote the distribution from which each valuation v_i is drawn. Let f denote the density function for this distribution and ϕ denote the virtual valuation function.

We consider the following pricing in this case: for all i , set $p_i = p = \max\{F^{-1}(1 - 1/n), \phi^{-1}(0)\}$. Let $q = 1 - F(p)$; note that $q \leq 1/n$.

Given this pricing, the probability that a sale is made is exactly $\Pr[\exists i : v_i \geq p] = 1 - (1 - q)^n$. Using Taylor's expansion and $q \leq 1/n$ we can simplify this expression as follows: $(1 - q)^n < 1 - qn + \frac{1}{2}q^2n^2 \leq 1 - \frac{1}{2}qn$. Noting that the expected revenue of \mathbf{p} is $\mathcal{R}^{\mathbf{p}} = p(1 - (1 - q)^n)$, we get the following.

LEMMA 3. $\mathcal{R}^{\mathbf{p}} \geq \frac{1}{2}pqn$.

We now analyze the expected revenue of Myerson's mechanism when each value is distributed according to F . We consider three distinct events. Let Ω_0 be the event that the value vector \mathbf{v} drawn from \mathbf{F} has $v_i < p$ for each i . Likewise, Ω_1 is the event that $v_i \geq p$ for exactly one index i , and $\Omega_{>1}$ is the event that $v_i \geq p$ for at least two indices i . Let ρ_0, ρ_1 and $\rho_{>1}$ denote the contribution of these three events respectively to the expected revenue obtained by Myerson's mechanism.

$$\widehat{\mathcal{R}}^{\mathcal{M}} = \rho_0 + \rho_1 + \rho_{>1}.$$

We now bound these three terms individually. ρ_1 is the easiest to bound:

LEMMA 4. $\rho_1 < \mathcal{R}^{\mathbf{p}}$.

PROOF. Suppose that the event Ω_1 happens and let \mathbf{v} denote the valuation vector drawn from \mathbf{F} . Note that the second highest virtual value among $\{\phi(v_i)\}$ is strictly less than $\phi(p)$. Therefore, the revenue generated from this vector is strictly less than p . On the other hand, this vector contributes p to the expected revenue of the pricing \mathbf{p} . Integrating over all such valuation vectors belonging to Ω_1 we get the result. \square

Next we analyze ρ_0 . We use the fact that the event Ω_0 happens with low probability.

LEMMA 5. $\rho_0 < \frac{2}{e}\mathcal{R}^{\mathbf{p}}$.

PROOF. We consider two cases. First, when $\phi(p) = 0$, then in the event Ω_0 , Myerson's mechanism does not allocate the item to any bidder (because all of them have negative virtual valuations). Therefore, $\rho_0 = 0$ and the lemma holds.

Next consider the case when $\phi(p) > 0$ (and $q = 1/n$). Then the contribution to the revenue $\widehat{\mathcal{R}}^{\mathcal{M}}$ by any valuation vector in the event Ω_0 is strictly less than p . On the other hand, $\Pr[\Omega_0] = (1 - q)^n \leq 1/e$. So we get $\rho_0 < p/e$. This along with Lemma 3 and $q = 1/n$ implies the lemma. \square

Finally, we analyze the contribution of the event $\Omega_{>1}$. Note that even though $\Omega_{>1}$ may be a very low probability event, its contribution to the expected revenue may be quite high, because \mathcal{M} may charge prices much higher than p in this event. We handle this issue by using the fact that the distribution F satisfies the MHR condition (Definition 3). The MHR condition, in particular, implies the following (we give a brief proof for completeness).

FACT 6. Let v be a random variable distributed according to distribution F and density f . Let $\phi(v)$ be defined as in Equation 1. If F satisfies the MHR condition (Definition 3) and $v_1 \geq v_2 \geq \phi^{-1}(0)$, then

$$v_1(1 - F(v_1)) \leq v_2(1 - F(v_2)).$$

PROOF. Let $G(v) = v(1 - F(v))$. Then $G'(v) = (1 - F(v)) - vf(v) = -f(v)\phi(v)$. The MHR condition implies that $\phi(\cdot)$ is an increasing function. Therefore, for $v > \phi^{-1}(0)$, $\phi(v) > 0$, implying that $G'(v) < 0$. \square

LEMMA 7. $\rho_{>1} \leq 2\mathcal{R}^{\mathbf{p}}$.

PROOF. Let $R_i(\mathbf{v}_{-i})$ denote the contribution of bidder i to the revenue $\rho_{>1}$ conditioned on the other values being given by \mathbf{v}_{-i} . In other words, fixing all values except v_i according to \mathbf{v}_{-i} , we look at those values v_i at which bidder i gets served and the event $\Omega_{>1}$ happens. The contribution to the revenue of all such values is termed $R_i(\mathbf{v}_{-i})$.

Fixing the values of all players except i , we know that Myerson essentially offers a take-it-or-leave-it price p' to bidder i . If $p' < p$, then regardless of the value v_i , the event $\Omega_{>1}$ does not hold; therefore, $R_i(\mathbf{v}_{-i}) = 0$. On the other hand, if $p' \geq p$, then the contribution is $R_i(\mathbf{v}_{-i}) = p'(1 - F(p'))$. Fact 6 along with $p' \geq p \geq \phi^{-1}(0)$ implies that $R_i(\mathbf{v}_{-i}) = p'(1 - F(p')) \leq p(1 - F(p)) = pq$. Removing the conditioning on \mathbf{v}_{-i} , we get that the contribution of bidder i to $\rho_{>1}$ is $R_i < pq$. The lemma follows by summing over all bidders and applying Lemma 3. \square

Lemmas 4, 5 and 7 together give the following.

THEOREM 8. When all valuations are distributed identically, the pricing \mathbf{p} given above satisfies $\widehat{\mathcal{R}}^{\mathcal{M}} \leq (3 + 2/e)\mathcal{R}^{\mathbf{p}} < 3.74\mathcal{R}^{\mathbf{p}}$.

3.2 The general case

Based on the simple pricing algorithm developed in the preceding section, it is very tempting to try setting a single price in the general case as well, for example by computing the distribution F_{\max} of the random variable $\max_i v_i$ and setting the price p which maximizes $p \cdot (1 - F_{\max}(p))$. However, the following example shows that algorithms which use a single price cannot achieve a constant-factor approximation to the profit of the optimal pricing.

EXAMPLE 1. Suppose the distribution of v_i is given by:

$$\begin{aligned} \Pr[v_i = \frac{n}{i}] &= \frac{1}{n} \\ \Pr[v_i = 1] &= 1 - \frac{1}{n}. \end{aligned}$$

The pricing \mathbf{p} which sets $p_i = n/i$ achieves a profit of $\Omega(\log n)$ because

$$\begin{aligned}\mathcal{R}^{\mathbf{P}} &= \sum_{i=1}^n \binom{n}{i} \Pr \left[v_i = \frac{n}{i} \text{ and } \forall j < i \ v_j = 1 \right] \\ &= \sum_{i=1}^n \binom{n}{i} \cdot \left(\frac{1}{n}\right) \cdot \left(1 - \frac{1}{n}\right)^{i-1} \\ &> \sum_{i=1}^n \binom{n}{i} \cdot \left(\frac{1}{n}\right) \cdot \left(\frac{1}{e}\right) \\ &= H_n/e,\end{aligned}$$

where H_n denotes the n -th harmonic number, $\sum_{i=1}^n \frac{1}{i}$.

On the other hand, if \mathbf{p} is any pricing which sets $p_i = p$ for some fixed value of $p > 1$, then

$$\begin{aligned}\mathcal{R}^{\mathbf{P}} &\leq \sum_{i=1}^n p \cdot \Pr[v_i \geq p] \\ &= \sum_{1 \leq i \leq n/p} p \cdot \left(\frac{1}{n}\right) + \sum_{i > n/p} p \cdot 0 \\ &\leq 1.\end{aligned}$$

This example establishes that no pricing which uses a single price can achieve a $o(\log n)$ -approximation to the profit of the optimal pricing.

Rather than using a single price, our solution to the general case is defined as follows. Let ν be the smallest non-negative number such that

$$\Pr[\exists i : \phi_i(v_i) \geq \nu] \leq \frac{1}{2}.$$

In other words, ν is defined such that the event that at least one virtual value is greater than ν has probability exactly $1/2$. If this virtual value ν turns out to be negative, we redefine it to be zero. We set the price of each item p_i to be $\phi_i^{-1}(\nu)$ and let $q_i = 1 - F_i(p_i)$ for all i . Let \mathbf{p} denote this collection of prices.

It is possible to analyze the expected revenue of this pricing using an extension of the technique which we applied to the i.i.d. case in Section 3.1. This leads to a proof of the bound $\widehat{\mathcal{R}}^{\mathcal{M}} \leq 7.62\mathcal{R}^{\mathbf{P}}$. We omit the details of this approach. Instead, in this section, we use a different technique to prove the stronger bound $\widehat{\mathcal{R}}^{\mathcal{M}} \leq 4\mathcal{R}^{\mathbf{P}}$.

Let \mathcal{M}_ν denote the following truthful mechanism. If all virtual valuations are less than ν , then the item is not sold, and no payments are charged to the bidders. Otherwise, the item is sold to the bidder i with the highest virtual valuation, at a price equal to that bidder's inverse-virtual-valuation of the second-highest virtual valuation, i.e. the minimum bid value that makes i the winner. Myerson's original paper on optimal auction design [12] proves that if an unsold item is worth ν to the seller, then \mathcal{M}_ν is the truthful auction mechanism which maximizes the seller's expected utility. In other words, if we use the notation $y(\mathcal{A})$ to denote the probability that the item is unsold when using a given auction mechanism \mathcal{A} , then for every truthful mechanism \mathcal{A} we have:

$$\text{FACT 9. } \widehat{\mathcal{R}}^{\mathcal{M}_\nu} + \nu \cdot y(\mathcal{M}_\nu) \geq \widehat{\mathcal{R}}^{\mathcal{A}} + \nu \cdot y(\mathcal{A}).$$

Let

$$y = y(\mathcal{M}_\nu), \quad x = 1 - y.$$

Our definition of ν ensures that either $x = y = \frac{1}{2}$ or $\nu = 0$, and in the latter case $y \geq \frac{1}{2}$.

$$\text{LEMMA 10. } \widehat{\mathcal{R}}^{\mathcal{M}_\nu} + \nu y \geq \widehat{\mathcal{R}}^{\mathcal{M}}.$$

PROOF. The first half of the lemma follows by invoking Fact 9 with $\mathcal{A} = \mathcal{M}$ and noting that $y(\mathcal{A}) \geq 0$ for every mechanism \mathcal{A} . \square

$$\text{LEMMA 11. } \mathcal{R}^{\mathbf{P}} \geq \nu x.$$

PROOF. For every i , we have $p_i = \phi_i^{-1}(\nu) \geq \nu$. Thus

$$\mathcal{R}^{\mathbf{P}} \geq \nu \cdot \Pr[\text{the item is sold}] = \nu \cdot x,$$

as when we post price \mathbf{p} the item is sold with the same probability as in \mathcal{M}_ν , i.e., x . \square

In the upcoming proofs, it will be useful to define

$$g = \sum_i p_i q_i.$$

$$\text{LEMMA 12. } \mathcal{R}^{\mathbf{P}} \geq gy.$$

PROOF. The revenue $\mathcal{R}^{\mathbf{P}}$ is bounded below by the summation, over all i , of p_i times the probability that i is the unique index satisfying $v_i \geq p_i$, i.e.

$$\mathcal{R}^{\mathbf{P}} \geq \sum_i p_i \cdot \left(q_i \prod_{j \neq i} (1 - q_j) \right). \quad (2)$$

Recall that

$$y = \Pr[\forall i : v_i \leq p_i] = \prod_i F_i(p_i) = \prod_i (1 - q_i).$$

Combining this with (2) we find that

$$\mathcal{R}^{\mathbf{P}} \geq \sum_i p_i q_i y = gy.$$

\square

Proceeding as in Section 3.1, we now partition the sample space into three events Ω_0 , Ω_1 , and $\Omega_{>1}$ according to whether the number of indices i satisfying $v_i \geq p_i$ is 0, 1, or greater than 1. The expected revenue of \mathcal{M}_ν , which is the integral of a random variable over the entire sample space, can be expressed as the sum of the integrals over the three pieces Ω_0 , Ω_1 , and $\Omega_{>1}$. We denote these three terms by ρ_0 , ρ_1 , and $\rho_{>1}$, respectively.

$$\widehat{\mathcal{R}}^{\mathcal{M}_\nu} = \rho_0 + \rho_1 + \rho_{>1}. \quad (3)$$

From the definition of \mathcal{M}_ν it is obvious that

$$\rho_0 = 0. \quad (4)$$

The following two lemmas bound the remaining two terms $\rho_1, \rho_{>1}$.

$$\text{LEMMA 13. } \rho_1 \leq \mathcal{R}^{\mathbf{P}}.$$

PROOF. Let the random variable R denote the revenue obtained from using the pricing \mathbf{p} , and let Q denote the revenue obtained from using the auction \mathcal{M}_ν . If i is the unique index satisfying $v_i \geq p_i$, then $Q = R = p_i$. Hence

$$\rho_1 = \int_{\Omega_1} Q f(\mathbf{v}) d\mathbf{v} = \int_{\Omega_1} R f(\mathbf{v}) d\mathbf{v} \leq \int_{\mathbf{v}} R f(\mathbf{v}) d\mathbf{v} = \mathcal{R}^{\mathbf{P}}.$$

\square

$$\text{LEMMA 14. } \rho_{>1} \leq g.$$

PROOF. As in Section 3.1, let $R_i(\mathbf{v}_{-i})$ denote the contribution of bidder i to the revenue $\rho_{>1}$ conditioned on the other values being given by \mathbf{v}_{-i} . Fixing the values of all players except i , \mathcal{M} offers a take-it-or-leave-it price $p' = \phi_i^{-1}(\nu')$ to bidder i , where $\nu' = \max_{j \neq i} \phi_j(v_j)$. If $p' < p_i$ (i.e. $\nu' < \nu$), then regardless of the value v_i , the event $\Omega_{>1}$ does not hold; so, $R_i(\mathbf{v}_{-i}) = 0$. On the other hand, if $p' \geq p_i$, then the contribution is $R_i(\mathbf{v}_{-i}) = p'(1 - F_i(p')) \leq p_i(1 - F_i(p_i)) = p_i q_i$, using Fact 6. Removing the conditioning on \mathbf{v}_{-i} , we get that the contribution of bidder i to $\rho_{>1}$ is $R_i < p_i q_i$. Summing over all the bidders and recalling that $g = \sum_i p_i q_i$, we get the result. \square

THEOREM 15. For the pricing \mathbf{p} defined above,

$$\widehat{\mathcal{R}}^{\mathcal{M}} \leq 4\mathcal{R}^{\mathbf{p}}.$$

PROOF. Let $\chi = 1$ if $\nu > 0$ and $\chi = 0$ otherwise. Notice that $\nu = \nu\chi$. We have

$$\begin{aligned} \widehat{\mathcal{R}}^{\mathcal{M}} &\leq \widehat{\mathcal{R}}^{\mathcal{M}\nu} + \nu y \chi && \text{[by Lemma 10]} \\ &= \rho_1 + \rho_{>1} + \nu y \chi && \text{[by (3) and (4)]} \\ &\leq \mathcal{R}^{\mathbf{p}} + g + \nu y \chi && \text{[by Lemmas 13,14]} \\ &\leq \mathcal{R}^{\mathbf{p}} + (\mathcal{R}^{\mathbf{p}}/y) + (\mathcal{R}^{\mathbf{p}}/x)y \chi && \text{[by Lemmas 11,12]} \\ &= \mathcal{R}^{\mathbf{p}} [1 + 1/y + (y/x)\chi]. \end{aligned}$$

If $\nu = 0$ then $\chi = 0$ and $y \geq \frac{1}{2}$, hence $\widehat{\mathcal{R}}^{\mathcal{M}} \leq 3\mathcal{R}^{\mathbf{p}}$. If $\nu > 0$ then $\chi = 1$ and $y = x = \frac{1}{2}$, hence $\widehat{\mathcal{R}}^{\mathcal{M}} \leq 4\mathcal{R}^{\mathbf{p}}$. \square

4. THE NON-REGULAR CASE

In our analysis in Section 3.2, we used the MHR condition to imply that the functions $\phi_i(v_i)$ are non-decreasing. When the MHR condition does not hold, Myerson applies a fix to the problem by smoothing out or “ironing” the virtual valuation function to make it a non-decreasing function of v_i . We now show that by picking a pricing based on ironed virtual valuations instead of the actual virtual valuations, we achieve the same guarantee as in the regular case—the revenue of our pricing is within a factor of 4 of the revenue of Myerson’s mechanism.

We briefly describe this ironing procedure below. The reader is referred to Myerson’s paper [12] and a survey of Bulow and Roberts [6] for more details.

The ironing procedure

The ironed virtual valuation function is defined as follows. Consider a single bidder with value v distributed according to function F . We assume that the density function $f(v)$ is non-zero for all $v \in [\ell, h]$. For $\alpha \in [0, 1]$, let $R(\alpha)$ denote the revenue generated from offering the item to this bidder at price $F^{-1}(\alpha)$:⁴

$$R(\alpha) = F^{-1}(\alpha)(1 - \alpha) = \int_{F^{-1}(\alpha)}^h \phi(t)f(t)dt$$

Let $\bar{R}(\alpha)$ be the least-valued concave function on $[0, 1]$ with $\bar{R}(\alpha) \geq R(\alpha)$ for all α in that range (see Figure 1). Since \bar{R} is concave, it is differentiable everywhere except at finitely many points. Let $\bar{r}(\alpha)$ denote the derivative of \bar{R} wherever defined. The ironed virtual valuation function is defined as below where ever \bar{r} is defined, and is extended to the full

⁴Note that $F^{-1}(\alpha)$ is well-defined because F is a strictly increasing function.

range of v by right continuity.

$$\bar{\phi}(v) = -\bar{r}(F(v))$$

Note that since $\bar{R}(\alpha)$ is concave and $F(v)$ is non-decreasing, $\bar{\phi}(v)$ is a non-decreasing function. Furthermore, observing that $\bar{R}(1) = R(1) = 0$, we get the following:

$$\int_{t=v}^{t=h} \bar{\phi}(t)f(t)dt = - \int_{t=F(v)}^{t=1} \bar{r}(t)dt = \bar{R}(F(v)) \quad (5)$$

For any ν , the inverse ironed virtual valuation $\bar{\phi}^{-1}(\nu)$ is defined to be the infimum over values v with $\bar{\phi}(v) = \nu$.

Approximate pricing

Myerson’s optimal mechanism in the non-regular case proceeds as follows. It first computes the ironed virtual valuations of the values of all bidders. It then allocates the item to the bidder with the highest ironed virtual valuation at a price equal to the inverse of the second highest one.

Our pricing is similarly defined. Let ν be the minimum non-negative number satisfying

$$\Pr[\exists i : \bar{\phi}_i(v_i) \geq \nu] \leq \frac{1}{2}.$$

Let $p_i = \bar{\phi}_i^{-1}(\nu)$. We offer the price p_i for item i . We obtain the following theorem.

THEOREM 16. In the non-regular case, for the pricing \mathbf{p} defined above, $\widehat{\mathcal{R}}^{\mathcal{M}} \leq 4\mathcal{R}^{\mathbf{p}}$.

The proof of this theorem is identical to that of Theorem 15, except that we need analogues of Fact 6 and Fact 9 for ironed virtual valuations. Fact 9 is proven in Myerson’s original paper on optimal auction design [12]. We prove the analogue of Fact 6 in Lemma 18 below. We omit a full proof of Theorem 16 for the sake of brevity, since it is essentially identical to the proof of Theorem 15.

We need the following property:

LEMMA 17. For any ν and $v = \bar{\phi}^{-1}(\nu)$,

$$\bar{R}(F(v)) = R(F(v)) = v(1 - F(v)).$$

PROOF. Note that by definition, \bar{R} is the boundary of the convex hull of the hypograph of R (see Figure 1). Theorem 18.3.1 in Rockafellar [13] implies that every extreme point of this convex hull lies on the boundary of the hypograph of R . That is, at all extreme points $(\alpha, \bar{R}(\alpha))$, we have $R(\alpha) = \bar{R}(\alpha)$. Furthermore, all points with non-zero curvature, as well as end-points of maximal linear segments in the graph of \bar{R} , are extreme points of the convex hull (i.e. they cannot be expressed as convex combinations of two other distinct points in the set).

Now consider some $v = \bar{\phi}^{-1}(\nu)$. If v is the unique value with $\bar{\phi}(v) = \nu$, then \bar{R} has non-zero curvature at $F(v)$ by definition, and so $\bar{R}(F(v)) = R(F(v))$. Otherwise, v is the infimum over all values x with $\bar{\phi}(x) = \nu$, in which case $F(v)$ is the left end-point of a maximal linear subsegment in \bar{R} . Again we have $\bar{R}(F(v)) = R(F(v))$. \square

The following lemma is a direct consequence of Lemma 17 by observing that for $v \geq \bar{\phi}^{-1}(0)$, $\bar{\phi}(v)$ is non-negative, $\bar{r}(F(v))$ is non-positive, and so $\bar{R}(F(v))$ is a non-increasing function of v .

LEMMA 18. Let $\nu_1 \geq \nu_2 \geq 0$, $v_1 = \bar{\phi}^{-1}(\nu_1)$, and $v_2 = \bar{\phi}^{-1}(\nu_2)$. Then $v_1(1 - F(v_1)) \leq v_2(1 - F(v_2))$.

We note that Theorem 16 only gives a characterization of an approximately optimal pricing in the non-regular case, and not a poly-time approximation algorithm. We leave open the problem of designing a polynomial time algorithm for this case (in particular, a polynomial time algorithm for computing ironed virtual valuations), noting that for the case when each of the distributions F_i is discrete and explicitly specified, a simple algorithm for computing ironed virtual valuations has been given by Elkind [8], and this implies a polynomial-time approximation algorithm for the non-regular case with discrete explicit distributions.

5. A POLYNOMIAL-TIME APPROXIMATION ALGORITHM

We now describe how to implement our algorithm for the regular case in the two computational models described in Section 2.

Implementation in the discrete explicit model is straightforward. Although we have focused on continuous distributions in Sections 3 and 4, we remark that virtual valuations and their inverses for discrete distributions can be defined and computed in much the same way as for continuous distributions. Our algorithm computes virtual valuations of all possible values for each item. During this process it keeps track of $F_i(\phi_i^{-1}(\nu))$. It then picks the least non-negative ν , with $\prod_i F_i(\phi_i^{-1}(\nu)) \geq 1/2$. The price of each item i is then defined to be the minimum value v_i at which $\phi_i(v_i) \geq \nu$. Each step of the algorithm takes at most linear time in n and the sizes of the supports. The resulting algorithm is a Monte Carlo randomized algorithm, i.e. it outputs a random pricing whose revenue, *in expectation*, approximates the revenue of the optimal pricing. It is natural to ask whether there is also a Las Vegas algorithm, i.e. one whose output is a good approximation to the revenue of the optimal pricing with probability 1. This question is especially natural given the problem's economic motivation: a firm is likely to feel much safer using an algorithm which always selects approximately optimal prices rather than one which has a small probability of setting disastrously suboptimal prices. However, in this paper we will not consider the question of whether there is a Las Vegas randomized algorithm to compute the optimal pricing.

In order to obtain an implementation in the continuous model with oracles, we use the following lemma from [3]. We present a brief proof for completeness.

LEMMA 19. Let \mathbf{p} be any price vector, and \mathbf{p}' be another price vector such that $p_i' \in [\beta, \alpha]p_i$ for all i with $\beta < \alpha < 1$. Then $\mathcal{R}^{\mathbf{p}'} \geq \frac{\beta(1-\alpha)}{(1-\beta)}\mathcal{R}^{\mathbf{p}}$.

PROOF. Consider any valuation vector \mathbf{v} , and let i be the index that maximizes $v_i - p_i$. In other words, when prices are given by \mathbf{p} and a consumer has values \mathbf{v} , the consumer buys item i . On the other hand, let j be the index that maximizes $v_j - p_j'$. That is, when the prices are given by \mathbf{p}' , the same consumer buys item j instead of i . The lemma follows from the claim that

$$p_j' \geq \frac{\beta(1-\alpha)}{(1-\beta)}p_i$$

To prove this claim, we first observe that $v_i - p_i \geq v_j - p_j$ and $v_j - p_j' \geq v_i - p_i'$. Rearranging terms and adding the

two we get $p_i - p_i' \leq p_j - p_j'$. Finally, using $p_i' \leq \alpha p_i$ and $p_j' \geq \beta p_j$, we get

$$p_j \geq \frac{(1-\alpha)}{(1-\beta)}p_i$$

The claim now follows by again using the fact that $p_j' \geq \beta p_j$. \square

Armed with this lemma, our algorithm essentially transforms the continuous case to a discrete version. Let $M = (\max_i h_i)/(\min_i \ell_i)$. Our algorithm will run in time polynomial in n and M . For each item i , we consider the set L_i of values that are powers of $\gamma = 1/(1-\epsilon)$ for some $\epsilon > 0$ in the range $[\ell_i, h_i]$. Note that $|L_i| = O(\log_\gamma \frac{h_i}{\ell_i}) = O(\log_\gamma M)$.

Our algorithm proceeds as follows:

1. For each i and each $v \in L_i$, compute $\phi_i(v)$ using the oracles for F_i and f_i and store these in a sorted list L' .
2. For each $\nu \in L'$, let $x_i(\nu)$ be $(1-\epsilon)$ times the largest value in L_i whose virtual value is at most ν . Note that $x_i(\nu) \in [(1-\epsilon)^2, (1-\epsilon)]\phi_i^{-1}(\nu)$. If $\prod_i F_i(x_i(\nu)) \geq 1/2$, remove ν from L' .
3. Add 0 to the set L' .
4. For each $\nu \in L'$, consider the pricing $\{x_i(\nu)\}$. Pick a sample $S(\nu)$ of size $\frac{64}{\epsilon^2}M^2 \log(nM/\epsilon^2)$ from the distribution \mathbf{F} . Let $\tilde{\mathcal{R}}(\nu)$ denote the expected revenue of the pricing $\{x_i(\nu)\}$ with respect to a uniform distribution over $S(\nu)$.
5. Let \mathbf{p} denote the pricing $\{x_i(\nu)\}$ for the virtual value ν that maximizes $\tilde{\mathcal{R}}(\nu)$. Output \mathbf{p} .

We first note that one of the pricings $\{x_i(\nu)\}$ for $\nu \in L'$ is a near optimal pricing.

LEMMA 20. Let ν^* be the minimum non-negative virtual valuation satisfying $\prod_i F_i(\phi_i^{-1}(\nu^*)) \leq 1/2$. Let \mathbf{p}^* be defined such that $p_i^* = \phi_i^{-1}(\nu^*)$. Then, there exists a virtual value $\nu \in L'$ such that

$$\mathcal{R}^{\{x_i(\nu)\}} \geq \frac{1}{2}(1-\epsilon)^2 \mathcal{R}^{\mathbf{p}^*}$$

PROOF. For every i let y_i be the largest power of $\gamma = 1/(1-\epsilon)$ smaller than p_i^* , and let $\nu_i = \phi_i(y_i)$. Let $\nu = \max_i \nu_i$. Note that we do not remove ν from the list L' in step 2 above. Consider $x_i(\nu)$. Note that $x_i(\nu) = (1-\epsilon)y_i$ by definition, which is less than $(1-\epsilon)p_i^*$. On the other hand, $y_i \geq (1-\epsilon)p_i^*$, so $x_i(\nu) \geq (1-\epsilon)^2 p_i^*$. Therefore, we have $x_i(\nu) \in [(1-\epsilon)^2, (1-\epsilon)]p_i^*$. Applying Lemma 19 with $\alpha = (1-\epsilon)$ and $\beta = (1-\epsilon)^2$ we get the result. \square

The next lemma shows that with a high probability, the estimates $\tilde{\mathcal{R}}(\nu)$ are good approximations to the true revenues $\mathcal{R}^{\{x_i(\nu)\}}$.

LEMMA 21. For any $\nu \in L'$,

$$\Pr \left[\left| \tilde{\mathcal{R}}(\nu) - \mathcal{R}^{\{x_i(\nu)\}} \right| \geq \epsilon \mathcal{R}^{\{x_i(\nu)\}} \right] \leq \frac{\epsilon^2}{n \log M}.$$

PROOF. Fix the pricing $\{x_i(\nu)\}$. For a valuation vector \mathbf{v} drawn from \mathbf{F} , let $Y_{\mathbf{v}}$ denote the revenue obtained by this pricing. Note that $\mathbf{E}[Y_{\mathbf{v}}] = \mathcal{R}^{\{x_i(\nu)\}}$. We also have $Y_{\mathbf{v}} \leq \max_i h_i$ for all \mathbf{v} , and $\mathcal{R}^{\{x_i(\nu)\}} \geq \frac{1}{4} \min_i \ell_i$ because we make a sale with probability at least $\frac{1}{4}$. Letting N denote the size of the sample $|S(\nu)|$, and applying the Chernoff bound, we get

$$\begin{aligned} & \Pr \left[\left| \tilde{\mathcal{R}}(\nu) - \mathcal{R}^{\{x_i(\nu)\}} \right| \geq \epsilon \mathcal{R}^{\{x_i(\nu)\}} \right] \\ &= \Pr \left[\left| \frac{\sum_{\mathbf{v} \in S(\nu)} Y_{\mathbf{v}}}{N} - \mathcal{R}^{\{x_i(\nu)\}} \right| \geq \epsilon \mathcal{R}^{\{x_i(\nu)\}} \right] \\ &\leq 2 \exp \left(-\frac{1}{4} \frac{\epsilon^2 (\mathcal{R}^{\{x_i(\nu)\}})^2 N}{(\max_i h_i)^2} \right) \\ &\leq 2 \exp \left(-\frac{1}{64} \frac{\epsilon^2 N}{M^2} \right) \end{aligned}$$

Using $N > \frac{64}{\epsilon^2} M^2 \log \left(\frac{2}{\epsilon^2} n \log M \right)$, we get the result. \square

THEOREM 22. *For any $\epsilon > 0$, the above algorithm gives a $8 + \epsilon$ approximation to the BUPP in time polynomial in n , M and $1/\epsilon$.*

PROOF. Using Lemma 21 and taking a union bound over all $\nu \in L'$, we get that with probability $1 - \epsilon$ all the estimates $\tilde{\mathcal{R}}(\nu)$ are within a $1 \pm \epsilon$ factor of the true revenues $\mathcal{R}^{\{x_i(\nu)\}}$. Therefore, the true revenue of the pricing picked by our algorithm is within a $(1 - \epsilon)^2$ factor of the maximum over $\nu \in L'$ of the true revenues $\mathcal{R}^{\{x_i(\nu)\}}$. Combining this with Lemma 20, Theorem 15 and the ϵ probability of failure, we get that the expected revenue obtained by our algorithm is a $8/(1 - \epsilon)^5$ approximation to the optimal revenue. An appropriate choice of ϵ implies the result. \square

6. CONCLUSIONS

Several interesting questions related to BUPP still remain open:

- Is the Bayesian unit-demand pricing problem with independently distributed values NP-hard to solve optimally? There is some evidence that this problem is indeed hard. For example, one can construct two-item instances with extremely simple distributions (e.g., a uniform distribution over some range), where the optimal price is irrational.
- Is our characterization tight? Can one construct an example where the revenue of Myerson's auction is indeed 4 times the revenue of the optimal pricing?

It is worth noting that there is a simple example in which the revenue of the pricing defined by our virtual valuation technique falls short of the optimal pricing by a factor of nearly 2, even in the i.i.d. case. Suppose that for each i , the distribution of v_i is given by

$$\begin{aligned} \Pr[v_i = n] &= \frac{1}{n^2} \\ \Pr[v_i = 1] &= 1 - \frac{1}{n^2}. \end{aligned}$$

The optimal pricing sets $p_1 = 1$ and $p_i = n$ for all $i > 1$. This achieves a revenue of $2 - o(1)$. However, for every ν the pricing which sets $p_i = \phi^{-1}(\nu)$ achieves a revenue of at most 1. (In this example the revenue of Myerson's

auction is nearly equal to 2, so the example does not prove any separation between the revenue of Myerson's auction and that of the optimal pricing.)

- Extending this work to accommodate combinatorial consumers seems tricky. An optimal pricing in that case may offer bundles at prices higher or lower than the sum of the prices of individual items in the bundle.
- Finally, a more general selling mechanism in the unit-demand case may offer lotteries to consumers. A lottery is a distribution over single items, sold at a price (typically) lower than the prices for the individual items. The revenue of the optimal collection of lotteries is not always bounded above by the revenue of Myerson's auction. In fact when values are correlated, the revenue of the optimal single-item pricing can be an exponential factor smaller than the revenue of the optimal collection of lotteries.

Acknowledgments

We thank Nina Balcan for several useful discussions.

7. REFERENCES

- [1] G. Aggarwal, T. Feder, R. Motwani, and A. Zhu. Algorithms for multi-product pricing. In *Proc. of ICALP*, 2004.
- [2] G. Aggarwal and J. Hartline. Knapsack auctions. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2006.
- [3] M.-F. Balcan, A. Blum, J. Hartline, and Y. Mansour. Mechanism Design via Machine Learning. In *Proc. of the 46th IEEE Symp. on Foundations of Computer Science*, 2005.
- [4] N. Balcan and A. Blum. Approximation algorithms and online mechanisms for item pricing. In *Proc. 8th ACM Conf. on Electronic Commerce*, 2006.
- [5] P. Briest and P. Krysta. Buying cheap is expensive: Hardness of non-parametric multi-product pricing. In *Proc. 18th ACM Symp. on Discrete Algorithms*. ACM/SIAM, 2007. To appear.
- [6] J. Bulow and J. Roberts. The simple economics of optimal auctions. *The Journal of Political Economy*, 97:1060–90, 1989.
- [7] E. Demaine, U. Feige, M. Hajiaghayi, and M. Salavatipour. Combination can be hard: Approximability of the unique coverage problem. In *Proc. 17th ACM Symp. on Discrete Algorithms*, 2006.
- [8] E. Elkind. Designing and learning optimal finite support auctions. In *Proc. 18th ACM Symp. on Discrete Algorithms*, 2007.
- [9] V. Guruswami, J. Hartline, A. Karlin, D. Kempe, C. Kenyon, and F. McSherry. On profit-maximizing envy-free pricing. In *Proc. 16th ACM Symp. on Discrete Algorithms*, 2005.
- [10] J. Hartline and V. Koltun. Near-Optimal Pricing in Near-Linear Time. In *Proc. of 9th Workshop on Algorithms and Data Structures*. Springer-Verlag, 2005.
- [11] V. Krishna. *Auction Theory*. Academic Press, 2002.
- [12] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.

[13] R. Tyrrell Rockafellar. *Convex Analysis*. Princeton University Press, 1970.

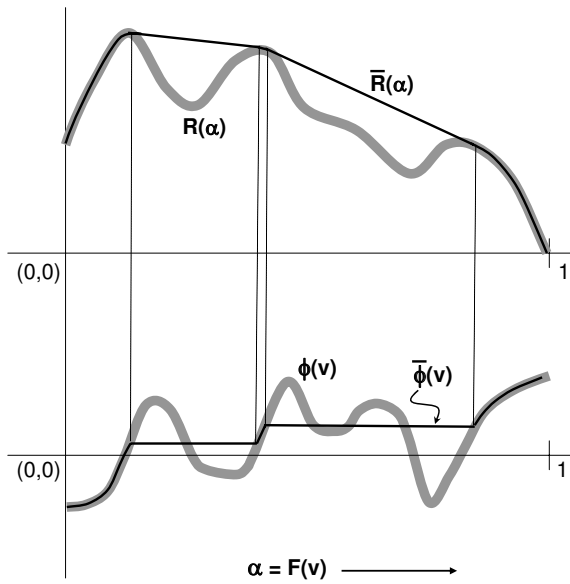


Figure 1: Converting a virtual valuation function ϕ to \bar{R} and $\bar{\phi}$