

Optimal and Efficient Parametric Auctions

Pablo Azar
MIT, CSAIL
Cambridge, MA 02139, USA
azar@csail.mit.edu

Silvio Micali
MIT, CSAIL
Cambridge, MA 02139, USA
silvio@csail.mit.edu

Constantinos Daskalakis*
MIT, CSAIL
Cambridge, MA, 02139, USA
costis@csail.mit.edu

S. Matthew Weinberg†
MIT, CSAIL
Cambridge, MA 02139, USA
smweinberg@csail.mit.edu

October 3, 2012

Abstract

Consider a seller who seeks to provide service to a collection of interested parties, subject to feasibility constraints on which parties may be simultaneously served. Assuming that a distribution is known on the value of each party for service—arguably a strong assumption—Myerson’s seminal work provides revenue optimizing auctions [12]. We show instead that, for very general feasibility constraints, only knowledge of the *median* of each party’s value distribution, or any other quantile of these distributions, or approximations thereof, suffice for designing simple auctions that simultaneously approximate both the optimal revenue and the optimal welfare. Our results apply to all downward-closed feasibility constraints under the assumption that the underlying, unknown value distributions are monotone hazard rate, and to all matroid feasibility constraints under the weaker assumption of regularity of the underlying distributions. Our results jointly generalize the single-item results obtained by Azar and Micali [2] on parametric auctions, and Daskalakis and Pierrakos [6] for simultaneously approximately optimal and efficient auctions.

1 Introduction

We study the problem of a seller who seeks to maximize her revenue when auctioning off service to a collection of interested buyers. Each buyer i is willing to pay some (private) v_i for receiving service, while the seller may have constraints on which buyers can receive service simultaneously. Traditionally, starting with the seminal work of Myerson [12], it is assumed that the values v_1, \dots, v_n are drawn from independent distributions F_1, \dots, F_n , and that these distributions are known to the seller. This assumption is often too strong, and we may instead want to construct mechanisms that guarantee good revenue, even when the seller has more limited knowledge about the buyers’ values. In this paper, we study the design of auctions under the much weaker assumption that the seller only knows some *parameters* of the distributions F_1, \dots, F_n , generalizing work by Azar and Micali [2] who construct parametric single-item and digital goods auctions which only use the mean and variance of each distribution. Our new results allow the seller to construct competitive auctions even when she only knows the medians of the distributions, or other quantiles, or approximations thereof, and apply to much more general settings beyond single-item and digital goods settings as discussed below. We complement these results by providing upper bounds on the attainable fraction of the optimal revenue by a seller who only knows the medians of the value distributions.

In addition, it is often desirable to guarantee good approximations to the optimal revenue *and* welfare simultaneously. We present two informal motivational examples where this is the case. The first example is a spectrum auction. The government does not just want to maximize the income it obtains from this sale. It also

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

†Supported by a NSF Graduate Research Fellowship and NSF award CCF-1101491.

wants to ensure that the spectrum frequencies are sold to companies who can use them best. A second example is a repeated auction. If the auction does not guarantee good welfare for the buyers, they will not return as customers. Therefore, a simultaneous revenue/welfare guarantee is desirable to any auctioneer, even one who only cares about revenue.

The auctions we construct will simultaneously approximate the optimal welfare and the optimal revenue achievable by a seller who knows the distributions F_1, \dots, F_n . This improves on the work of Daskalakis and Pierrakos [6], who show how to construct a single-item auction that simultaneously guarantees 20% of the optimal revenue and the optimal welfare, when the seller knows the distributions. We show how to construct such auctions when the seller only knows parameters of F_1, \dots, F_n , and our results hold for much more general auction settings. In particular, our results provide *a simple, parametric, and approximately revenue- and welfare- optimal auction for arbitrary matroid environments when the underlying distributions are regular, and arbitrary downward-closed environments when the underlying distributions are monotone hazard rate.*

1.1 Techniques Our main technical contribution is a reduction from downwards-closed single-dimensional multi-bidder auction design to pricing a single item for a single bidder. Our reduction will apply even when the seller has limited information. Formally we show that, if P_i is a (possibly randomized) pricing scheme that obtains an α -fraction of the optimal revenue for selling a single item to bidder i , then running a VCG auction with a separate reserve for each bidder i defined according to P_i obtains a $(\alpha/2)$ -fraction of the optimal revenue for the single-dimensional downwards-closed multi-bidder setting. The only necessary condition for our reduction to hold is that the VCG auction with optimal monopoly reserve prices obtain a $\frac{1}{2}$ approximation to the optimal revenue. It is shown in [7] that this holds in all downwards-closed settings when each F_i is MHR, and in all matroid settings when each F_i is regular. So in these settings our reduction can be readily applied.

While our results build upon those of [7], our techniques are orthogonal. Our main result is a reduction from the n bidder auction problem with downward closure constraints to the single-bidder pricing problem. Our reduction not only simultaneously preserves approximations for revenue and welfare, but also works even when the seller has arbitrary limited information about the distribution.

1.2 Related Work The Wilson doctrine [13] suggests that practical mechanisms should be *detail-free*. That is, they should not depend on prior information that the designer has about the buyers participating in the mechanism. Our work is in the spirit of the Wilson doctrine, in that we attempt to use as little information as possible about the buyers' valuations.

Goldberg, Hartline, Karlin, Saks and Wright [8] construct prior-free multi-unit auctions. They show that, even when buyers' values are not drawn from a distribution, they can still guarantee revenue equal to at least a constant fraction of some benchmark $\mathcal{F}_2(v_1, \dots, v_n)$. As made explicit by Hartline and Roughgarden [10], any auction that guarantees a constant fraction of this benchmark for all value vectors (v_1, \dots, v_n) , also guarantees a constant fraction of the optimal revenue when the values are drawn from identical and independent distributions. However, when the distributions are not identical, there are examples where no prior-free auction can guarantee a constant fraction of the optimal revenue.

Our auctions follow the framework of simple, approximately optimal auctions proposed by Hartline and Roughgarden [11], where it is shown that VCG auctions with monopoly reserve prices obtain a constant fraction of the optimal revenue.

Our work is most closely related to that of Dhangwatnotai, Roughgarden and Yan [7], who show how to guarantee approximately optimal revenue when the seller only has access to one sample from the joint distribution $F_1 \times \dots \times F_n$. This is already a strong result in the area of maximizing revenue with limited information. Our results build on theirs in a few ways: First, their techniques are inherently limited to sampling-based information, as they obtain their results through Bulow-Klemperer type inequalities. Our techniques apply to any collection of "good" pricing schemes for selling a single item to each single bidder. Second, our auctions obtain a simultaneous revenue/welfare guarantee, which is often desirable in practice.

In addition, while there are certainly many practical settings where it is easy to access a single sample from each value distribution, it is not hard to imagine a setting where this is difficult. Realistically, such a sample would be obtained by sampling a bidder from each population F_i and learning their value for obtaining service (possibly from existing data). However, it is often the case that bidders do not know exactly their value for obtaining service. It is generally regarded as *much* easier for a bidder to evaluate whether they are willing to purchase service at a given price. In this sense, it may be easier in practice to (approximately) learn quantiles of each value distribution by asking

questions of this form, rather than obtaining even a single sample. It is also easy to imagine settings where it is easier to obtain samples, or learn the means, or other useful, albeit limited information. Rather than argue which type of information is most accessible in practice, our results accommodate all types: if the seller has enough information to give a good approximation to the optimal revenue and welfare for the single-bidder problem, then she has enough information to give a good approximation to the optimal revenue and welfare for multi-unit auctions with many bidders. Finally, our results come with a nice robustness guarantee: any approximation error in estimating the statistics can be directly absorbed into the approximation error for revenue and welfare.

2 Preliminaries

Single-Dimensional Environments In our model, there is one seller who can provide a service to n buyers. The seller has some constraints over the sets of buyers she can serve simultaneously. These constraints are represented by a collection of sets $\mathcal{I} \subset 2^{\{1, \dots, n\}}$. A set $S \subseteq \{1, \dots, n\}$ of buyers can only be served simultaneously if $S \in \mathcal{I}$. We call this setting a *downward closed environment* if for every $S \in \mathcal{I}$ and every $T \subset S$, we have $T \in \mathcal{I}$. We call such a setting a *matroid environment* if the collection of sets \mathcal{I} forms a matroid.¹

Valuations Each buyer has a private value, v_i , for being served, sampled independently from a continuous distribution over positive real numbers. We denote by $F_i(x) = \Pr[v_i \leq x]$ the cumulative density function of this distribution, and by $f_i(x)$ its probability density function. We denote by $v = (v_1, \dots, v_n)$ the vector of values, and by $F = F_1 \times \dots \times F_n$ the joint product distribution from which it is drawn. When we want to emphasize the value v_i , we will write $v = (v_i, v_{-i})$.

Truthful Auctions Let $b = (b_1, \dots, b_n)$ denote a vector of bids. An auction is given by a pair of vector functions (x, p) , where $x_i(b) \in [0, 1]$ denotes the probability that bidder i will receive service, and $p_i(b) \in \mathbb{R}^+$ denotes the expected price that bidder i will be charged under bid vector b (where the probability and the expectation are taken with respect to the randomness in the auction). Given the bid vector b , bidder i 's utility from participating in the auction is $x_i(b) \cdot v_i - p_i(b)$, where v_i is his value. We say that the auction is *dominant strategy truthful* if, for each bidder i , all v_i 's in the support of F_i , and all bid vectors $b = (b_i, b_{-i})$, we have $x_i(v_i, b_{-i}) \cdot v_i - p_i(v_i, b_{-i}) \geq x_i(b_i, b_{-i}) \cdot v_i - p_i(b_i, b_{-i})$.

¹We recall that \mathcal{I} is a matroid if it is downward closed, and furthermore, for any $S, T \in \mathcal{I}$ with $|T| < |S|$, there exists an element $e \in S - T$ such that $T \cup \{e\} \in \mathcal{I}$.

$v_i - p_i(b_i, b_{-i})$. It is well known [12, 1] that an auction is dominant strategy truthful if and only if x_i is increasing in b_i , and $p_i(b_i, b_{-i}) = x_i(b_i, b_{-i})b_i - \int_0^{b_i} x_i(z; b_{-i})dz$. When the auction is deterministic (i.e. $x(b)$ is a 0/1 vector for all b), this implies that there exists a *reserve price* $p_i(b_{-i})$ such that bidder i obtains service if and only if $b_i \geq p_i(b_{-i})$, and pays the reserve price only if he obtains service.

Revenue and Social Welfare Given a valuation vector v , a bid vector b and auction $A = (x, p)$, its revenue is $\sum_{i=1}^n p_i(b)$, and its social welfare is $\sum_{i=1}^n x_i(b)v_i$. Given an auction A , we denote by $Rev(A, F)$ its expected revenue and by $SW(A, F)$ its expected social welfare when the bidder values are drawn from distribution F and the bidders bid their values truthfully.

Monotone Hazard Rate and Regular Distributions Given a distribution F over the real numbers, its *hazard rate* is $h_i(v) = \frac{f_i(v)}{1 - F_i(v)}$. The distribution has a monotone hazard rate (MHR) if $h_i(v)$ is increasing. The virtual valuation function associated with the distribution is $\phi_i(v_i) = v_i - \frac{1}{h_i(v)}$. We say that the distribution F is regular if its associated virtual valuation function is increasing. Clearly, any monotone hazard rate distribution is also regular.

Virtual Surplus and Optimal Auctions Given a product distribution $F = F_1 \times \dots \times F_n$ and a valuation vector (v_1, \dots, v_n) , the virtual surplus is defined as $\sum_{i=1}^n \phi_i(v_i)$. When the distributions F_1, \dots, F_n are regular, Myerson's optimal auction, which maximizes virtual surplus on every profile obtains the maximum expected revenue over all truthful auctions with knowledge of F_i [12].² We denote by $Rev(\text{OPT}, F)$ the expected revenue of the optimal auction. Given an auction A , its *revenue competitive ratio* is defined as $\frac{Rev(A, F)}{Rev(\text{OPT}, F)}$.

VCG auctions Myerson's auction maximizes revenue, but it may not always serve the bidders who have the highest value for service. The VCG auction is a truthful auction that maximizes social welfare. It serves a set of bidders S such that $S \in \arg\max_{T \in \mathcal{I}} \sum_{i \in T} b_i$. If bidder i is served, her price is $\max_{T \in \mathcal{I}} \sum_{j \in T - \{i\}} b_j$. We denote by $SW(\text{OPT}, F)$ the social welfare generated by the VCG auction when the bidders' values are drawn from F . Given any auction A and distribution F , its *social welfare competitive ratio* is defined as $\frac{SW(A, F)}{SW(\text{OPT}, F)}$. Our goal is to design truthful auctions whose revenue competitive ratio and social welfare competitive ratio are simultaneously bounded below by a constant.

²When the distributions are not necessarily regular, Myerson shows that the optimal auction maximizes *ironed* virtual valuations. We refer the reader to [12] for more details.

VCG auctions with Lazy Reserves Finally, we will use the following type of auction from [7]. For the following definition, each P_i denotes some randomized pricing scheme.³

DEFINITION 1. *VCG- L_P (Lazy VCG with reserves $P = (P_1, \dots, P_n)$): First, run the VCG auction to determine the set of (candidate) winners, W , who would receive service under VCG. Let VCG_i denote the price VCG would charge bidder i . Then, for each $i \in W$, sample a price p_i from P_i and offer service to bidder i at price $\max\{p_i, VCG_i\}$. (The bidder chooses whether or not to accept service at this price.)*

3 A reduction from downward closed single-dimensional environments to pricing

In this section, we will show a generic reduction from single-dimensional downward-closed environments to pricing a single item for a single bidder. Theorem 3.1 below states this formally. First, we must prove two simple facts about $VCG-L_P$. These facts are immediate when P is a deterministic pricing scheme, and are not hard to see in general. We include them below for completeness.

LEMMA 3.1. *VCG- L_P is truthful for all P .*

Proof. Bidders cannot control the price offered to them once they are in W . Therefore, it is clear that bidders have nothing to gain by changing their bid but staying above/below VCG_i . Changing the bid from above VCG_i to below VCG_i can only cost bidder i service when they might have been willing to pay for it. Changing from below VCG_i to above VCG_i also cannot increase the buyers' utility because the price offered for receiving service is always at least VCG_i . \square

LEMMA 3.2. *Let $q'_i(v_i)$ denote the probability that bidder i receives service in VCG- L_P , conditioned on v_i and $i \in W$. Let also $q_i(v_i)$ denote the probability that bidder i receives service in P_i conditioned on v_i . Then $q'_i(v_i) = q_i(v_i)$ for all i, v_i .*

Proof. Let p_i be a random variable denoting the price drawn from P_i . Then the probability that bidder i receives service in VCG- L_P conditioned on v_i and $i \in W$ is exactly the probability that $v_i \geq \max\{p_i, VCG_i\}$. As $i \in W$, we must have $v_i \geq VCG_i$. So the probability that $v_i \geq \max\{p_i, VCG_i\}$ is exactly the probability that $v_i \geq p_i$, which is exactly $q_i(v_i)$. \square

³A randomized pricing scheme samples a price p_i from P_i and offers price p_i .

THEOREM 3.1. *Let $F = F_1 \times \dots \times F_n$ be a product distribution, where each marginal F_i is regular. Let $P = (P_1, \dots, P_n)$ be any collection of randomized single-item, single-bidder pricing schemes that each obtain an α -fraction of the optimal expected revenue (for selling a single item to bidder i respectively). Then in all environments where VCG-L with Myerson reserves obtains expected revenue at least $\frac{1}{2} \text{Rev}(\text{OPT}, F)$, VCG- L_P obtains expected revenue at least $\frac{\alpha}{2} \text{Rev}(\text{OPT}, F)$.*

Proof. We will prove the theorem by showing directly that the expected virtual surplus of VCG- L_P is at least an α -fraction of the expected virtual surplus of Lazy VCG with Myerson reserves. As expected virtual surplus equals exactly expected revenue, and Lazy VCG with Myerson reserves obtains half the expected virtual surplus of Myerson's auction by hypothesis, this suffices to prove the theorem.

Let's first write an expression for the expected virtual surplus of Lazy VCG with Myerson reserves. For each possible value v_i , there is some expected probability that $i \in W$, conditioned on v_i .⁴ Denote this value by $\pi_i(v_i)$. If $i \in W$ and additionally v_i exceeds the Myerson reserve, M_i , bidder i receives service. Otherwise, he receives nothing. So the expected virtual surplus of Lazy VCG with Myerson reserves is:

$$\sum_i \int_{M_i}^{\infty} f_i(v_i) \cdot \pi_i(v_i) \cdot \phi_i(v_i) dv_i.$$

This is simply integrating over all $i, v_i \geq M_i$, the probability that bidder i 's value is v_i , times the probability that bidder i is in W conditioned on v_i , times the virtual value of bidder i with value v_i .

Let's now write an expression for the expected virtual surplus of VCG- L_P . Below, $q_i(v_i)$ denotes the probability that bidder i receives service in VCG- L_P conditioned on v_i and $i \in W$:

$$\sum_i \int_0^{\infty} f_i(v_i) \cdot \pi_i(v_i) \cdot q_i(v_i) \cdot \phi_i(v_i) dv_i$$

This again is integrating over all i, v_i the probability that bidder i 's value is v_i , times the probability that bidder i is in W conditioned on v_i , times the probability that VCG- L_P awards them service conditioned on v_i and $i \in W$, times the virtual value of bidder i with value v_i .

By Lemma 3.2, $q_i(v_i)$ is exactly the probability that bidder i with value v_i is awarded service in P_i . As P_i obtains an α -fraction of the optimal expected revenue

⁴This expectation is taken over the randomness in b_{-i} assuming every bidder reports truthfully.

when selling a single item to bidder i , we have the following inequalities for all i :

$$(3.1) \quad \int_0^\infty f_i(v_i) \cdot q_i(v_i) \cdot \phi_i(v_i) dv_i \geq \alpha \int_{M_i}^\infty f_i(v_i) \cdot \phi_i(v_i) dv_i.$$

Equation (3.1) simply states that the expected virtual surplus of P_i is at least α times that of using Myerson's reserve. We can rewrite Equation (3.1) as:

$$(3.2) \quad \int_0^{M_i} f_i(v_i) \cdot q_i(v_i) \cdot \phi_i(v_i) dv_i + \int_{M_i}^\infty f_i(v_i) \cdot (q_i(v_i) - \alpha) \cdot \phi_i(v_i) dv_i \geq 0.$$

Now, observe first that $f_i(v_i)$ and $q_i(v_i)$ are both non-negative, and $\phi_i(v_i) \leq 0$ for all $v_i \leq M_i$. So the first term only integrates over negative values. In addition, we may denote by x_i the smallest value of v_i such that $q(x_i) \geq \alpha$ and $x_i \geq M_i$. Then for all $M_i \leq v_i \leq x_i$, $q_i(v_i) - \alpha \leq 0$, $f_i(v_i) \geq 0$ and $\phi_i(v_i) \geq 0$, so the second term only integrates negative values from M_i through x_i . For all $v_i \geq x_i$, $q_i(v_i) - \alpha \geq 0$, so the second term only integrates positive values from x_i through ∞ . Therefore, we may again rewrite Equation (3.2) as:

$$(3.3) \quad \int_{x_i}^\infty f_i(v_i) \cdot (q_i(v_i) - \alpha) \cdot \phi_i(v_i) dv_i - \int_0^{M_i} f_i(v_i) \cdot q_i(v_i) \cdot \phi_i(v_i) dv_i - \int_{M_i}^{x_i} f_i(v_i) \cdot (q_i(v_i) - \alpha) \phi_i(v_i) dv_i.$$

Now, observe further that π_i is an increasing function of v_i . As each side of Equation (3.3) integrates only positive terms, we obtain the following two inequalities:

$$(3.4) \quad \int_{x_i}^\infty \pi_i(v_i) \cdot f_i(v_i) \cdot (q_i(v_i) - \alpha) \cdot \phi_i(v_i) dv_i \geq \pi_i(x_i) \cdot \left(\int_{x_i}^\infty f_i(v_i) \cdot (q_i(v_i) - \alpha) \cdot \phi_i(v_i) dv_i \right);$$

$$(3.5) \quad -\pi_i(x_i) \int_0^{M_i} f_i(v_i) \cdot q_i(v_i) \cdot \phi_i(v_i) dv_i - \pi(x_i) \int_{M_i}^{x_i} f_i(v_i) \cdot (q_i(v_i) - \alpha) \phi_i(v_i) dv_i \geq - \int_0^{M_i} \pi_i(v_i) \cdot f_i(v_i) \cdot q_i(v_i) \cdot \phi_i(v_i) dv_i - \int_{M_i}^{x_i} \pi_i(v_i) \cdot f_i(v_i) \cdot (q_i(v_i) - \alpha) \phi_i(v_i) dv_i.$$

Putting Equations (3.3), (3.4) and (3.5) together, we obtain the following inequality:

$$(3.6) \quad \int_{x_i}^\infty \pi_i(v_i) \cdot f_i(v_i) \cdot (q_i(v_i) - \alpha) \cdot \phi_i(v_i) dv_i \geq - \int_0^{M_i} \pi_i(v_i) \cdot f_i(v_i) \cdot q_i(v_i) \cdot \phi_i(v_i) dv_i - \int_{M_i}^{x_i} \pi_i(v_i) \cdot f_i(v_i) \cdot (q_i(v_i) - \alpha) \phi_i(v_i) dv_i.$$

Finally, we may rearrange Equation (3.6) back to obtain:

$$(3.7) \quad \int_0^\infty \pi_i(v_i) \cdot f_i(v_i) \cdot q_i(v_i) \cdot \phi_i(v_i) dv_i \geq \alpha \left(\int_{M_i}^\infty \pi_i(v_i) \cdot f_i(v_i) \cdot \phi_i(v_i) dv_i \right).$$

After summing over all i , Equation (3.7) exactly says that the expected virtual surplus of $VCG-L_P$ is at least an α -fraction of the expected virtual surplus of Lazy VCG with Myerson reserves and completes the proof. \square

Formally proving Theorem 3.1 requires analyzing virtual surplus and carefully moving around equations. But there is a clean intuition as to why Theorem 3.1 holds. Assume we are forced to use a Lazy VCG mechanism with some reserves, and are just trying to pick good reserves. Conditioned on $i \in W$ and fixing the remaining bids, this is exactly coming up with a good pricing scheme for a single bidder sampled from F_i , conditioned on $v_i \geq VCG_i$. It is not hard to see that the optimal solution for this single-bidder pricing problem is to set price $\max\{M_i, VCG_i\}$. What Theorem 3.1 is claiming is that if P_i is a good pricing scheme for F_i , then sampling p_i from P_i and setting price $\max\{p_i, VCG_i\}$ is a good pricing scheme for F_i , conditioned on $v_i \geq VCG_i$.

Virtually the same proof yields a similar theorem for welfare:

THEOREM 3.2. *Let $F = F_1 \times \dots \times F_n$ be any product distribution. Let also $P = (P_1, \dots, P_n)$ be any collection of randomized single-item, single-bidder pricing schemes that each obtain an α -fraction of the optimal expected welfare (for selling a single item to bidder i respectively). Then $VCG-L_P$ obtains expected welfare at least $\alpha SW(\text{OPT}, F)$.*

Proof. Virtually the same proof as Theorem 3.1 works after some modifications. First, instead of comparing to Lazy VCG with Myerson reserves, we just compare to VCG. Second, we compare expected welfare rather than expected virtual welfare. Moreover, as VCG always achieves exactly optimal expected welfare without any assumptions, we do not need to make any assumptions on F_i or lose anything in the approximation factor. Also, since we are comparing to VCG instead, we should replace M_i with 0 everywhere (think of 0 as VCG's reserve). Finally, since we are computing expected welfare instead of virtual welfare, we need to replace $\phi_i(v_i)$ with v_i everywhere. After making these replacements, the proof is identical. \square

4 Optimal and Efficient Parametric Auctions

Our reduction from downward-closed settings to single-bidder problems not only preserves approximately optimal revenue and welfare, but also does so when the seller has limited information about what buyers are willing to pay. In particular, if the seller only knows some parameters about each bidder i 's distribution, our reduction allows us to build new parametric auctions for downward-closed and matroid environments by solving the simpler single bidder versions.

Median, quantiles, and approximations thereof. We show that, if the underlying distributions are regular, then knowledge of the median of these distributions allows for a simultaneous $\frac{1}{4}$ -approximation of the optimal revenue and welfare. This guarantee is robust to approximate knowledge of the median or some other quantile of the distributions. We begin by recalling a lemma from Cai and Daskalakis [5].

LEMMA 4.1. [5] *Let F be a regular distribution and let $R_F(x) = x \cdot F^{-1}(1-x)$ be the revenue curve in quantile space (i.e. $R_F(x)$ is the expected revenue for setting price $F^{-1}(1-x)$ which has probability of sale exactly equal to x). Then for all $0 < \bar{q} \leq q < 1$,*

$$R_F(q) \geq (1-q)R_F(\bar{q}).$$

Using this lemma, we can easily obtain the following corollary:

COROLLARY 4.1. *Let F be a regular distribution, and let $p = F^{-1}(1-q)$. Then the single-item, single-bidder pricing scheme that sets price p obtains at least a $\min\{q, 1-q\}$ -fraction of the optimal revenue and at least a q -fraction of the optimal welfare.*

Proof. That a q -fraction of the optimal welfare is achieved is trivial: the pricing scheme achieves welfare exactly equal to $q \cdot \mathbb{E}[v|v \geq p]$. As conditioning on $v \geq p$ can only increase the expected welfare, this is at least $q \cdot \mathbb{E}[v]$, q times the maximum possible expected welfare.

That a $\min\{q, 1-q\}$ -fraction of the optimal revenue is achieved comes from the following argument, where M denotes Myerson's reserve. Maybe $p \geq M$. In this case, Myerson's pricing scheme clearly makes at most M revenue. As $p \geq M$, the revenue obtained by setting price p is exactly $qp \geq qM$. Maybe $p \leq M$. In this case, we can use Lemma 4.1 to observe that the revenue obtained by setting price p is at least $(1-q)$ times the revenue obtained by Myerson's pricing scheme. Therefore, no matter where Myerson's reserve lies with respect to p , we obtain a $\min\{q, 1-q\}$ -fraction of the optimal revenue. \square

Using Corollary 4.1, Theorems 3.1 and 3.2, and [7] we immediately obtain the following:⁵

THEOREM 4.1. *Let P be any collection of single-bidder deterministic pricing schemes such that for all i , the probability of sale q_i satisfies $x \leq q_i \leq 1-x$. Then in all regular matroid environments and all MHR downward-closed environments, the auction $VCG-L_P$ simultaneously obtains an $(x/2)$ -fraction of the optimal revenue and an x -fraction of the optimal welfare. In particular, if $x = 1/2$ (i.e. we use the median as a reserve for each bidder), this is $\frac{1}{4}$ of the optimal revenue and $\frac{1}{2}$ of the optimal welfare.*

Observe that Theorem 4.1 is inherently robust: as long as we can learn some price with a reasonably good probability of sale, we obtain good revenue. For instance, if we can find for every bidder some price p_i such that the probability of sale lies in $[1/4, 3/4]$, then we may take $x = 1/4$ in Theorem 4.1 and get $\frac{1}{8}$ of the optimal expected revenue and $\frac{1}{4}$ of the optimal expected welfare.

We also show that, for a large class of distributions, a similar result can be obtained with just knowledge of the mean and standard deviation. We first recall what it means for a distribution to be c -informative and define a variant of symmetric distributions called c -symmetric.

⁵Recall that it was shown in [7] that $VCG-L$ with Myerson reserves obtains a $\frac{1}{2}$ approximation to the optimal revenue in all downwards-closed settings, when each F_i is MHR, and in all matroid settings, when each F_i is regular.

c-informative and c-symmetric distributions

Let V be a positive real random variable with mean μ and variance σ^2 . We say that the distribution of V is c -informative if $\frac{\mu}{\sigma} \geq c$. We denote by \mathcal{I}_c the class of all c -informative distributions, and remark that any monotone hazard rate distribution is 1-informative [3]. We say that the distribution of V is c -symmetric if $\Pr[V \geq \mu] \geq c$, and note that any symmetric distribution is $\frac{1}{2}$ -symmetric. We denote by \mathcal{S}_c the class of all c -symmetric distributions. We remark that any monotone hazard rate distribution is $\frac{1}{e}$ -symmetric [4].

We show that, when a buyer's distribution belongs to \mathcal{F}_c or \mathcal{S}_c for some constant c , we can guarantee a constant fraction of the optimal revenue in the pricing problem for that buyer with only knowledge of the mean and variance of the buyer's distribution. As a consequence, when the underlying distributions are regular (monotone hazard rate) we are able to guarantee a constant fraction of the optimal revenue in multi-bidder auctions with matroid (downward-closed) constraints.

LEMMA 4.2. *If V is a random variable with a c -symmetric distribution, then knowledge of the mean $\mu = \mathbb{E}[V]$ suffices to guarantee a c -approximation to the optimal revenue and welfare for the single-bidder single-item problem where the buyer's value has the same distribution as V .*

Proof. It is clear that the optimal pricing mechanism cannot obtain more than μ revenue in expectation. A mechanism that sets a posted price of μ sells to the buyer when $V \geq \mu$. Since $\Pr[V \geq \mu] \geq c$, we obtain a competitive ratio for revenue bounded from below by

$$\frac{\Pr[V \geq \mu] \cdot \mu}{\mu} \geq c.$$

Moreover, the optimal welfare is μ , while the welfare achieved by a posted price at μ is at least as much as the revenue achieved by a posted price at μ , which we have already shown is at least a c -fraction of μ . \square

Notice that we did not require regularity in the proof of the above lemma. Using our reductions (Theorems 3.1 and 3.2) from the many-bidder auction to the single-bidder pricing problem, we conclude then:

THEOREM 4.2. *Let $\mu = (\mu_1, \dots, \mu_n)$ be the vector of means of the buyers' value distributions, and assume that these distributions are c -symmetric. For regular matroid environments and MHR downward closed environments, the auction $VCG-L_\mu$ simultaneously guarantees a $\frac{c}{2}$ -fraction of the optimal revenue and a c -fraction of the optimal welfare.*

As noted earlier, it is known [4] that if a variable V is distributed according to a monotone hazard distribution, then it is $\frac{1}{e}$ -symmetric. Thus, we obtain the following corollary,

COROLLARY 4.2. *For any downward closed environment with monotone hazard rate distributions, the mechanism $VCG-L_\mu$ that sets the reserve price of bidder i at the mean μ_i of bidder i 's distribution guarantees a $\frac{1}{2e} \geq 18\%$ fraction of the optimal revenue and a $\frac{1}{e}$ -fraction of the optimal welfare.*

We remark that Hartline, Mirrokni and Sundararajan [9] obtain a similar result for digital good auctions with means as reserves in their analysis of marketing a digital good over a network. Our mechanism generalizes theirs to downward closed environments.

It is possible to construct a distribution with heavy left tails for which the mechanism $VCG-L_\mu$ guarantees arbitrarily low revenue and welfare. As an example, consider a single-bidder whose value takes a low value $\mu - \sigma t$ with probability $\frac{1}{1+t^2}$, and a high value $\mu + \frac{\sigma}{t}$ with probability $\frac{t^2}{1+t^2}$. As $t \rightarrow 0$, the expected welfare and revenue of $VCG-L_\mu$ become negligible. However, if we use information about both the mean and standard deviation to set appropriate reserve prices, we can still guarantee a constant fraction of the optimal revenue and welfare, as long as $\frac{\mu}{\sigma} > c$ for some constant c .

LEMMA 4.3. [2] *Let V be a random variable with mean μ and standard deviation σ . Let $k = k(\frac{\mu}{\sigma})$ be the unique real solution to the cubic equation $\frac{\mu}{\sigma} = \frac{1}{2}(k^3 + 3k)$, and let $\rho(\frac{\mu}{\sigma}) = 1 - \frac{\sigma}{\mu}k(\frac{\mu}{\sigma})$. Then a mechanism that sets a reserve price of $\mu - \sigma k$ guarantees a $\rho(\frac{\mu}{\sigma})$ -approximation to the optimal revenue and welfare of the single-bidder pricing problem where the buyer's value has the same distribution as V .*

One can show that $\rho(\cdot)$ is an increasing function. Thus, if $\frac{\mu}{\sigma} > c$ for some constant c , then the single-bidder pricing mechanism that sets a reserve price of $\mu - \sigma k(\frac{\mu}{\sigma})$ guarantees a $\rho(c)$ -approximation to the optimal revenue and welfare. Using the reduction from matroid and downward closed environments to the single-bidder pricing problem, we can derive the following theorem, which generalizes the single-item auction from [2] to matroid and downward closed environments.

THEOREM 4.3. *Let $F = F_1 \times \dots \times F_n$ be a product distribution with mean vector $\mu = (\mu_1, \dots, \mu_n)$ and standard deviation vector $\sigma = (\sigma_1, \dots, \sigma_n)$, and assume that the distributions F_1, \dots, F_n are c -informative. Let $\mu - \sigma k = (\mu_1 - \sigma_1 k_1, \dots, \mu_n - \sigma_n k_n)$ be a vector of reserve prices as in lemma 4.3. For regular matroid*

environments and MHR downward closed environments, the auction $VCG-L_{\mu-\sigma_k}$ simultaneously guarantees a $\frac{1}{2}\rho(c)$ -fraction of the optimal revenue and a $\rho(c)$ -fraction of the optimal welfare.

5 Upper Bounds on the Competitive Ratio

We provide some upper bounds on the best achievable revenue competitive ratio for the single-bidder pricing problem with only knowledge of the median. First, we show that the regularity assumption cannot be relaxed: for any deterministic pricing scheme P that only uses knowledge of the median and all $c > 0$, there is a (non-regular) distribution under which P does not obtain a c -fraction of the optimal revenue. Second, we show that no deterministic pricing scheme P that only has knowledge of the median can guarantee a worst-case approximation ratio better than $\frac{1}{2}$, even when the (unknown) distribution is regular. We highlight that all mechanisms we construct in this paper are deterministic and thus our bounds are tight for deterministic mechanisms.

LEMMA 5.1. *Let $P(m)$ be a (single-bidder single-item) deterministic pricing scheme that only uses the median m of the bidder's value distribution as input. Then for all $c > 0$ and $m > 0$, there exists a distribution F with median m such that $Rev(P(m), F) \leq c Rev(OPT, F) < \infty$.*

Proof. Fix m . We separate into two cases.

If $P(m) > m$, then the adversary can choose F to be uniform in $[m - \epsilon, m + \epsilon]$ where $m + \epsilon < P(m)$. This is a distribution with median m , but where the bidder has zero chance of being willing to buy the good for price $P(m)$. Hence, the expected revenue of a mechanism that sets $P(m) > m$ is zero, and the competitive ratio of P is zero.

If $P(m) \leq m$, the adversary can choose an arbitrarily large H and have the value drawn from the same distribution as the random variable $Y = \frac{m(1 + \frac{1}{\sqrt{H}})^2}{4} \cdot X$, where $X \in [1, H]$ is distributed according to $G(x) = \frac{1}{1 - \frac{1}{\sqrt{H}}} \cdot (1 - \frac{1}{\sqrt{x}})$. Note that the median of X is $\frac{4}{(1 + \frac{1}{\sqrt{H}})^2}$, and the median of Y is m . We have $1 - G(x) = \frac{\frac{1}{\sqrt{x}} - \frac{1}{\sqrt{H}}}{1 - \frac{1}{\sqrt{H}}}$. So the revenue curve associated with the random variable Y is

$$\begin{aligned} R\left(y = \frac{m(1 + \frac{1}{\sqrt{H}})^2}{4} \cdot x\right) &:= y \cdot (1 - G(y)) \\ &= \frac{m(1 + \frac{1}{\sqrt{H}})^2}{4} \cdot x \cdot \frac{\frac{1}{\sqrt{x}} - \frac{1}{\sqrt{H}}}{1 - \frac{1}{\sqrt{H}}}, \end{aligned}$$

which is maximized by setting $x = \frac{H}{4}$ and takes the maximum value $\frac{m(1 + \frac{1}{\sqrt{H}})^2}{4} \cdot \frac{\sqrt{H}}{4 \cdot (1 - \frac{1}{\sqrt{H}})}$, which grows

arbitrarily large with H . The revenue of an auction which sets a reserve price below the median is at most m . Thus, by choosing $H \gg m$, we can make the competitive ratio as small as we want. \square

LEMMA 5.2. *Let $P(m)$ be a (single-bidder single-item) deterministic pricing scheme that only uses the median m of the bidder's value distribution as input. For all $m > 0$, there is a regular distribution F with median m such that $Rev(P(m), F) \leq \frac{1}{2} Rev(OPT, F)$.*

Proof. Fix m . If $P(m) > m$, set $\epsilon = (P(m) - m)/2$ and let F be the uniform distribution on $[m - \epsilon, m + \epsilon]$. Then $P(m)$ makes expected revenue 0, but the optimal revenue is clearly at least $m/2$.

If $P(m) \leq m$, consider the distribution $F(x) = 1 - \frac{m}{m+x}$. Then F is regular.⁶ With knowledge of this distribution, it is easy to see that one can obtain revenue arbitrarily close to m by setting a sufficiently high reserve price. It is also easy to see that $Rev(P(m), F) = \frac{mP(m)}{m+P(m)}$. As $P(m) \leq m$, this is at most $m/2$. \square

6 Conclusion

We showed that one can reduce the problem of approximating revenue and social welfare in single-dimensional multi-bidder settings to the problem of approximating revenue and social welfare when selling one item to a single bidder. This provides a unified framework for the problem of designing auctions with limited information: If a seller has enough information to sell a single item to each bidder separately, then they have enough information to design an auction for all bidders simultaneously in very broad settings.

References

- [1] A. Archer and É. Tardos. Truthful mechanisms for one-parameter agents. In *the 42nd IEEE Symposium on Foundations of Computer Science (FOCS)*, 2001.
- [2] P. Azar and S. Micali. Optimal parametric auctions. Technical Report MIT-CSAIL-TR-2012-015, MIT CSAIL, 2012. In submission.
- [3] R.E. Barlow and A.W. Marshall. Bounds for distributions with monotone hazard rate, i. *The Annals of Mathematical Statistics*, 35(3):1234–1257, 1964.
- [4] R.E. Barlow, A.W. Marshall, and F. Proschan. Properties of probability distributions with monotone hazard rate. *The Annals of Mathematical Statistics*, 34(2):375–389, 1963.

⁶ F is such that $\phi(x) = -m$ for all $x \in [0, \infty)$. If one prefers to have an example where virtual values are strictly increasing, set $F(x) = 1 - \left(\frac{m}{m+x}\right)^{1+\epsilon}$ for arbitrarily small $\epsilon > 0$, and adjust so that the median stays exactly at m .

- [5] Y. Cai and C. Daskalakis. Extreme-Value Theorems for Optimal Multidimensional Pricing. In *the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2011.
- [6] C. Daskalakis and G. Pierrakos. Simple, Optimal and Efficient Auctions. In *the 7th Workshop on Internet and Network Economics (WINE)*, 2011.
- [7] P. Dhangwatnotai, T. Roughgarden, and Q. Yan. Revenue maximization with a single sample. In *the 11th ACM conference on Electronic commerce (EC)*, 2010.
- [8] A.V. Goldberg, J.D. Hartline, A.R. Karlin, M. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, 2006.
- [9] J. Hartline, V. Mirrokni, and M. Sundararajan. Optimal marketing strategies over social networks. In *the 17th international conference on World Wide Web (WWW)*, 2008.
- [10] J. D. Hartline and T. Roughgarden. Optimal mechanism design and money burning. In *the 40th annual ACM symposium on Theory of Computing (STOC)*, 2008.
- [11] J. D. Hartline and T. Roughgarden. Simple versus optimal mechanisms. In *the 10th ACM conference on Electronic commerce (EC)*, 2009.
- [12] R. B. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [13] R. Wilson. Game-theoretic analysis of trading processes, 1985.