

Near-Optimal Multi-Unit Auctions with Ordered Bidders

SAYAN BHATTACHARYA, Max-Planck Institute für Informatik, Saarbrücken

ELIAS KOUTSOUPIS, University of Oxford and University of Athens

JANARDHAN KULKARNI, Duke University

STEFANO LEONARDI, Sapienza University of Rome

TIM ROUGHGARDEN, Stanford University

XIAOMING XU, Duke University

We construct prior-free auctions with constant-factor approximation guarantees with ordered bidders, in both unlimited and limited supply settings. We compare the expected revenue of our auctions on a bid vector to the monotone price benchmark, the maximum revenue that can be obtained from a bid vector using supply-respecting prices that are nonincreasing in the bidder ordering and bounded above by the second-highest bid. As a consequence, our auctions are simultaneously near-optimal in a wide range of Bayesian multi-unit environments.

Categories and Subject Descriptors: F.0 [Theory of Computation]: General

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Prior-free Auctions

1. INTRODUCTION

The goal in prior-free auction design is to design auctions that have robust, input-by-input performance guarantees. Traditionally, auctions are evaluated using average-case or Bayesian analysis, and expected auction performance is optimized with respect to a prior distribution over inputs (i.e., bid vectors). The Bayesian versions of the problems we consider are completely solved [Myerson 1981]. Worst-case guarantees are desirable when, for example, good prior information is expensive or impossible to acquire, and when a single auction is to be re-used several times, in settings with different or not-yet-known input distributions.

Prior-free auctions were first studied in [Goldberg et al. 2006; Goldberg et al. 1999]. They focused on symmetric settings, where goods and bidders are identical, and sought auctions with expected revenue close to the fixed-price benchmark $\mathcal{F}^{(2)}$, defined as the maximum revenue that can be obtained from a given bid vector by offering every bidder a common posted price (i.e., take-it-or-leave-it offer) that is at most the second-highest bid. [Goldberg et al. 2006] showed that no auction has expected revenue more than a $\approx .42$ fraction of $\mathcal{F}^{(2)}$ for every bid vector, and constructed auctions with expected revenue at least a constant fraction of this benchmark on every input. See [Hartline and Karlin 2007] for a survey of further work in this vein.

[Hartline and Roughgarden 2008] proposed a framework for defining meaningful performance benchmarks much more generally — when bidders or feasibility constraints are

This research was supported in part by NSF grants IIS-0964560, CCF-0745761, CCF-1008065, CCF-1016885, and CCF-1215965, an ONR PECASE Award, an AFOSR MURI grant, Google Faculty Award 2012, EU FET project MULTIPLEX 317532, EU ERC project PAAI 259515.

Permission to make digital or hardcopies of part or all 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 show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credits permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

EC'13, June 16–20, 2013, Philadelphia, USA. Copyright © 2013 ACM 978-1-4503-1962-1/13/06...\$15.00

asymmetric, and for objective functions other than revenue. The first step of this framework is a “Bayesian thought experiment” — if bidders’ valuations were drawn from a prior distribution in some class, what would the optimal auction be? The second step is to characterize the collection \mathcal{C} of all optimal auctions that can arise, ranging over all permissible prior distributions. Finally, given a bid vector \mathbf{b} , the performance benchmark is defined as the maximum objective function value obtained by an auction in \mathcal{C} on the input \mathbf{b} . This framework regenerates the $\mathcal{F}^{(2)}$ benchmark (modulo the technically necessary upper bound on prices) and has been used for several other objective functions and asymmetric environments [Devanur and Hartline 2009; Hartline and Roughgarden 2008; 2009; Hartline and Yan 2011; Leonardi and Roughgarden 2012]. Every benchmark generated by this framework is automatically well motivated in the following sense: if the performance of an auction is within a constant factor of such a benchmark for every input, then in particular it is simultaneously near-optimal in every Bayesian environment with valuations drawn from one of the permissible prior distributions.¹

[Leonardi and Roughgarden 2012] studied the design and analysis of prior-free digital goods (i.e., unlimited supply) auctions with asymmetric bidders. They pointed out that the framework in [Hartline and Roughgarden 2008] can be applied successfully to non-identical bidders only if sufficient qualitative information about bidder asymmetry is publicly known. They proposed a model of *ordered bidders*. Earlier bidders are in some sense expected to have higher valuations. This information could be derived from, for example, zip codes, eBay bidding histories, credit history, previous transactions with the seller, and so on. [Leonardi and Roughgarden 2012] defined the *monotone price benchmark* $\mathcal{M}^{(2)}(\mathbf{b})$ for every bid vector \mathbf{b} as the maximum revenue obtainable via a monotone price vector — meaning prices are nonincreasing in the bidder ordering — in which every price is at most the second-highest bid.² The value of this benchmark is always at least that of the fixed-price benchmark $\mathcal{F}^{(2)}$, and can be a factor of $\Theta(\log n)$ larger, where n is the number of bidders. Essentially by construction, a digital goods auction that always has revenue at least a constant fraction of $\mathcal{M}^{(2)}$ is simultaneously near-optimal in every Bayesian environment with ordered distributions (where monopoly prices are nonincreasing in the bidder ordering), or when the valuation distribution of each bidder stochastically dominates that of the next one in the ordering (see [Leonardi and Roughgarden 2012] for details). Examples include uniform distributions with intervals $[0, h_i]$ and nonincreasing h_i ’s; exponential distributions with nondecreasing rates; Gaussian distributions with nonincreasing means; and so on. The main result in [Leonardi and Roughgarden 2012] is a prior-free digital goods auction with ordered bidders with expected revenue $\Omega(\mathcal{M}^{(2)}(\mathbf{b})/\log^* n)$ for every input \mathbf{b} , where n is the number of bidders and $\log^* n$ denotes the number of times that the \log_2 operator can be applied to n before the result drops below a fixed constant.³

1.1. Our Results

We give the first digital goods auction that is $O(1)$ -competitive with the monotone price benchmark $\mathcal{M}^{(2)}$. Our auction is simple and natural. It follows the standard approach of randomly partitioning the bidders into two groups, using one group of bidders to set prices for the other. We restrict prices to be (essentially) all powers of 2, but otherwise our prices are simply the optimal monotone ones for the first bidder group. Finally, to handle inputs

¹This weaker goal of good *prior-independent* auctions can also be studied in its own right [Devanur et al. 2011; Dhangwatnotai et al. 2010; Roughgarden et al. 2012]. See [Azar et al. 2013; Chen and Micali 2011; Lopomo et al. 2009] for other interpolations between average-case and worst-case analysis of auctions.

²This benchmark was also considered earlier, with a different motivation and application, by [Aggarwal and Hartline 2006].

³[Aggarwal and Hartline 2006] previously obtained an incomparable guarantee of $\Omega(\mathcal{M}^{(2)}(\mathbf{b})) - O(h \log \log h)$, where h is the ratio between the maximum and minimum bids.

where the monotone price benchmark derives most of its revenue from a small number of bidders, with constant probability we invoke an auction that is $O(1)$ -competitive with the fixed-price benchmark $\mathcal{F}^{(2)}$.

We extend our results to multi-unit auctions, where the number of items k can be less than the number of bidders. We consider the analog $\mathcal{M}^{(2,k)}$ of the monotone price benchmark, which maximizes only over (monotone) price vectors that sell at most k units. We prove that every auction that is $O(1)$ -competitive with the benchmark $\mathcal{M}^{(2,k)}$ implies simultaneously near-optimal for a range of Bayesian multi-unit environments — roughly, those in which the (ironed) virtual valuation functions of the bidders form a pointwise total ordering. We also give a general reduction, showing how to build a limited-supply auction that is $O(1)$ -competitive w.r.t. $\mathcal{M}^{(2,k)}$ from an unlimited-supply auction that is $O(1)$ -competitive w.r.t. $\mathcal{M}^{(2)}$.

2. PRELIMINARIES

In a *multi-unit auction*, there is one seller, n bidders, and k identical items. Each bidder wants only one good, and has a private — i.e., unknown to the seller — *valuation* v_i . We call the special case where $k = n$ *unlimited supply* or *digital goods*. We study direct-revelation auctions, in which the bidders report bids \mathbf{b} to the seller, and the seller then decides who wins a good and at what price.⁴ For a fixed (randomized) auction, we use $X_i(\mathbf{b})$ and $P_i(\mathbf{b})$ to denote the winning probability and expected payment of bidder i when the bid profile is \mathbf{b} . As in previous works on prior-free auction design, we consider only auctions that are individually rational — meaning $P_i(\mathbf{b}) \leq v_i \cdot X_i(\mathbf{b})$ for every i and \mathbf{b} — and truthful, meaning that for each bidder i and fixed bids \mathbf{b}_{-i} by the other bidders, bidder i maximizes its quasi-linear utility $v_i \cdot X_i(b_i, \mathbf{b}_{-i}) - P_i(b_i, \mathbf{b}_{-i})$ by setting $b_i = v_i$. Since we consider only truthful auctions, from now on we use bids \mathbf{b} and valuations \mathbf{v} interchangeably.

Truthful and individually rational digital goods auctions have a nice canonical form: for every bidder i there is a (possibly randomized) function $t_i(\mathbf{v}_{-i})$ that, given the valuations \mathbf{v}_{-i} of the other bidders, gives bidder i a “take-it-or-leave-it offer” at the price $t_i(\mathbf{v}_{-i})$. This means that bidder i is given a good if and only if $v_i \geq t_i(\mathbf{v}_{-i})$, in which case it is charged the price $t_i(\mathbf{v}_{-i})$. It is clear that every choice (t_1, \dots, t_n) of such functions defines a truthful, individually rational digital goods auction; conversely, every such auction is equivalent to a choice of (t_1, \dots, t_n) [Goldberg et al. 2006]. A special case of such an auction is a *price vector* p , in which each t_i is the constant function $t_i(\mathbf{v}_{-i}) = p_i$. When the supply is limited (i.e., there are $k < n$ copies of the good), truthful auctions induce functions (t_1, \dots, t_n) with the property that, on every input, at most k bidders win.

The *revenue* of an auction on the valuation profile \mathbf{v} is the sum of the payments collected from the winners. Let $v^{(2)}$ denote the second-highest valuation of a profile \mathbf{v} . The *fixed-price benchmark* $\mathcal{F}^{(2)}$ is defined, for each valuation profile \mathbf{v} , as the maximum revenue that can be obtained from a constant price vector whose price is at most $v^{(2)}$:

$$\mathcal{F}^{(2)}(\mathbf{v}) = \max_{p \leq v^{(2)}} \left(\sum_{i: v_i \geq p} p \right).$$

Now suppose there is a known ordering on the bidders, say $1, 2, \dots, n$. The *monotone-price benchmark* $\mathcal{M}^{(2)}$ is defined analogously to $\mathcal{F}^{(2)}$, except that non-constant monotone price vectors are also permitted:

⁴For the questions we ask, the “Revelation Principle” (see, e.g., Nisan [Nisan 2007]) ensures that there is no loss of generality by considering only direct-revelation auctions.

$$\mathcal{M}^{(2)}(\mathbf{v}) = \max_{v^{(2)} \geq p_1 \geq p_2 \geq \dots \geq p_n} \left(\sum_{i: v_i \geq p_i} p_i \right). \quad (1)$$

Clearly, $\mathcal{M}^{(2)}(\mathbf{v}) \geq \mathcal{F}^{(2)}(\mathbf{v})$ for every input \mathbf{v} .

The monotonicity and upper-bound constraints are enforced only in the computation of the benchmark $\mathcal{M}^{(2)}$. Auctions, while obviously not privy to the private valuations, can employ whatever prices they see fit. This is natural for prior-free auctions and also necessary for non-trivial results [Goldberg and Hartline 2003].

Finally, when we say that an auction is α -competitive with or has *approximation factor* α for a benchmark, we mean that the auction's expected revenue is at least a $1/\alpha$ fraction of the benchmark for every input \mathbf{v} .

3. THE AUCTION FOR UNLIMITED SUPPLY OF ITEMS

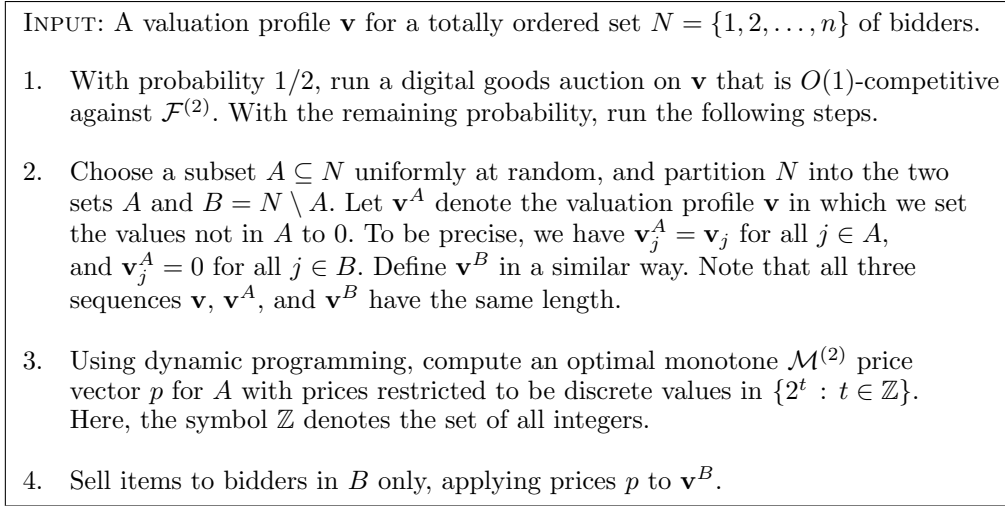


Fig. 1. The auction OPTIMAL PRICE SCALING (OPS).

Let $\text{REV}^A(p)$ denote the revenue extracted by the price vector p from the bidders in A . Similarly, define the notation $\text{REV}^B(p)$. Let $\text{REV}(p) = \text{REV}^A(p) + \text{REV}^B(p)$. Note that:

$$\text{REV}^A(p) = \sum_{j \in A: v_j \geq p_j} p_j, \quad \text{and} \quad \text{REV}^B(p) = \sum_{j \in B: v_j \geq p_j} p_j.$$

We bound the expected revenue of our auction (see Figure 1) by considering two cases.

Case 1. The ratio $\mathcal{F}^{(2)}/\mathcal{M}^{(2)}$ is at least some constant. Note that with probability $1/2$, we execute an auction which is $O(1)$ -competitive against $\mathcal{F}^{(2)}$. Hence, in this case, our revenue is clearly within a constant factor of $\mathcal{M}^{(2)}$.

Case 2. The ratio $\mathcal{F}^{(2)}/\mathcal{M}^{(2)}$ is very small. If this is the case, then we prove that the expected value of $\text{REV}^B(p)$ is within a constant factor of $\mathcal{M}^{(2)}$. Note that with probability $1/2$, we run the general scheme whose revenue is given by $\text{REV}^B(p)$. Hence, the auction's revenue remain $O(1)$ -competitive against $\mathcal{M}^{(2)}$.

We introduce the following notations and terminologies.

DEFINITION 3.1. For any integer $l \geq 0$, the l -th price level is the price q in $\{2^t : t \in \mathbb{Z}\}$ which lies in the range: $\mathcal{M}^{(2)}/2^{l+1} < q \leq \mathcal{M}^{(2)}/2^l$.

Since the prices in the set $\{2^t : t \in \mathbb{Z}\}$ are powers of 2, the l -th price level is unique. Throughout the paper, we reserve the symbol $p_{(l)}$ for the l -th price level.

DEFINITION 3.2. Fix any two bidders $i < j$, and any integer $l \geq 0$. If both the bidders' valuations are at least $p_{(l)}$, then we say that (i, j, l) is a level- l -triple.

The concept of a triple is linked with the ordering of the bidders. Thus, a bidder $k \in N$ belongs to a triple (i, j, l) iff $i \leq k \leq j$. The bidder is *winning* iff $v_k \geq p_{(l)}$.

DEFINITION 3.3. The set of winning bidders in a triple (i, j, l) is defined as:

$$W_{ijl} = \{k \in N : i \leq k \leq j \text{ and } v_k \geq p_{(l)}\}.$$

A triple is *balanced* iff its winning bidders are evenly partitioned among A and B .

DEFINITION 3.4. A triple (i, j, l) is balanced iff we have

$$\frac{1}{3} \times |W_{ijl}| \leq |A \cap W_{ijl}|, |B \cap W_{ijl}| \leq \frac{2}{3} \times |W_{ijl}|.$$

A triple is *large* if it contains sufficiently many winning bidders.

DEFINITION 3.5. A level- l -triple (i, j, l) is large iff we have $|W_{ijl}| \geq 288l$.

In Section 3.1, we show that certain important events occur with constant probability. In Section 3.2, we show that conditioned on these important events, our auction generates good revenue.

3.1. Important Events

We define the event \mathcal{E}_1 where $\text{REV}(p) \geq \mathcal{M}^{(2)}/6$. Next, we define the event $\mathcal{E}_2(l)$ where every large level- l -triple is balanced. Further, we define the event \mathcal{E}_2 as follows.

$$\mathcal{E}_2 = \bigcap_{l \geq 24} \mathcal{E}_2(l) \tag{2}$$

We show that the events \mathcal{E}_1 and \mathcal{E}_2 occur simultaneously with constant probability.

LEMMA 3.1. We have: $\Pr[\mathcal{E}_1] \geq 1/16$.

PROOF. Let OPT^A denote the maximum revenue of any monotone $\mathcal{M}^{(2)}$ price vector from the bidders in A . Leonardi et al. proved that $\text{OPT}^A \geq \mathcal{M}^{(2)}/3$ with probability at least $1/16$ (see Lemma 3.2 in [Leonardi and Roughgarden 2012]). Since p is the optimal monotone $\mathcal{M}^{(2)}$ price vector for A with prices restricted to powers of 2, we get $\text{REV}^A(p) \geq \text{OPT}^A/2$. Now, the lemma follows from the observation that $\text{REV}(p) \geq \text{REV}^A(p)$. \square

CLAIM 3.1. For every integer $l \geq 0$, the number of level- l -triples is at most 2^{2l+2} .

PROOF. Consider a bidder k whose valuation v_k is at least $p_{(l)}$. Since $p_{(l)} > \mathcal{M}^{(2)}/2^{l+1}$, we infer that $v_k > \mathcal{M}^{(2)}/2^{l+1}$. Thus, there are at most 2^{l+1} such bidders. Since a level- l -triple (i, j, l) is uniquely determined by two bidders $i < j$ having valuations at least $p_{(l)}$, we infer that there can be at most $(2^{l+1})^2 = 2^{2l+2}$ level- l -triples. \square

We use the following version of the Chernoff bound.

THEOREM 3.2. *Let T_1, \dots, T_m be i.i.d random variables such that $T_i \in \{0, 1\}$ for all $i \in \{1, \dots, m\}$. Define their sum as $T = \sum_{i=1}^m T_i$, and let $\mu = E[T]$. For all $0 < \delta < 1$:*

$$Pr[(1 - \delta)\mu \leq T \leq (1 + \delta)\mu] \geq 1 - 2 \times \exp\left(-\frac{\mu\delta^2}{4}\right).$$

CLAIM 3.2. *For all $l \geq 24$, we have: $Pr[\mathcal{E}_2(l)] \geq 1 - 1/2^l$.*

PROOF. Fix any large level- l -triple (i, j, l) . By definition, the number of winning bidders in (i, j, l) is at least $288l$. Since each of these bidders is included in the set A independently and uniformly at random, Theorem 3.2 implies that the triple (i, j, l) is not balanced with probability at most $2/e^{4l}$. By Claim 3.1, there are at most 2^{2l+2} level- l -triples. Applying union bound, the probability that some level- l -triple is not balanced is at most $2^{2l+2} \times 2/e^{4l} \leq 1/2^l$, for $l \geq 24$. \square

LEMMA 3.3. *We have: $Pr[\mathcal{E}_2] \geq 31/32$.*

PROOF. Applying union-bound, we infer that

$$1 - Pr[\mathcal{E}_2] \leq \sum_{l \geq 24} (1 - Pr[\mathcal{E}_2(l)]) \leq \sum_{l \geq 24} \frac{1}{2^l} \leq \frac{1}{32}.$$

\square

THEOREM 3.4. *We have: $Pr[\mathcal{E}_1 \cap \mathcal{E}_2] \geq 1/32$.*

PROOF. Follows from applying union bound on Lemma 3.1 and Lemma 3.3. \square

3.2. Main Analysis

Let $I_l(p)$ denote the *interval* of bidders whose prices lie at the l -th level, under the price vector p . To be more specific, we define $I_l(p) = \{j \in N : p_j = p_{(l)}\}$. Since the price vector p is monotone, the bidders in the set $I_l(p)$ are contiguous to one another.

Let $W_l(p)$ denote the set of winning bidders in $I_l(p)$. To be more precise, we have $W_l(p) = \{j \in I_l(p) : v_j \geq p_{(l)}\}$. Let $\text{REV}_l(p)$ denote the contribution towards $\text{REV}(p)$ by the interval $I_l(p)$. Note that $\text{REV}_l(p) = |W_l(p)| \times p_{(l)}$, and $\text{REV}(p) = \sum_{l \geq 0} \text{REV}_l(p)$.

DEFINITION 3.6. *A interval $I_l(p)$ is good if $|W_l(p)| \geq 288l$, and bad otherwise.*

We show that the bad intervals $I_l(p)$, with $l \geq 24$, contribute relatively little revenue.

CLAIM 3.3. *We have:*

$$\sum_{l \geq 24 : I_l(p) \text{ is bad}} \text{REV}_l(p) \leq \frac{1}{18} \times \mathcal{M}^{(2)}.$$

PROOF. Fix any bad interval $I_l(p)$. Since $|W_l(p)| < 288l$ and $p_{(l)} \leq \mathcal{M}^{(2)}/2^l$, we have:

$$\text{REV}_l(p) = |W_l(p)| \times p_{(l)} < \frac{288l}{2^l} \times \mathcal{M}^{(2)}.$$

Summing over all bad intervals $I_l(p)$ with $l \geq 24$, we get:

$$\sum_{l \geq 24 : I_l(p) \text{ is bad}} \text{REV}_l(p) \leq \sum_{l \geq 24} \frac{288l}{2^l} \times \mathcal{M}^{(2)} \leq \frac{1}{18} \times \mathcal{M}^{(2)}.$$

\square

Now, we are ready to prove the revenue guarantee.

THEOREM 3.5. *The expected revenue of the auction in Figure 1 is within a constant factor of the benchmark $\mathcal{M}^{(2)}$.*

PROOF. We shall consider two mutually exclusive and exhaustive cases.

Case 1. $432 \times \mathcal{F}^{(2)} \geq \mathcal{M}^{(2)}$.

With probability $1/2$, we execute an auction that is $O(1)$ -competitive against the benchmark $\mathcal{F}^{(2)}$. Hence, the expected revenue of our auction is at least $\mathcal{F}^{(2)}/O(1)$, which in turn, is at least $\mathcal{M}^{(2)}/O(1)$.

Case 2. $432 \times \mathcal{F}^{(2)} < \mathcal{M}^{(2)}$.

Here, we claim that the first few intervals contribute little revenue.

$$\sum_{l=0}^{23} \text{REV}_l(p) \leq \mathcal{M}^{(2)}/18 \quad (3)$$

For the sake of contradiction, suppose that the above equation is not true. Then there is some interval $I_{l^*}(p)$ with $l^* \in [0, 23]$ such that:

$$\text{REV}_{l^*}(p) = |W_{l^*}(p)| \times p_{(l^*)} > \mathcal{M}^{(2)}/(18 \times 24).$$

Consider the price vector p' which offers the item at price $p_{(l^*)}$ to every bidder, so that we have $p'_j = p_{(l^*)}$ for all $j \in N$. Next, recall that p is a monotone $\mathcal{M}^{(2)}$ price vector, and note that the set $W_{l^*}(p)$ is non-empty. Thus, there should be at least two bidders in A whose valuations are at least $p_{(l^*)}$. We infer that p' is a uniform $\mathcal{F}^{(2)}$ price vector, and:

$$\mathcal{F}^{(2)} \geq \text{REV}(p') \geq \text{REV}_{l^*}(p) \geq \mathcal{M}^{(2)}/(18 \times 24).$$

This contradicts our assumption that $432 \times \mathcal{F}^{(2)} < \mathcal{M}^{(2)}$. Hence, equation (3) must hold.

For the rest of the proof, we condition on the event $\mathcal{E}_1 \cap \mathcal{E}_2$.

First, recall that under the event \mathcal{E}_1 , we have $\text{REV}(p) \geq \mathcal{M}^{(2)}/6$. Combining this with Claim 3.3 and equation 3, we see that the latter good intervals give large revenue.

$$\sum_{l \geq 24: I_l(p) \text{ is good}} \text{REV}_l(p) \geq \left(\frac{1}{18}\right) \times \mathcal{M}^{(2)} \quad (4)$$

Fix any good interval $I_l(p)$ with $l \geq 24$. Let the first (resp. last) bidder in $W_l(p)$ be denoted by i (resp. j), that is, for all $k \in W_l(p)$, we have $i \leq k \leq j$. Since $p_i = p_j = p_{(l)}$ and $v_i, v_j \geq p_{(l)}$, we infer that (i, j, l) is a level- l -triple. The number of winning bidders in this triple is $|W_l(p)| \geq 288l$. We conclude that the triple (i, j, l) is large. Since we condition on the event $\mathcal{E}_1 \cap \mathcal{E}_2$ and $l \geq 24$, it follows that the triple (i, j, l) is balanced. The bidders in $W_l(p)$ are evenly partitioned among the sets A and B , so that we have:

$$|W_l(p) \cap B| \geq \left(\frac{1}{3}\right) \times |W_l(p)|.$$

Thus, the revenue from the bidders in $I_l(p) \cap B$ is at least $(1/3) \times \text{REV}_l(p)$. Summing over all good intervals $I_l(p)$ with $l \geq 24$, and applying equation (4), we get:

$$\text{REV}^B(p) \geq \sum_{l \geq 24: I_l(p) \text{ is good}} \left(\frac{1}{3}\right) \times \text{REV}_l(p) \geq \left(\frac{1}{54}\right) \times \mathcal{M}^{(2)} \quad (5)$$

To summarize, we recall that with probability $1/2$, the expected revenue of our auction is exactly $\text{REV}^B(p)$. Under this scenario, the event $\mathcal{E}_1 \cap \mathcal{E}_2$ occurs with probability at least $1/32$ (see Theorem 3.4), and conditioned on this event, we have $\text{REV}^B(p) \geq \mathcal{M}^{(2)}/54$ (see equation 5). Putting all these observations together, we find that the expected revenue of our auction is at least

$$\frac{1}{2} \times \frac{1}{32} \times \frac{\mathcal{M}^{(2)}}{54} = \frac{\mathcal{M}^{(2)}}{O(1)}.$$

□

4. MULTI-UNIT AUCTIONS

In this section we extend our results to multi-unit auctions with limited supply. To develop this theory, we extend the monotone price benchmark $\mathcal{M}^{(2)}$ to the case of an arbitrary number $k \geq 2$ of units for sale. We call a price vector p *feasible* for the valuation profile \mathbf{v} and supply limit k if: (i) $p_1 \geq p_2 \geq \dots \geq p_n$; (ii) all prices are at most the second-highest valuation of \mathbf{v} ; and (iii) there are at most k bidders i with $v_i > p_i$. We allow our benchmark to break ties in an optimal way. More precisely, the revenue earned by a feasible price vector is $\sum_{i: v_i > p_i} p_i$ plus, if there are ℓ items remaining, the sum of the prices offered to up to ℓ bidders i with $v_i = p_i$. We define the *k-unit monotone price benchmark* $\mathcal{M}^{(2,k)}(\mathbf{v})$ as the maximum revenue obtained by a price vector that is feasible for \mathbf{v} and k .

There are two main issues to address. The first issue is to identify a class of priors \mathcal{F}_i such that $\mathcal{M}^{(2,k)}(\mathbf{v})$ is a meaningful benchmark for prior-free approximation, i.e., it simultaneously approximates all optimal auctions in multi-unit Bayesian settings with priors drawn from the class. The challenge, relative to the unlimited-supply setting introduced in [Leonardi and Roughgarden 2012], is that limited-supply Bayesian optimal auctions exhibit more complex behavior than unlimited-supply ones. Section 4.1 shows, essentially, that the benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$ is meaningful for any valuation distributions that have pointwise ordered ironed virtual valuations. The second issue is to design auctions competitive with the benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$. We accomplish this through a general reduction, showing how to build a limited-supply auction that is $O(1)$ -competitive w.r.t. $\mathcal{M}^{(2,k)}(\mathbf{v})$ from a digital goods auction that is $O(1)$ -competitive w.r.t. $\mathcal{M}^{(2)}$.

4.1. Justifying the k-Unit Monotone Price Benchmark

The goal of this section is to prove that every prior-free auction that is $O(1)$ -competitive with the benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$ has expected revenue at least a constant fraction of optimal in every Bayesian multi-unit environment with valuation distributions lying in a prescribed class. Making this precise requires some terminology and facts from the theory of Bayesian optimal auction design, as developed in [Myerson 1981]. See also the exposition in [Hartline 2012].

Consider a bidder with valuation drawn from a prior distribution \mathcal{F} with positive and continuous density f on some interval. The *virtual value* v at a point v in the support is defined as

$$\phi(v) = v - \frac{1 - \mathcal{F}(v)}{f(v)}.$$

For example, if \mathcal{F} is the uniform distribution on $[0, a]$, then the corresponding virtual valuation function is $\phi(v) = 2v - a$.

For clarity, we first discuss the case of *regular* distributions, meaning distributions with nondecreasing virtual valuation functions. In this case, the Bayesian optimal auction awards items to the (at most k) bidders with the highest positive virtual valuations. The payment of a winning bidder is the minimum bid at which it would continue to win (keeping others' bids

the same). That is, if the $(k+1)$ th highest virtual valuation is z , then every winning bidder i pays $\phi_i^{-1}(\max\{0, z\})$. For these prices to be related to the monotone price benchmark, we need to impose conditions on the $\phi_i^{-1}(z)$'s. This contrasts with unlimited-supply settings, where restricting the $\phi_i^{-1}(0)$'s — that is, the monopoly reserve prices — to be nonincreasing in i is enough to justify the monotone-price benchmark [Leonardi and Roughgarden 2012]. Since the $(k+1)$ th highest virtual valuation could be anything, the natural extension of the condition in [Leonardi and Roughgarden 2012] is to restrict $\phi_i^{-1}(z)$ to be nonincreasing in i for every non-negative number z .

Accommodating irregular distributions, for which the optimal Bayesian auction is more complicated, presents additional complications. Each virtual valuation function ϕ_i is replaced by the “nearest nondecreasing approximation”, called the *ironed virtual valuation function* $\bar{\phi}_i$. The optimal auction awards the items to the (at most k) bidders with the highest positive ironed virtual valuations. Since ironed virtual valuation functions typically have non-trivial constant regions, ties can occur, and we assume that ties are broken randomly. That is, if there are k items, a group S of bidders with identical ironed virtual values $z > 0$, $\ell < k$ bidders with ironed virtual value greater than z , and $\ell + |S| > k$, then $k - \ell$ winners from S are chosen uniformly at random.

We call valuation distributions $\mathcal{F}_1, \dots, \mathcal{F}_n$ *pointwise ordered* if $\bar{\phi}_i^{-1}(z)$ is nonincreasing in i for every non-negative z .⁵ The motivating parametric examples discussed earlier — uniform distributions with intervals $[0, h_i]$ and nonincreasing h_i 's, exponential distributions with nondecreasing rates, and Gaussian distributions with nonincreasing means — are pointwise ordered in this sense.

We also require a second condition, which we inherit from the standard i.i.d. unlimited-supply setting. The issue is that, with arbitrary irregular distributions, no prior-free auction can be simultaneously near-optimal in all Bayesian environments, even with i.i.d. bidders and unlimited supply.⁶ Various mild conditions are sufficient to rule out this problem; see [Hartline and Roughgarden 2008] for a discussion. Here, for simplicity, we restrict attention to *well-behaved* Bayesian multi-unit environments, meaning that the Bayesian optimal auction derives at most a constant fraction (90%, say) of its revenue from outcomes in which some winner is charged a price higher than the second-highest valuation. (Such a winner is necessarily the bidder with the highest valuation.) Standard distributions always yield well-behaved environments. Even pathological distributions produce well-behaved environments provided the market is sufficiently large (e.g., there are enough bidders drawn i.i.d. from each of the distributions).

Our main result in this section is that approximating the k -unit monotone price benchmark guarantees simultaneous approximation of the optimal auction in all well-behaved Bayesian multi-unit environments with pointwise ordered distributions.

THEOREM 4.1. *If the expected revenue of the multi-unit auction \mathcal{A} is at least a constant fraction of $\mathcal{M}^{(2,k)}(\mathbf{v})$ on every input, then, in every well-behaved multi-unit Bayesian environment with pointwise ordered distributions, the expected revenue of \mathcal{A} is at least a constant fraction of that of the optimal auction for the environment.*

PROOF. Fix an auction that is β -competitive with $\mathcal{M}^{(2,k)}(\mathbf{v})$ on every input. Fix a well-behaved Bayesian multi-unit environment with pointwise ordered valuation distribu-

⁵Since $\bar{\phi}_i$ is continuous and nondecreasing, $\bar{\phi}_i^{-1}(z)$ is an interval. If the inverse image has multiple points, we define $\bar{\phi}_i^{-1}(z)$ by the infimum. If the inverse image is empty, we define $\bar{\phi}_i^{-1}(z)$ as the left or right endpoint of the distribution's support, as appropriate.

⁶Informally, consider valuation distributions that take on only two values, one very large (say M) and the other 0. Suppose the probability of having a very large valuation is very small (say $1/n^2$). If the distribution is known, the optimal auction uses a reserve price of M for each bidder. Elementary arguments, as in [Hartline and Roughgarden 2008], show that no single auction is near-optimal for all values of M .

tions $\mathcal{F}_1, \dots, \mathcal{F}_n$. Let \mathcal{A}^* be the optimal auction for this environment. We claim that, for every input \mathbf{v} in which the revenue collected by \mathcal{A}^* from the bidder with the highest valuation is at most the second-highest valuation, the benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$ is at least half the expected revenue of \mathcal{A}^* on \mathbf{v} . This implies that the expected revenue of \mathcal{A} is at least $1/2\beta$ times that of \mathcal{A}^* on this input. Since the environment is well behaved, the theorem follows.

To prove the claim, fix an input \mathbf{v} , as above. Recall that \mathcal{A}^* , as a Bayesian optimal auction, awards items to the (at most k) bidders with the highest positive ironed virtual valuations, breaking ties randomly. The tricky case of the proof is when ties occur. Assume there are k items, a group S of bidders with common ironed virtual value $z > 0$, and a group T of $\ell \in (k - |S|, k)$ bidders with ironed virtual value greater than ℓ (so $|S| > k - \ell$). We next explicitly compute the payments collected by \mathcal{A}^* on this input, using the standard payment formula for incentive-compatible mechanisms (see [Myerson 1981] or [Hartline 2012]). Let a_i and b_i denote the left and right endpoints, respectively, of the interval of values v that satisfy $\phi_i(v) = z$. Since the distributions are pointwise ordered, the a_i 's and the b_i 's are nonincreasing in i . Let $q = (k - \ell)/|S|$ denote the winning probability of a bidder in S . Define $q' = (k - \ell + 1)/(|S| + 1)$ as the hypothetical winning probability of a bidder in T if it lowered its bid to the value $\phi_i^{-1}(z)$. The expected payment of a bidder i in S is qa_i (i.e., a_i in the event that it wins). The expected payment of a bidder i in T (who wins with certainty) is $q'a_i + (1 - q')b_i$. To complete the proof, we argue that $\mathcal{M}^{(2,k)}(\mathbf{v})$ is at least the expected revenue collected by \mathcal{A}^* from the bidders in S , and also at least that from the bidders in T .

Projecting onto a subset of bidders only decreases the value of the k -unit monotone price benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$ (see Lemma A.1 for the formal argument). First, project onto the k bidders of S with the highest a_i values. Consider charging each such bidder the price a_i . This is a monotone price vector. By our assumption on the input \mathbf{v} , all of these prices are at most the second-highest valuation in \mathbf{v} . By the definitions, $v_i \geq a_i$ for every bidder $i \in S$ so every offer will be accepted. The resulting revenue is at least the expected revenue earned by \mathcal{A}^* on \mathbf{v} , and the value of the monotone price benchmark can only be higher. This shows that $\mathcal{M}^{(2,k)}(\mathbf{v})$ is at least the expected revenue collected by \mathcal{A}^* from bidders in S .

Similarly, project onto the (at most k) bidders of T , and consider charging each such bidder i the price $q'a_i + (1 - q')b_i$. Again, this is a monotone price vector with all prices bounded above by the second-highest valuation of \mathbf{v} , and every offer will be accepted. The value of the monotone price benchmark can only be larger, so $\mathcal{M}^{(2,k)}(\mathbf{v})$ is also at least the expected revenue collected by \mathcal{A}^* from bidders in T . The proof is complete. \square

4.2. Reduction from Limited to Unlimited Supply

Having justified the k -unit monotone price benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$, we turn to designing auctions that approximate it well. We show that competing with this benchmark reduces to competing with the benchmark $\mathcal{M}^{(2)}$ in unlimited-supply settings. The reduction from limited to unlimited supply for ordered bidders was given in [Aggarwal and Hartline 2006] for knapsack auction. This reduction is also a generalization of the one of [Goldberg et al. 2006] for identical bidders. The idea is to first identify the k “most valuable” bidders, and then run an unlimited-supply auction on them. Observe that the most valuable bidders with an ordering are not necessarily those with the highest valuations. For example, a high-valuation bidder late in the ordering need not be valuable, because extracting high revenue from it might necessitate excluding many moderate-valuation bidders earlier in the ordering. We report the “black-box reduction” of [Aggarwal and Hartline 2006], in Figure 2.

THEOREM 4.2. *If \mathcal{A} is a truthful unlimited-supply auction with ordered bidders that is β -competitive with $\mathcal{M}^{(2)}$, then the BLACK-BOX REDUCTION (BBR) auction is a truthful limited-supply auction with ordered bidders that is 2β -competitive with $\mathcal{M}^{(2,k)}(\mathbf{v})$.*

Input: A valuation profile \mathbf{v} for a totally ordered set $N = \{1, 2, \dots, n\}$ of bidders and k identical items. A truthful digital goods (unlimited supply) auction \mathcal{A} for ordered bidders.

- (1) Let p^* achieve the optimum monotone price benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$ for \mathbf{v} and k . Let $S = \{i \in N : v_i \geq p_i^*\}$ be the set of winners under p^* .
- (2) Run the unlimited supply auction \mathcal{A} on the bidders S , with the induced bidder ordering.
- (3) Charge suitable prices so that truthful reporting is a dominant strategy for every bidder.

Fig. 2. The auction BLACK-BOX REDUCTION (BBR).

PROOF. As pointed out in [Aggarwal and Hartline 2006] the mechanism produces a feasible outcome since the BLACK-BOX REDUCTION (BBR) auction has at most k winners. Also observe that the first step can be implemented efficiently using dynamic programming, so if \mathcal{A} runs in polynomial time, then so does the BLACK-BOX REDUCTION (BBR) auction. [Aggarwal and Hartline 2006] also shows that the composition of the two mechanisms yields a truthful mechanisms. Crucial to this conclusion is that the set S of winners is unchanged whenever the bid of a winner is increased.

We finally prove the performance guarantee by arguing the following two statements: (i) the unlimited supply benchmark $\mathcal{M}^{(2)}$ applied to S is at least half of the limited-supply benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$ applied to the original bidder set; and (ii) the revenue of BLACK-BOX REDUCTION (BBR) on the original bidder set is at least that of the unlimited-supply auction \mathcal{A} with the bidders S . The second statement follows immediately from the facts that the winners of BLACK-BOX REDUCTION (BBR) are the same as those of \mathcal{A} , and that the winners' payments are only higher. For statement (i), consider prices p^* that determine the benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$. The projection p_S^* of this price vector onto the set S of bidders has revenue exactly $\mathcal{M}^{(2,k)}(\mathbf{v})$. If p_S^* is feasible, then it certifies that the benchmark $\mathcal{M}^{(2)}$ is at least $\mathcal{M}^{(2,k)}(\mathbf{v})$. The only issue is if the second-highest bidder is excluded from S , in which case p_S^* might use a price larger than the second-highest valuation in S (which is not permitted by the benchmark $\mathcal{M}^{(2)}$). But such a price can only extract revenue from the bidder with the highest valuation, and every price of p^* is at most the second-highest valuation $v^{(2)}$ of the original bidders. Thus, we can restore feasibility to p_S^* by lowering at most one price to the second-highest valuation of S , and we lose revenue at most $v^{(2)}$. Since $\mathcal{M}^{(2,k)}(\mathbf{v}) \geq 2v^{(2)}$ — consider the price vector that offers $v^{(2)}$ to everybody — we retain at least half the revenue of p_S^* . Statement (i) and the theorem follow. \square

Of course, we can use the OPTIMAL PRICE SCALING (OPS) auction from Section 3 in Theorem 4.2 to obtain a truthful limited-supply auction that is $O(1)$ -competitive with the benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$. Theorem 4.1 implies that the resulting auction also enjoys a strong simultaneous approximation guarantee in Bayesian environments.

REFERENCES

- AGGARWAL, G. AND HARTLINE, J. D. 2006. Knapsack auctions. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1083–1092.
- AZAR, P., DASKALAKIS, C., MICALI, S., AND WEINBERG, S. M. 2013. Optimal and efficient parametric auctions. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
- CHEN, J. AND MICALI, S. 2011. Mechanism design with set-theoretic beliefs. In *Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 87–96.
- DEVANUR, N. AND HARTLINE, J. D. 2009. Limited and online supply and the Bayesian foundations of prior-free mechanism design. In *Proceedings of the 10th ACM Conference on Electronic Commerce (EC)*. 41–50.

- DEVANUR, N., HARTLINE, J. D., KARLIN, A. R., AND NGUYEN, T. 2011. A prior-independent mechanism for profit maximization in unit-demand combinatorial auctions. In *Proceedings of 7th Workshop on Internet & Network Economics*. To appear.
- DHANGWATNOTAI, P., ROUGHGARDEN, T., AND YAN, Q. 2010. Revenue maximization with a single sample. In *Proceedings of the 11th ACM Conference on Electronic Commerce (EC)*. 129–138.
- GOLDBERG, A. AND HARTLINE, J. 2003. Envy-free auctions for digital goods. In *Proceedings of the 4th ACM Conference on Electronic Commerce (EC)*. 29–35.
- GOLDBERG, A. V., HARTLINE, J. D., KARLIN, A., SAKS, M., AND WRIGHT, A. 2006. Competitive auctions. *Games and Economic Behavior* 55, 2, 242–269.
- GOLDBERG, A. V., HARTLINE, J. D., AND WRIGHT, A. 1999. Competitive auctions and digital goods. Tech. Rep. STAR-TR-99.09.01, STAR Laboratory, InterTrust Tech. Corp., Santa Clara, CA.
- HARTLINE, J. D. 2012. Approximation in economic design. Book in preparation.
- HARTLINE, J. D. AND KARLIN, A. 2007. Profit maximization in mechanism design. In *Algorithmic Game Theory*, N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani, Eds. Cambridge University Press, Chapter 13, 331–362.
- HARTLINE, J. D. AND ROUGHGARDEN, T. 2008. Optimal mechanism design and money burning. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC)*. 75–84. Full version last revised August, 2012.
- HARTLINE, J. D. AND ROUGHGARDEN, T. 2009. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM Conference on Electronic Commerce (EC)*. 225–234.
- HARTLINE, J. D. AND YAN, Q. 2011. Envy, truth, and optimality. In *Proceedings of the 12th ACM Conference on Electronic Commerce (EC)*. 243–252.
- LEONARDI, S. AND ROUGHGARDEN, T. 2012. Prior-free auctions with ordered bidders. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC)*. 427–434.
- LOPOMO, G., RIGOTTI, L., AND SHANNON, C. 2009. Uncertainty in mechanism design. Submitted.
- MYERSON, R. 1981. Optimal auction design. *Mathematics of Operations Research* 6, 1, 58–73.
- NISAN, N. 2007. Introduction to mechanism design (for computer scientists). In *Algorithmic Game Theory*, N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, Eds. Cambridge University Press, Chapter 9, 209–241.
- ROUGHGARDEN, T., TALGAM-COHEN, I., AND YAN, Q. 2012. Supply-limiting mechanisms. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*. 844–861.

A. MISSING PROOFS

LEMMA A.1. *For every valuation profile \mathbf{v} , $k \geq 2$, and subset S of the bidders with induced profile \mathbf{v}^S , $\mathcal{M}^{(2,k)}(\mathbf{v}) \geq \mathcal{M}^{(2)}(k, \mathbf{v}^S)$.*

PROOF. (Sketch.) Fix an input \mathbf{v} , with monotone prices p^* determining $\mathcal{M}^{(2,k)}(\mathbf{v})$. By induction, we only need to show that adding a single new bidder i can only increase the value of the benchmark. Start by offering i the same price q as its predecessor in the ordering (or the second-highest valuation, if there is no predecessor). If i rejects (i.e., $v_i < q$), this extended price vector is feasible and we are done (the optimal feasible price vector is only better). If i accepts (i.e., $v_i \geq q$) then the price vector is infeasible (with $k + 1$ winners) and we argue as follows. Go through the bidders after i one by one, increasing the offer price to q . This preserves monotonicity. If a previously winning bidder ever rejects this higher offer price, we are done (feasibility is restored and the overall revenue is higher). If not, there is now a “suffix” of bidders with the common offer price q . (This case only occurs if i is after all of the winners in p^* .) We now increase their common offer price until it equals that of the previous bidder, thereby increasing the number of bidders in the suffix. Eventually a bidder that was winning under p^* will reject the new offer price (otherwise it would contradict the optimality of p^*), leaving us with a feasible monotone price vector with revenue at least that of the original one. \square