

Optimal Competitive Auctions

Ning Chen

Nanyang Technological University
Singapore
ningc@ntu.edu.sg

Nick Gravin

Microsoft Research
ngravin@microsoft.com

Pinyan Lu

Microsoft Research
pinyanl@microsoft.com

ABSTRACT

We study the design of truthful auctions for selling identical items in unlimited supply (e.g., digital goods) to n unit demand buyers. This classic problem stands out from profit-maximizing auction design literature as it requires no probabilistic assumptions on buyers' valuations and employs the framework of competitive analysis. Our objective is to optimize the worst-case performance of an auction, measured by the ratio between a given benchmark and revenue generated by the auction.

We establish a sufficient and necessary condition that characterizes competitive ratios for *all* monotone benchmarks. The characterization identifies the worst-case distribution of instances and reveals intrinsic relations between competitive ratios and benchmarks in the competitive analysis. With the characterization at hand, we show optimal competitive auctions for two natural benchmarks.

The most well-studied benchmark $\mathcal{F}^{(2)}(\cdot)$ measures the envy-free optimal revenue where at least two buyers win. Goldberg et al. [13] showed a sequence of lower bounds on the competitive ratio for each number of buyers n . They conjectured that all these bounds are tight. We show that optimal competitive auctions match these bounds. Thus, we confirm the conjecture and settle a central open problem in the design of digital goods auctions. As one more application we examine another economically meaningful benchmark, which measures the optimal revenue across all limited-supply Vickrey auctions. We identify the optimal competitive ratios to be $(\frac{n}{n-1})^{n-1} - 1$ for each number of buyers n , that is $e - 1$ as n approaches infinity.

Categories and Subject Descriptors

F.0 [Theory of Computation]: General; J.4 [Social and Behavioral Sciences]: Economics

General Terms

Theory, Economics

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. Request permissions from Permissions@acm.org. *STOC'14*, May 31 – June 03, 2014, New York, NY, USA. Copyright 2014 ACM 978-1-4503-2710-7/14/05 ...\$15.00 <http://dx.doi.org/10.1145/2591796.2591855>.

Keywords

Auction, mechanism design, digital goods, competitive analysis

1. INTRODUCTION

A central question in auction theory is to design optimal truthful auctions that maximize the revenue of the auctioneer. A truthful auction must be incentive compatible with the selfish behavior of bidders and encourage them to reveal their private information truthfully. In the classic single-parameter setting, the auctioneer sells multiple copies of an item to n unit-demand bidders, each bidder having a private value v_i for the item. The revenue maximization problem in economics is traditionally analyzed in the Bayesian framework. The seminal Myerson's auction [24] provides an optimal design that extracts the maximum expected revenue for a given distribution of values.

While Myerson's auction gives an optimal design in the Bayesian framework, in many scenarios determining or estimating the prior distribution in advance is impossible. Without prior information on the distribution, how should an auction be designed? The auctioneer not being certain about the prior distribution is likely to resort to truthful auctions that generate a good revenue on every possible instance of bidders' valuations. We therefore employ the worst-case competitive analysis, which characterizes the worst-case performance of an auction across all possible bidders' valuations. We follow the framework summarized in [14] where a truthful auction is required to generate a profit comparable to the value of a certain economically meaningful benchmark on every input of values. The motivation of this framework comes from the analysis of on-line algorithms, where the performance of an on-line algorithm, which is unaware of the future, is measured in terms of the performance of the optimal off-line algorithm that knows the future. The assumption that an on-line algorithm does not know the future in advance corresponds to the assumption that an auction does not know the bidders' valuations in advance. How well an auction can perform in the worst-case? We will answer this question in the present paper.

A particularly interesting single-parameter setting is *digital goods auctions* in which the number of units for sale is unlimited or is greater than or equal to the number of bidders. Digital goods auctions are motivated by the applications of selling digital goods like downloadable software on the Internet or pay-per-view television where there is a negligible cost for producing a copy of the item. The design of the optimal Bayesian auction for digital goods is trivial: given a prior

distribution of values, one may independently offer each bidder a fixed price p_i tailored to the value's distribution v_i so as to maximize the expected revenue $p_i \cdot \Pr[v_i \geq p_i]$. From the worst-case perspective, digital goods auctions attracted considerable attention over the last decade [15, 10, 12, 13, 19, 9, 2, 21]. However, the optimal design was still unknown. The question about designing optimal digital goods auctions remained widely open and is fundamental to our understanding of the optimal competitive auction design in the general single-parameter setting.

Formally, in the competitive analysis framework, given a benchmark function $f(\cdot)$, the *competitive ratio* (with respect to $f(\cdot)$) of a truthful auction \mathcal{A} is defined as $\max_{\mathbf{v}} \frac{f(\mathbf{v})}{\mathbf{E}[\mathcal{A}(\mathbf{v})]}$, where $\mathbf{E}[\mathcal{A}(\mathbf{v})]$ is the expected revenue of the auction \mathcal{A} on valuation vector $\mathbf{v} = (v_1, \dots, v_n)$. The objective is to design an auction that minimizes the competitive ratio with respect to a given benchmark.

We note that a number of functions can serve as meaningful benchmarks and that they may have different optimal competitive ratios. Such flexibility in the choice of a target benchmark function may prove to be helpful for modeling different objectives of the auctioneer. Conversely, it may allow us to make meaningful conclusions and prescriptions about target benchmarks. However, most of the work done along the line of competitive analysis in algorithmic game theory [25] is devoted to the competitive analysis of a well-motivated but fixed concrete benchmark function. Namely, while there are basic guidelines for choosing a good benchmark function (e.g., the benchmark should have a strong economic motivation and should match the performance of the optimal auction as closely as possible), there are no formal criteria for distinguishing different benchmarks. In this paper, instead of justifying what a good benchmark is, we focus on the intrinsic relation between benchmarks and their corresponding competitive ratios. We give a complete characterization for almost *all* possible benchmark functions.

Theorem 1. (Characterization) For any non-negative and monotonically increasing function $f(\cdot)$, there is a truthful digital goods auction that achieves a competitive ratio of λ with respect to $f(\cdot)$ if and only if

$$\int_S f(\mathbf{v}) \cdot w(\mathbf{v}) \, d\mathbf{v} \leq \lambda \cdot \sum_{i=1}^n \int_{S_{\downarrow_i}} w(\mathbf{v}_{-i}) \, d\mathbf{v}_{-i}, \quad (1)$$

where S is any upward closed set in the support of valuation vectors, S_{\downarrow_i} is the projection of S along the i -th coordinate, $w(\mathbf{v}) = \prod_i 1/v_i^2$, and $\mathbf{v}_{-i} = (v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_n)$.

The fact that the weight function $w(\mathbf{v}) = \prod_i 1/v_i^2$ appears in the above theorem is not a coincidence. The corresponding single-parameter distribution with the density function $w(v_i) = 1/v_i^2$ for $v_i \in [1, \infty)$ is a common tool to provide bounds on the performance of auctions. It is called *equal-revenue distribution*, as it enjoys a remarkable property that if the value of a bidder i is drawn from this distribution, then any fixed price p_i offered to i generates the same expected revenue $p_i \cdot \Pr[v_i \geq p_i] = 1$. In particular, the product $w(\mathbf{v})$ of independent equal-revenue distributions $w(v_i)$ was used in [13] to obtain the best known lower bounds on the competitive ratios of digital goods auctions; it has also been used in the auction analysis of other models, e.g., in [17].

As we know from Yao's minimax principle, for any benchmark function $f(\cdot)$, there is a distribution of instances such

that on average for that distribution no truthful auction can beat the worst-case competitive ratio with respect to $f(\cdot)$. The inequality (1) in the theorem states that the equal revenue distribution is indeed the worst-case distribution to estimate the competitive ratio of an auction. In other words, in the context of digital goods auctions, the worst-case distribution can be described only by its support, and the actual density of the distribution is given by the equal-revenue distribution. The theorem gives a sufficient and necessary condition for a benchmark $f(\cdot)$ to admit a competitive ratio of λ . It implies that the optimal competitive ratio with respect to $f(\cdot)$ is the smallest value of λ for which the inequality (1) holds. The theorem indicates that all benchmarks and their competitive ratios are tied to the equal revenue distribution: it is the worst distribution not only for the benchmark considered in [13] but also for *all* monotone benchmarks.

We note that our characterization is provided by the set of inequalities (1) that only involves a function $f(\cdot)$ and a ratio λ , but does not describe an actual auction. This is similar in spirit to the characterization of truthfully implementable allocation functions, where an allocation function is *truthfully implementable* if there exists a payment function that makes the allocation function truthful. To characterize a truthfully implementable allocation function, one uses a set of inequalities such as weak-monotonicity [26, 6, 3] that only specifies an allocation function but says nothing about payments. Hence, without describing any payment function, one can determine whether there is a truthful auction with a specified allocation. Our characterization of benchmarks with a given competitive ratio shares a similar philosophy: the condition determines whether there exists a truthful auction with a certain competitive ratio with respect to $f(\cdot)$, but does not explicitly describe the auction.

The characterization theorem provides us with a powerful tool to analyze the optimal competitive ratios of auctions for different benchmarks. We first consider the most well-studied benchmark $\mathcal{F}^{(2)}$ introduced in [15], which is defined as the maximum revenue achieved in an envy-free allocation provided that at least two bidders receive the item. We next study another natural benchmark, denoted by MAXV , which is defined as the maximal revenue of the k -item Vickrey auction across all possible values of k . We have the following results on the optimal competitive ratios for these two benchmarks.

Theorem 2. There are truthful digital goods auctions that achieve the optimal competitive ratios of λ_n and γ_n for any $n \geq 2$ with respect to the benchmarks $\mathcal{F}^{(2)}$ and MAXV , respectively, where

$$\lambda_n = 1 - \sum_{i=2}^n \left(\frac{-1}{n}\right)^{i-1} \frac{i}{i-1} \binom{n-1}{i-1}$$

and

$$\gamma_n = \left(\frac{n}{n-1}\right)^{n-1} - 1.$$

It was shown in [13] that λ_n gives a lower bound on the competitive ratio of any auction with respect to $\mathcal{F}^{(2)}$ for any $n \geq 2$. These lower bounds were obtained by calculating the expected value of the benchmark $\mathcal{F}^{(2)}(\mathbf{v})$ when \mathbf{v} is drawn from the equal revenue distribution. In particular, $\lambda_2 = 2$, $\lambda_3 = 13/6$, and in general, $\{\lambda_n\}_{n=2}^{\infty}$ is an increasing sequence with a limit of roughly 2.42. Goldberg et al. [13]

conjectured that the lower bounds given by $\{\lambda_n\}$ are tight. Indeed, for $n = 2$ bidders, the second price auction gives a matching competitive ratio of λ_2 ; for $n = 3$ bidders, [19] gave a sophisticated auction with a competitive ratio that matches the lower bound of λ_3 . These are the only cases for which optimal competitive auctions were known. Our result confirms the conjecture of [13] and settles the long standing open problem of designing optimal digital goods auctions with respect to the benchmark $\mathcal{F}^{(2)}$.

For the benchmark MAXV , we calculate the expected value of $\text{MAXV}(\mathbf{v})$ when \mathbf{v} is drawn from the equal revenue distribution, and derive a sequence of optimal competitive ratios γ_n for each number n of bidders in the auction. We note that $\gamma_2 = 1$, $\gamma_3 = 5/4$, and $\{\gamma_n\}_{n=1}^\infty$ is an increasing sequence with the limit of $e - 1$ as n approaches infinity.

Finally, as another application of the characterization theorem, we consider optimal competitive auctions for the multi-unit limited supply setting where there is an item with k units of supply for sale to n unit-demand bidders. It was observed in [14] that there is a competitive ratio preserving reduction from unlimited supply to limited supply. Namely, given a digital goods auction with a competitive ratio λ with respect to $\mathcal{F}^{(2)}$ for k bidders, one can construct a truthful auction for n bidders with the same competitive ratio λ with respect to $\mathcal{F}^{(2,k)}$ (where $\mathcal{F}^{(2,k)}$ is the optimal fixed price revenue provided that at least 2 and at most k bidders receive the item). Our analysis continues to hold for those benchmarks that only depend on the k highest values, and thus, gives optimal competitive auctions in the competitive analysis framework.

1.1 Related Work

The study of competitive digital goods auctions was coined by Goldberg et al. [15], where the authors introduced the random sampling optimal price auction and showed that it has a constant competitive ratio with respect to $\mathcal{F}^{(2)}$. Later on the competitive ratio of the auction was shown to be 15 and 4.68, by Feige et al. [9] and Alaei et al. [2], respectively. Since the pioneer work of [15], a sequence of work has been devoted to design auctions with improved competitive ratios: the random sampling cost sharing auction [10] with ratio 4; the consensus revenue estimate auction [12] with ratio 3.39; the aggregation auction [19] with ratio 3.25; the best known ratio 3.12 is attained by the averaging auction [21].

Fiat et al. [10] formulated the prior-free analysis framework for digital goods auctions. The framework was further developed to general symmetric auction problems and connected with the Bayesian framework in [20]. The relation between envy-freeness and prior-free mechanism design was further investigated in [18]. Aggarwal et al. [1] showed that every randomized auction in the digital goods environment can be derandomized in polynomial time with an extra additive error that depends on the maximal range of values.

Leonardi and Roughgarden [22] introduced another benchmark, namely monotone-price benchmark $\mathcal{M}^{(2)}$. Later it was shown in [5] that digital goods auctions have a constant competitive ratio with respect to $\mathcal{M}^{(2)}$.

A number of variants of digital goods auctions have been investigated, including, e.g., online auctions [4, 7], limited supply (k -unit auctions) [8], online auctions with unknown

limited supply [23], externalities between bidders [16], and matroid permutations and position environments [18].

2. PRELIMINARIES

In a digital goods auction, an auctioneer sells multiple copies of an item in unlimited supply to n bidders. Each bidder i is interested in a single unit of the item and values it at a privately known value v_i . We consider a single-round auction, where each bidder submits a sealed bid b_i to the auctioneer. Upon receiving submitted bids $\mathbf{b} = (b_1, \dots, b_n)$ from all bidders, the auctioneer decides on whether each bidder i receives an item and the amount that i pays. If bidder i wins an item, his *utility* is the difference between his value v_i and his payment; otherwise, the bidder pays 0 and his utility is 0. The auctioneer's *revenue* is the total payment of the bidders.

We assume that all bidders are self-motivated and aim to maximize their own utility. We say that an auction is *truthful* or incentive compatible if it is a dominant strategy for every bidder i to submit his private value, i.e., $b_i = v_i$, no matter how other bidders behave. A randomized auction is (universally) truthful if it is randomly distributed over deterministic truthful auctions.

An auction is called *bid-independent* if, for each bidder i , the auctioneer computes a threshold price p_i according to the bids of the rest $n-1$ bidders $\mathbf{b}_{-i} = (b_1, \dots, b_{i-1}, ?, b_{i+1}, \dots, b_n)$. In other words, there is a function g_i such that $p_i = g_i(\mathbf{b}_{-i})$. It was shown in [15] that an auction is truthful if and only if it is bid-independent. Thus, it is sufficient to consider bid-independent auctions in order to design truthful auctions.

To evaluate the performance of an auction, we need to have a reasonable benchmark function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, where $f(\mathbf{b})$ measures our target revenue for the bid vector \mathbf{b} . Given a benchmark function $f(\cdot)$, we say that an auction \mathcal{A} has a *competitive ratio* of λ with respect to $f(\cdot)$ if

$$\frac{f(\mathbf{v})}{\mathbf{E}[\mathcal{A}(\mathbf{v})]} \leq \lambda, \quad \forall \mathbf{v} = (v_1, \dots, v_n)$$

where $\mathbf{E}[\mathcal{A}(\mathbf{b})]$ is the expected revenue of auction \mathcal{A} on the bid vector \mathbf{b} . The focus of our paper is to design truthful auctions that minimize the competitive ratio with respect to different benchmarks.

In this paper, we assume that a benchmark function $f(\cdot)$ is non-negative and monotone. These are natural conditions for a function to serve as a reasonable benchmark. Specifically, we will focus on the following benchmark functions. (Given a bid vector $\mathbf{b} = (b_1, \dots, b_n)$ we reorder bids so that $b_{(1)} \geq b_{(2)} \geq \dots \geq b_{(n)}$.)

- $\mathcal{F}^{(2)}(\mathbf{b}) = \max_{2 \leq k \leq n} k \cdot b_{(k)}$. That is, $\mathcal{F}^{(2)}$ gives the largest possible revenue obtained in a fixed price auction given that there are at least two winners. $\mathcal{F}^{(2)}$ was denoted sometimes as $\mathcal{F}^{(2)}$ in the previous literature and provides the optimal envy-free revenue conditioned on that at least two bidders receive the item.
- $\text{MAXV}(\mathbf{b}) = \max_{1 \leq k < n} k \cdot b_{(k+1)}$. We note that $k \cdot b_{(k+1)}$ is the revenue of the k -item Vickrey auction with a fixed supply of k items. Hence, MAXV gives the largest revenue obtained in the Vickrey auction for selling k items for all possible values of the limited supply k .

THEOREM 1 (GOLDBERG ET AL. [13]). *The competitive ratio with respect to $\mathcal{F}^{(2)}$ of any truthful randomized auction is at least*

$$\lambda_n = 1 - \sum_{i=2}^n \left(\frac{-1}{n}\right)^{i-1} \frac{i}{i-1} \binom{n-1}{i-1}.$$

We note that $\lambda_2 = 2$, $\lambda_3 = 13/6$, and $\{\lambda_n\}_{n=1}^\infty$ is an increasing sequence with a limit of 2.42 when n approaches infinity. In the proof of the above theorem, the authors of [13] constructed a so-called equal revenue distribution where all values v_i are drawn identically and independently with probability $\Pr[v_i > x] = \frac{1}{x}$ for any $x \geq 1$. A remarkable property of this distribution is that any truthful auction has the same expected revenue n . It was shown that the expected value of $\mathcal{F}^{(2)}(\mathbf{v})$ is $n \times \lambda_n$. Thus, the theorem follows since

$$\max_{\mathbf{v}} \frac{\mathcal{F}^{(2)}(\mathbf{v})}{\mathbf{E}[\mathcal{A}(\mathbf{v})]} \geq \frac{\mathbf{E}[\mathcal{F}^{(2)}(\mathbf{v})]}{\mathbf{E}[\mathcal{A}(\mathbf{v})]} = \frac{n \times \lambda_n}{n} = \lambda_n.$$

3. CHARACTERIZATION OF BENCHMARKS

We introduce the following central definition for a Benchmark function.

DEFINITION 1 (ATTAINABILITY). *A benchmark function $f(\cdot)$ is λ -attainable if there exists a truthful auction that has a competitive ratio of λ with respect to $f(\cdot)$.*

We shall give a sufficient and necessary condition of attainability of a benchmark in this section with a set of inequalities which only involves the function f and ratio λ but not any auction. After having the characterization, we analyze the attainability of two well-studied benchmarks in the next section.

For technical simplicity, we consider a discrete and bounded domain for all bids: for any bidder i , we assume that $b_i \in \{(1+\delta)^t \mid t = 0, 1, 2, \dots, N\}$, where $\delta > 0$ is any fixed small constant. Thus, all bids are between 1 and $(1+\delta)^N$, and are multiples of $1+\delta$. Let

$$\Lambda = \{(1+\delta)^t \mid t = 0, 1, 2, \dots, N\}$$

denote the support of a single bidder's bids and Λ^n denote the support of bid vectors of all n bidders. We note that such a multiplicative discretization is not critical for our characterization. (Alternatively, we may consider an additive discretization with bids being integer multiples of δ .) We will discuss how to generalize our analysis to continuous and unbounded domains at the end of this section.

For the domain Λ^n we assume without loss of generality that for any \mathbf{b}_{-i} , the price offered to bidder i is also from Λ , which is of the form $(1+\delta)^t$. For $\mathbf{b}_{-i} \in \Lambda^{n-1}$ and $p_i \in \Lambda$, let $z_i(\mathbf{b}_{-i}, p_i)$ denote the probability that the auctioneer offers price p_i to bidder i when observing others' bids \mathbf{b}_{-i} . There exists a truthful auction that is λ -competitive with respect to the benchmark $f(\cdot)$ if and only if the following linear system is feasible:

$$\text{LS}_1 : \begin{cases} \lambda \cdot \sum_i \sum_{p_i=1}^{b_i} p_i \cdot z_i(\mathbf{b}_{-i}, p_i) \geq f(\mathbf{b}), & \forall \mathbf{b} \\ \sum_{p_i=1}^{(1+\delta)^N} z_i(\mathbf{b}_{-i}, p_i) \leq 1, & \forall i, \mathbf{b}_{-i} \\ z_i(\mathbf{b}_{-i}, p_i) \geq 0, & \forall i, \mathbf{b}_{-i}, p_i \end{cases}$$

Note that the summation over p_i in LS_1 is taken in the domain $\Lambda = \{(1+\delta)^t \mid t = 0, 1, 2, \dots, N\}$.

We remark that the correspondence between $\{z_i(\mathbf{b}_{-i}, p_i)\}_i$ and truthful auctions is not one-to-one but one-to-many. For any auction, there is a corresponding probability profile $\{z_i(\mathbf{b}_{-i}, p_i)\}_i$ that satisfies LS_1 . On the other hand, different auctions may have the same probability profile $\{z_i(\mathbf{b}_{-i}, p_i)\}_i$ (thus, they have the same expected revenue). Note that for any given $\{z_i(\mathbf{b}_{-i}, p_i)\}_i$, we can construct at least one corresponding truthful auction, which independently offers the threshold price p_i to each bidder i with probability $z_i(\mathbf{b}_{-i}, p_i)$.

We define

$$x_i(\mathbf{b}) = x_i(\mathbf{b}_{-i}, b_i) = \sum_{p_i=1}^{b_i} p_i \cdot z_i(\mathbf{b}_{-i}, p_i).$$

Intuitively, $x_i(\mathbf{b})$ gives the expected revenue obtained from bidder i when the bid vector is \mathbf{b} . We further define

$$w((1+\delta)^t) = \begin{cases} \frac{\delta}{(1+\delta)^{t+1}}, & t = 0, 1, 2, \dots, N-1 \\ \frac{1}{(1+\delta)^N}, & t = N \end{cases}$$

We note that $w(\cdot)$ can be viewed as a equal revenue distribution over Λ , which satisfies the following nice property:

$$\sum_{t=k}^N w((1+\delta)^t) = \frac{1}{(1+\delta)^k}.$$

Let $w(\mathbf{b}) = \prod_{k=1}^n w(b_k)$ and $w(\mathbf{b}_{-i}) = \prod_{k \neq i} w(b_k)$.

Given these definitions, the aforementioned linear system LS_1 can be rewritten as follows.

$$\text{LS}_2 : \begin{cases} \lambda \cdot \sum_i x_i(\mathbf{b}_{-i}, b_i) \geq f(\mathbf{b}), & \forall \mathbf{b} \\ \sum_{t=0}^N w((1+\delta)^t) \cdot x_i(\mathbf{b}_{-i}, (1+\delta)^t) \leq 1, & \forall i, \mathbf{b}_{-i} \\ x_i(\mathbf{b}_{-i}, (1+\delta)^t) \leq x_i(\mathbf{b}_{-i}, (1+\delta)^{t+1}), & \forall i, \mathbf{b}_{-i}, t \\ x_i(\mathbf{b}_{-i}, b_i) \geq 0, & \forall \mathbf{b} \end{cases}$$

The third constraint in LS_2 requires monotonicity of $x_i(\cdot)$, which is a necessary condition for the equivalence of the two linear systems. Indeed, $z_i(\mathbf{b}_{-i}, b_i) = \frac{1}{b_i}(x_i(\mathbf{b}_{-i}, b_i) - x_i(\mathbf{b}_{-i}, \frac{b_i}{1+\delta}))$ must be non-negative, where $x_i(\mathbf{b}_{-i}, \frac{1}{1+\delta}) = 0$. To summarize, we have the following claim.

PROPOSITION 1. *A monotone function f is λ -attainable if and only if the linear system LS_2 has a feasible solution $\mathbf{x} = \{x_i(\mathbf{b})\}_{i, \mathbf{b}}$ (or equivalently, LS_1 has a feasible solution $\mathbf{z} = \{z_i(\mathbf{b})\}_{i, \mathbf{b}}$).*

For any set $S \subseteq \Lambda^n$, let $S_{\downarrow i}$ denote the projections of S along the i -th coordinate. Formally,

$$S_{\downarrow i} = \{\mathbf{b}_{-i} \mid \exists b_i \text{ s.t. } (\mathbf{b}_{-i}, b_i) \in S\}.$$

For a non-negative and monotone function f , we have the following characterization.

THEOREM 2. *Over a domain Λ , a non-negative and monotone function f is λ -attainable if and only if for any upward closed set $S \subset \Lambda^n$,*

$$\sum_{\mathbf{b} \in S} w(\mathbf{b}) \cdot f(\mathbf{b}) \leq \lambda \cdot \sum_i \sum_{\mathbf{b}_{-i} \in S_{\downarrow i}} w(\mathbf{b}_{-i}). \quad (2)$$

PROOF. *Only if (necessity).* If $f(\cdot)$ is λ -attainable, then there exists a solution $x_i(\mathbf{b}_{-i}, b_i)$ which satisfies all constraints in LS_2 . Thus,

$$\begin{aligned} \sum_{\mathbf{b} \in S} w(\mathbf{b}) \cdot f(\mathbf{b}) &\leq \lambda \cdot \sum_{\mathbf{b} \in S} w(\mathbf{b}) \sum_i x_i(\mathbf{b}_{-i}, b_i) \\ &= \lambda \cdot \sum_i \sum_{\mathbf{b}_{-i} \in S_{\downarrow i}} w(\mathbf{b}_{-i}) \sum_{b_i: \mathbf{b} \in S} w(b_i) \cdot x_i(\mathbf{b}_{-i}, b_i) \\ &\leq \lambda \cdot \sum_i \sum_{\mathbf{b}_{-i} \in S_{\downarrow i}} w(\mathbf{b}_{-i}) \sum_{b_i \in \Lambda} w(b_i) \cdot x_i(\mathbf{b}_{-i}, b_i) \\ &\leq \lambda \cdot \sum_i \sum_{\mathbf{b}_{-i} \in S_{\downarrow i}} w(\mathbf{b}_{-i}). \end{aligned}$$

If (sufficiency). Our goal is to find a feasible solution $\mathbf{x} = \{x_i(\mathbf{b})\}_{i, \mathbf{b}}$ to the above linear system LS_2 , given the system of inequalities (2). Our proof is constructive: we provide a procedure of continuously increasing $\{x_i(\mathbf{b})\}_{i, \mathbf{b}}$, starting from 0, to a point where all constraints of LS_2 are satisfied. In order to do that, we write a slightly more general linear system of the following form.

$$\text{LS}_3 : \begin{cases} \lambda \cdot \sum_i x_i(\mathbf{b}_{-i}, b_i) \geq f(\mathbf{b}), & \forall \mathbf{b} \\ \sum_{t=0}^N w((1+\delta)^t) \cdot x_i(\mathbf{b}_{-i}, (1+\delta)^t) \leq g_i(\mathbf{b}_{-i}), & \forall i, \mathbf{b}_{-i} \\ x_i(\mathbf{b}_{-i}, (1+\delta)^t) \leq x_i(\mathbf{b}_{-i}, (1+\delta)^{t+1}), & \forall i, \mathbf{b}_{-i}, t \\ x_i(\mathbf{b}_{-i}, b_i) \geq 0, & \forall \mathbf{b} \end{cases}$$

The only difference between LS_2 and LS_3 is $g_i(\mathbf{b}_{-i})$ in the second constraint. Intuitively, $g_i(\mathbf{b}_{-i})$ represents a total mass along direction i at the point \mathbf{b}_{-i} . Initially, $g_i(\mathbf{b}_{-i}) = 1$, for all i and \mathbf{b}_{-i} , and in which case, the two linear systems LS_2 and LS_3 are identical. We will however decrease the values of $g_i(\mathbf{b}_{-i})$'s in the process of the proof while maintaining the following condition, derived from (2), for all upward closed sets $S \subset \Lambda^n$.

$$\sum_{\mathbf{b} \in S} w(\mathbf{b}) \cdot f(\mathbf{b}) \leq \lambda \cdot \sum_i \sum_{\mathbf{b}_{-i} \in S_{\downarrow i}} g_i(\mathbf{b}_{-i}) w(\mathbf{b}_{-i}). \quad (3)$$

Let Γ be the collection of upward closed sets for which the above inequality (3) is tight. It turns out that Γ has a nice structure summarized in the following claim.

CLAIM 1. *If $S_1, S_2 \in \Gamma$, then $S_1 \cup S_2, S_1 \cap S_2 \in \Gamma$.*

The high-level idea of our proof is to continuously increase $\{x_i(\mathbf{b})\}_{i, \mathbf{b}}$, starting from 0, to a point where all constraints of LS_2 are satisfied. To implement the idea, we identify two special sets S_0 and S_1 , where $S_0 \supseteq S_1$, and a special set $T_i \subseteq S_{0\downarrow i} \setminus S_{1\downarrow i}$ for a specific coordinate i . Then for each $\mathbf{b}_{-i} \in T_i$, we find a threshold point $c_i(\mathbf{b}_{-i})$ and increase $x_i(\mathbf{b}_{-i}, b)$ by ε for all $b \geq c_i(\mathbf{b}_{-i})$. (In the proof, $c_i(\mathbf{b}_{-i})$ turns out to be the boundary of S_0 .) However, in the process, instead of increasing x_i 's, we decrease the values of $f(\cdot)$ and $g_i(\cdot)$ to simplify our analysis (thus, the values of x_i 's do not change): For each $\mathbf{b}_{-i} \in T_i$, we decrease $f(\mathbf{b}_{-i}, b)$ by $\lambda \cdot \varepsilon$ for $b \geq c_i(\mathbf{b}_{-i})$ to have an equivalent effect on the first constraint of LS_3 ; further, we subtract $\frac{\varepsilon}{c_i(\mathbf{b}_{-i})}$ from $g_i(\mathbf{b}_{-i})$ to balance the update of x_i 's in the second constraint of LS_3 . The process continues until all $f(\cdot)$'s become 0, from which

point we get an equivalent solution $\{x_i(\mathbf{b})\}_{i, \mathbf{b}}$ to the original problem. We next describe the formal proof.

Let $R = \{\mathbf{b} \mid f(\mathbf{b}) > 0\}$ be the support of $f(\cdot)$. Since $f(\cdot)$ is a monotone function, we know that R is an upward closed set. If $R = \emptyset$, then we are done; thus, we assume that $R \neq \emptyset$. For any $S \in \Gamma$, the following chain of inequalities

$$\begin{aligned} \sum_{\mathbf{b} \in S} w(\mathbf{b}) \cdot f(\mathbf{b}) &= \sum_{\mathbf{b} \in S \cap R} w(\mathbf{b}) \cdot f(\mathbf{b}) \\ &\leq \lambda \cdot \sum_i \sum_{\mathbf{b}_{-i} \in (S \cap R)_{\downarrow i}} w(\mathbf{b}_{-i}) \\ &\leq \lambda \cdot \sum_i \sum_{\mathbf{b}_{-i} \in S_{\downarrow i}} w(\mathbf{b}_{-i}) \end{aligned}$$

is in fact an equality. Hence, the above two inequalities are tight and $S \cap R \in \Gamma$. Hence, in the following, we will only consider those tight sets in Γ that are contained in R .

Let $S_0 = R$. We will maintain a chain of upward closed sets $S_1, \dots, S_m \in \Gamma$ with the following structure:

$$R = S_0 \supseteq S_1 \supseteq S_2 \supseteq \dots \supseteq S_m = \emptyset.$$

During our process, we preserve the following key properties (\dagger):

Invariant properties (\dagger)

1. Inequality (3) holds for all upward closed sets.
2. All sets S_1, \dots, S_m in the chain $R = S_0 \supseteq S_1 \supseteq S_2 \supseteq \dots \supseteq S_m = \emptyset$ remain tight.
3. $f(\cdot)$ is nonnegative and monotonically increasing.
4. $g_i(\mathbf{b}_{-i})$ is nonnegative and monotonically decreasing on $(S_{j-1})_{\downarrow i} \setminus S_{j\downarrow i}$ for every $1 \leq i \leq n$ and $1 \leq j \leq m$.

We note that all these properties hold at the beginning of the process (e.g., we may simply choose $m = 1$ and $S_m = \emptyset$). Since $R \supseteq S_1$, we have

$$\sum_{\mathbf{b} \in R} w(\mathbf{b}) \cdot f(\mathbf{b}) > \sum_{\mathbf{b} \in S_1} w(\mathbf{b}) \cdot f(\mathbf{b}).$$

Since R satisfies condition (3) and $S_1 \in \Gamma$, we have

$$\begin{aligned} \sum_{\mathbf{b} \in R} w(\mathbf{b}) \cdot f(\mathbf{b}) &\leq \lambda \cdot \sum_i \sum_{\mathbf{b}_{-i} \in R_{\downarrow i}} g_i(\mathbf{b}_{-i}) \cdot w(\mathbf{b}_{-i}) \\ \sum_{\mathbf{b} \in S_1} w(\mathbf{b}) \cdot f(\mathbf{b}) &= \lambda \cdot \sum_i \sum_{\mathbf{b}_{-i} \in S_{1\downarrow i}} g_i(\mathbf{b}_{-i}) \cdot w(\mathbf{b}_{-i}). \end{aligned}$$

Hence,

$$\sum_i \sum_{\mathbf{b}_{-i} \in R_{\downarrow i}} g_i(\mathbf{b}_{-i}) \cdot w(\mathbf{b}_{-i}) > \sum_i \sum_{\mathbf{b}_{-i} \in S_{1\downarrow i}} g_i(\mathbf{b}_{-i}) \cdot w(\mathbf{b}_{-i}). \quad (4)$$

Thus, we can find i and \mathbf{b}_{-i} such that $\mathbf{b}_{-i} \in R_{\downarrow i} \setminus S_{1\downarrow i}$ and $g_i(\mathbf{b}_{-i}) > 0$. From now on to the end of the proof, for notational simplicity we use i to denote this particular index rather than a generic one. Let

$$T_i = \{\mathbf{b}_{-i} \in R_{\downarrow i} \setminus S_{1\downarrow i} \mid g_i(\mathbf{b}_{-i}) > 0\}$$

Note that $T_i \neq \emptyset$.

For each $\mathbf{b}_{-i} \in T_i$ we consider the smallest $c_i(\mathbf{b}_{-i}) \in \Lambda$ such that $f(\mathbf{b}_{-i}, c_i(\mathbf{b}_{-i})) \neq 0$. We note that $c_i(\mathbf{b}_{-i})$ is well

defined for each $\mathbf{b}_{\cdot i} \in T_i$ since $T_i \subset R_{\downarrow i}$. For the fixed index i , we intend to update \mathbf{x} as follows:

$$x_i(\mathbf{b}_{\cdot i}, b) \leftarrow x_i(\mathbf{b}_{\cdot i}, b) + \varepsilon, \text{ for all } \mathbf{b}_{\cdot i} \in T_i \text{ and } b \geq c_i(\mathbf{b}_{\cdot i}),$$

for some fixed $\varepsilon > 0$. In our process, instead of increasing \mathbf{x} we decrease the values of $f(\cdot)$ and $g_i(\cdot)$ as follows:

$$(\diamond) \begin{cases} f(\mathbf{b}_{\cdot i}, b) \leftarrow f(\mathbf{b}_{\cdot i}, b) - \lambda \cdot \varepsilon, & \forall \mathbf{b}_{\cdot i} \in T_i \text{ and } b \geq c_i(\mathbf{b}_{\cdot i}), \\ g_i(\mathbf{b}_{\cdot i}) \leftarrow g_i(\mathbf{b}_{\cdot i}) - \frac{\varepsilon}{c_i(\mathbf{b}_{\cdot i})}, & \forall \mathbf{b}_{\cdot i} \in T_i. \end{cases}$$

The decrements are with respect to the fixed index i only¹ and are implemented by continuously increasing ε from 0 until the value of one of $f(\mathbf{b})$ and $g_i(\mathbf{b}_{\cdot i})$ drops down to 0, or one more inequality (3) becomes tight for a new upward closed set.

Before describing how to proceed with the process, we establish some observations for the above updates (\diamond). Their proofs can be found in the full version of the paper.

CLAIM 2. *For any ε and upward closed set $S \subseteq S_1$, the two sides of the inequality (3) for S remain unchanged. In particular, the inequality is still tight for all sets S_1, \dots, S_m in the chain $R = S_0 \supseteq S_1 \supseteq S_2 \supseteq \dots \supseteq S_m = \emptyset$.*

CLAIM 3. *For any upward closed set $S \supset R \setminus S_1$ and ε , the condition (3) is still satisfied.*

CLAIM 4. *The function $f(\cdot)$ remains monotonically increasing.*

CLAIM 5. *The function $g_i(\cdot)$ remains monotonically decreasing on $(S_{j-1})_{\downarrow i} \setminus S_{j\downarrow i}$ for each set S_j in the chain $R = S_0 \supseteq S_1 \supseteq S_2 \supseteq \dots \supseteq S_m = \emptyset$.*

Combining Claim 2, 3, 4 and 5 together, we conclude that after update (\diamond) the aforementioned four key properties (\dagger) still hold. Thus, we continuously increase ε until a threshold point when one of the three boundary conditions (i.e., $f(\cdot) = 0$, $g_i(\cdot) = 0$, or inequality (3) becomes tight for a new set) becomes tight, which forbids us from further increasing ε . If more than one conditions become tight simultaneously, we consider them according to the following order.

- At a new point $f(\cdot)$ becomes 0. We then redefine set $R = \{\mathbf{b} \mid f(\mathbf{b}) > 0\}$. Note that by Claim 4, R remains to be an upward closed set. Further, we let $S_i \leftarrow S_i \cap R$ for all sets in the chain $R = S_0 \supseteq S_1 \supseteq S_2 \supseteq \dots \supseteq S_m = \emptyset$. (If some of the sets become the same, we contract the chain into a shorter one. With respect to this refined chain, $g_i(\cdot)$ is still monotonically decreasing on each $(S_{j-1})_{\downarrow i} \setminus S_{j\downarrow i}$.) Then we start over the whole process unless $R = \emptyset$, in which case the process terminates.
- At a new point $g_i(\cdot)$ becomes 0. We start over by finding a new index i and set T_i for inequality (4). (The set R and the chain remain the same.)
- There is a new upward closed set, say $S' \subset \Lambda^n$, for which inequality (3) becomes tight. Similarly, we assume without loss of generality that $S' \subset R$. First, S'

¹Because the process updates values only for one dimension at one step, the auction generated by our approach may not be symmetric. See the example in the next section.

cannot contain $R \setminus S_1$ due to Claim 3. Second, S_1 cannot contain S' due to Claim 2. By Claim 1, we know that $S^* \triangleq S' \cup S_1$ is also an upward closed tight set. We observe that $S^* = S' \cup S_1 \supseteq S_1$ and $R = S_0 \supseteq S^*$. Thus, we can plug S^* into the sequence

$$R = S_0 \supseteq S^* \supseteq S_1 \supseteq S_2 \supseteq \dots \supseteq S_m = \emptyset.$$

Since $g_i(\cdot)$ is a monotonically decreasing function on set $S_{0\downarrow i} \setminus S_{1\downarrow i}$, it is a monotonically decreasing function on both $S_{0\downarrow i} \setminus S^*_{\downarrow i}$ and $S^*_{\downarrow i} \setminus S_{1\downarrow i}$. All the inequalities (3) and conditions in (\diamond) continue to hold with respect to this new chain; hence, we can start over the whole process.

In each step of the update, we either increase the length of the tight set chain or decrease the support of $f(\cdot)$ or $g_i(\cdot)$. Thus, the process will eventually terminate as there are only finitely many $\mathbf{b} \in \Lambda^n$ and finitely many upward closed sets. We note that if there is $\mathbf{b} \in \Lambda^n$ with $f(\mathbf{b}) \neq 0$, then not all $g_i(\cdot)$'s can become 0, because otherwise (3) would be violated for $S = \Lambda^n$. Therefore, the process must terminate when $f(\mathbf{b}) = 0$ for every $\mathbf{b} \in \Lambda^n$.

Therefore, the respective sequence of updates of $\{x_i(\mathbf{b})\}_{i,\mathbf{b}}$ is a feasible solution to LS₂. Indeed, we have checked that the first and second constraints in LS₂ are satisfied. The third and fourth constraints hold true, because in every update (\diamond) we may only increase $x_i(\mathbf{b})$, and for each fixed $\mathbf{b}_{\cdot i}$ we increase $x_i(\mathbf{b}_{\cdot i}, b)$ monotonically in b . Hence, the benchmark function $f(\cdot)$ is attainable. \square

We further have the following claim, which will be used in the next section to prove that $\mathcal{F}^{(2)}$ is attainable.

CLAIM 6. *If $f(\cdot)$ is a symmetric monotone function (i.e., $f(\mathbf{b})$ is invariant under permutations of coordinates in \mathbf{b}), then it is sufficient to verify inequality (2) only for all symmetric upward closed sets for $f(\cdot)$ to be attainable.*

3.1 Unbounded and Continuous Domains

Our above analysis works for the discrete and bounded domain where all bids are in the set $\Lambda = \{(1 + \delta)^t \mid t = 0, 1, 2, \dots, N\}$. Our analysis continues to hold when $N = \infty$. Indeed, the argument for the necessity part of Theorem 2 for $N = \infty$ works in exactly the same way as the finite and bounded case. For the sufficiency part, we apply a standard argument from measure theory and mathematical analysis as follows.

For the unbounded domain, if inequality (2) in the characterization theorem holds true for all upward closed sets, then every finite version of LS₂ must have a feasible solution \mathbf{x}_N for every finite size N of the support $\Lambda(N) = \{1, 1 + \delta, \dots, (1 + \delta)^N\}$, where all outside densities are accumulated at the boundary of $\Lambda(N)$. We consider a corresponding feasible solution $\mathbf{z}_N = \{z_i(\mathbf{b})\}_{i,\mathbf{b}}$ to LS₁ for each $\Lambda(N)$. We treat the price distribution $z_i(\mathbf{b}_{\cdot i})$ for each bidder i as a measure defined on the σ -algebra generated by the sets $\{[1, 1 + \delta), \dots, [(1 + \delta)^N, \infty)\}^n$. Every solution \mathbf{z}_N to LS₁ is a vector of n measures $z_i(\mathbf{b})$, of which we altogether further regard as a measure defined on the σ -algebra σ_N generated by the sets $\{[1, 1 + \delta), \dots, [(1 + \delta)^N, \infty)\}^{n \times n}$. In a few steps below we show how to construct a measure \mathbf{z}^* that generates a feasible solution of LS₁ on every σ -algebra σ_N for every N .

1. We notice that for any $N_2 > N_1$, every solution \mathbf{z}_{N_2} to LS₁ for N_2 is also a solution of LS₁ for N_1 . Indeed, as a solution of LS₁ for N_1 we may just use the same measure \mathbf{z}_{N_2} on a smaller σ -algebra σ_{N_1} .
2. We observe that the set of all feasible solutions \mathbf{z}_N to LS₁ for any fixed N is a compact set (bounded and closed). We recall that one of the definitions of a compact set says that every infinite sequence \mathbf{z}_N ($N = 1, 2, \dots$) of points in a compact C must contain an infinite subsequence $\mathbf{z}_{N(j)}$ ($j = 1, 2, \dots$) that converges to a point $\bar{\mathbf{z}}_N \in C$.
3. For every fixed $N = \ell$, we consider a sequence of solutions \mathbf{z}_k to LS₁ for each $k = \ell, \ell + 1, \dots$. For each \mathbf{z}_k we get a feasible solution $\mathbf{z}_k|_\ell$ to LS₁ for $N = \ell$. Further, since the set of the solutions to LS₁ for $N = \ell$ is a compact set, we may choose an infinite subsequence that converges to $\bar{\mathbf{z}}_\ell$.
4. We note that by our construction, the measure $\bar{\mathbf{z}}_\ell$ can be extended to a feasible solution of LS₁ for every $k \geq \ell$, i.e., we can find a feasible solution $\hat{\mathbf{z}}_k$ to LS₁ for $N = k$ such that as a measure $\bar{\mathbf{z}}_\ell = \hat{\mathbf{z}}_k|_\ell$. Indeed, we can take the infinite sequence $\{N(j)\}$ for which $\mathbf{z}_{N(j)}|_\ell \rightarrow \bar{\mathbf{z}}_\ell$ and consider another infinite sequence of measures $\mathbf{z}_{N(j)}|_k$. Within the latter sequence we can choose an infinite subsequence converging to a feasible solution $\hat{\mathbf{z}}_k$ of LS₁ for $N = k$ (note that $\hat{\mathbf{z}}_k|_\ell = \bar{\mathbf{z}}_\ell$).
5. In fact, when constructing each $\bar{\mathbf{z}}_\ell$ we can ensure that $\bar{\mathbf{z}}_\ell|_j = \bar{\mathbf{z}}_j$ for every $j \leq \ell$.
6. Finally, we may define our measure \mathbf{z}^* on the unbounded domain Λ as a limit of $\bar{\mathbf{z}}_\ell$, where $\ell \rightarrow \infty$.

In a similar way by taking $\delta \rightarrow 0$, we can extend Theorem 2 to the case of continuous support. Hence, in the continuous and unbounded domain, the sufficient and necessary condition (2) in Theorem 2 translates into the following.

THEOREM 3. *A monotone function $f(\cdot)$ is λ -attainable if and only if for any measurable upward closed set $S \subset \mathbb{R}_{\geq 1}^n$*

$$\int_S f(\mathbf{b}) \cdot w(\mathbf{b}) \, d\mathbf{b} \leq \lambda \cdot \sum_{i=1}^n \int_{S_{\downarrow i}} w(\mathbf{b}_{-i}) \, d\mathbf{b}_{-i}, \quad (5)$$

where

$$w(t) = \frac{1}{t^2}, \quad w(\mathbf{b}) = \prod_{k=1}^n w(b_k), \quad w(\mathbf{b}_{-i}) = \prod_{k \neq i} w(b_k).$$

4. $\mathcal{F}^{(2)}$ IS ATTAINABLE

We show in this section that $\mathcal{F}^{(2)}$ is λ_n -attainable, for λ_n being the number conjectured in [14].

THEOREM 4. *$\mathcal{F}^{(2)}(\cdot)$ is λ_n -attainable for any $n \geq 2$, where*

$$\lambda_n = 1 - \sum_{i=2}^n \left(\frac{-1}{n}\right)^{i-1} \frac{i}{i-1} \binom{n-1}{i-1}.$$

For the sake of notational convenience, we employ $F_n(\mathbf{b})$ to denote $\mathcal{F}^{(2)}(\mathbf{b})$ when the input vector \mathbf{b} contains n bids. We let $\Lambda = [1, +\infty)$ and $\Lambda^n = \mathbb{R}_{\geq 1}^n$ denote the support of

bid vectors. Our proof is by induction on n . The base case $n = 2$ is trivial (where $\lambda_2 = 2$ and $F_2(\cdot)$ is 2-attainable given by the Vickrey auction). We assume that for all $k \leq n$, each $F_k(\cdot)$ is λ_k -attainable. In the rest of this section, we show that $F_{n+1}(\cdot)$ is λ_{n+1} -attainable.

For the equal revenue distribution with density $w(\cdot)$, it was shown in [13] that

$$\mathbb{E}_{\mathbf{b} \sim w} [F_k(\mathbf{b})] = \int_{\Lambda^k} F_k(\mathbf{b}) \cdot w(\mathbf{b}) \, d\mathbf{b} = \lambda_k \cdot k. \quad (6)$$

We note that $1 = \int_1^\infty w(t) \, dt$ and, therefore, we have

$$\int_{\Lambda^k} F_k(\mathbf{b}) \cdot w(\mathbf{b}) \, d\mathbf{b} = \lambda_k \cdot \sum_{i=1}^k \int_{\Lambda^{k-1}} w(\mathbf{b}_{-i}) \, d\mathbf{b}_{-i}. \quad (7)$$

Comparing to the inequality (5) in Theorem 3, we have the following claim.

CLAIM 7. *For $k = 2, \dots, n$, each λ_k is chosen so that inequality (5) is tight for the set Λ^k .*

Let

$$G_n(\mathbf{z}) = F_{n+1}(1, \mathbf{z}), \quad \forall \mathbf{z} \in \Lambda^n.$$

By the following claim, to prove that $F_{n+1}(\mathbf{b})$ with $\mathbf{b} \in \Lambda^{n+1}$ is λ_{n+1} -attainable, it suffices to show that $G_n(\mathbf{z})$ with $\mathbf{z} \in \Lambda^n$ is λ_{n+1} -attainable.

CLAIM 8. *If $G_n(\mathbf{z})$ with $\mathbf{z} \in \Lambda^n$ is λ_{n+1} -attainable, then $F_{n+1}(\mathbf{b})$ with $\mathbf{b} \in \Lambda^{n+1}$ is λ_{n+1} -attainable.*

Given the above claim, we now show that $G_n(\mathbf{z})$ behaves similarly as $F_{n+1}(\mathbf{b})$ in (6) (the proofs are deferred to the full version of the paper).

CLAIM 9. *The set Λ^n is tight for the benchmark $G_n(\mathbf{z})$ and competitive ratio λ_{n+1} in (2), i.e.,*

$$\int_{\Lambda^n} G_n(\mathbf{z}) \cdot w(\mathbf{z}) \, d\mathbf{z} = \lambda_{n+1} \cdot n.$$

Let

$$H_n(\mathbf{z}) = \max(0, n + 1 - F_n(\mathbf{z})), \quad \forall \mathbf{z} \in \Lambda^n.$$

We observe that $G_n(\mathbf{z}) = F_{n+1}(1, \mathbf{z}) = \max(n + 1, F_n(\mathbf{z}))$ for $\mathbf{z} \in \Lambda^n$; thus, $G_n(\mathbf{z}) = F_n(\mathbf{z}) + H_n(\mathbf{z})$. We note that due to the induction hypothesis, $F_n(\mathbf{z})$ is λ_n -attainable. This means that for every symmetric upward closed set $S \subset \Lambda^n$, we have

$$\int_S F_n(\mathbf{z}) \cdot w(\mathbf{z}) \, d\mathbf{z} \leq \lambda_n \cdot \sum_{i=1}^n \int_{S_{\downarrow i}} w(\mathbf{z}_{-i}) \, d\mathbf{z}_{-i}.$$

We show in Claim 10 below that for every symmetric upward closed set $S \subset \Lambda^n$,

$$\int_S H_n(\mathbf{z}) \cdot w(\mathbf{z}) \, d\mathbf{z} \leq (\lambda_{n+1} - \lambda_n) \cdot \sum_{i=1}^n \int_{S_{\downarrow i}} w(\mathbf{z}_{-i}) \, d\mathbf{z}_{-i}.$$

If we combine the above two inequalities, then we get for every symmetric upward closed set $S \subset \Lambda^n$,

$$\int_S G_n(\mathbf{z}) \cdot w(\mathbf{z}) \, d\mathbf{z} \leq \lambda_{n+1} \cdot \sum_{i=1}^n \int_{S_{\downarrow i}} w(\mathbf{z}_{-i}) \, d\mathbf{z}_{-i}.$$

We observe that $F_n(\cdot), G_n(\cdot)$ and $H_n(\cdot)$ all are symmetric functions. By Theorem 6, we know that $G_n(\mathbf{z})$ is λ_{n+1} -attainable. Hence, we are only left to prove the following Claim 10 in order to show λ_{n+1} -attainability of $F_{n+1}(\mathbf{b})$ and complete the proof of Theorem 4.

CLAIM 10. *For every symmetric upward closed set $S \subset \Lambda^n$,*

$$\int_S H_n(\mathbf{z}) \cdot w(\mathbf{z}) \, d\mathbf{z} \leq (\lambda_{n+1} - \lambda_n) \cdot \sum_{i=1}^n \int_{S_{\downarrow i}} w(\mathbf{z}_i) \, d\mathbf{z}_i. \quad (8)$$

PROOF. Since S is a symmetric set, all projections $S_{\downarrow i}$ are in fact equal to the same set, denoted by $X \subset \Lambda^{n-1}$. We notice that S can be divided into n (almost) disjoint sets $S = S_1 \cup \dots \cup S_n$, where each $S_i = \{(\mathbf{z}_i, z_i) \mid \mathbf{z}_i \in X, z_i \geq \max(\mathbf{z}_i)\}$ and $\max(\mathbf{z}_i)$ denotes the largest value of all coordinates of \mathbf{z}_i . Furthermore, note that for any $z_i \in [\max(\mathbf{z}_i), \infty)$, the value of the benchmark $F_n(\mathbf{z}_i, z_i)$ is fixed; thus, the value of $H_n(\mathbf{z}_i, z_i)$ is also fixed. Hence, we have the following equalities.

$$\begin{aligned} & \int_S H_n(\mathbf{z}) \cdot w(\mathbf{z}) \, d\mathbf{z} \\ &= \sum_{i=1}^n \int_{S_i} H_n(\mathbf{z}) \cdot w(\mathbf{z}) \, d\mathbf{z} \\ &= n \cdot \int_X w(\mathbf{z}_n) \int_{\max(\mathbf{z}_n)}^{\infty} H_n(\mathbf{z}_n, t) \cdot w(t) \, dt \, d\mathbf{z}_n \\ &= n \cdot \int_X \frac{H_n(\mathbf{z}_n, \max(\mathbf{z}_n))}{\max(\mathbf{z}_n)} w(\mathbf{z}_n) \, d\mathbf{z}_n. \end{aligned} \quad (9)$$

By Claim 9, we know that

$$\int_{\Lambda^n} G_n(\mathbf{z}) \cdot w(\mathbf{z}) \, d\mathbf{z} = \lambda_{n+1} \cdot n = \lambda_{n+1} \cdot \sum_{i=1}^n \int_{\Lambda^{n-1}} w(\mathbf{z}_i) \, d\mathbf{z}_i.$$

Further, equality (7) says

$$\int_{\Lambda^n} F_n(\mathbf{z}) \cdot w(\mathbf{z}) \, d\mathbf{z} = \lambda_n \cdot \sum_{i=1}^n \int_{\Lambda^{n-1}} w(\mathbf{z}_i) \, d\mathbf{z}_i.$$

Thus,

$$\begin{aligned} \int_{\Lambda^n} H_n(\mathbf{z}) \cdot w(\mathbf{z}) \, d\mathbf{z} &= \int_{\Lambda^n} (G_n(\mathbf{z}) - F_n(\mathbf{z})) \cdot w(\mathbf{z}) \, d\mathbf{z} \\ &= (\lambda_{n+1} - \lambda_n) \cdot \sum_{i=1}^n \int_{\Lambda^{n-1}} w(\mathbf{z}_i) \, d\mathbf{z}_i. \end{aligned}$$

That is, Λ^n is a tight set for (8). This implies that

$$\begin{aligned} (\lambda_{n+1} - \lambda_n) \cdot n &= (\lambda_{n+1} - \lambda_n) \cdot \sum_{i=1}^n \int_{\Lambda^{n-1}} w(\mathbf{z}_i) \, d\mathbf{z}_i \\ &= \int_{\Lambda^n} H_n(\mathbf{z}) \cdot w(\mathbf{z}) \, d\mathbf{z} \\ &= n \int_{\Lambda^{n-1}} \frac{H_n(\mathbf{z}_n, \max(\mathbf{z}_n))}{\max(\mathbf{z}_n)} w(\mathbf{z}_n) \, d\mathbf{z}_n \end{aligned} \quad (10)$$

Therefore,

$$\begin{aligned} (\lambda_{n+1} - \lambda_n) \sum_{i=1}^n \int_{S_{\downarrow i}} w(\mathbf{z}_i) \, d\mathbf{z}_i &= (\lambda_{n+1} - \lambda_n) \cdot n \cdot \int_X w(\mathbf{z}_n) \, d\mathbf{z}_n \\ &= n \cdot \int_{\Lambda^{n-1}} \frac{H_n(\mathbf{z}_n, \max(\mathbf{z}_n))}{\max(\mathbf{z}_n)} w(\mathbf{z}_n) \, d\mathbf{z}_n \cdot \int_X w(\mathbf{z}_n) \, d\mathbf{z}_n \end{aligned} \quad (11)$$

Combining (9) and (11), we are left to show that

$$\begin{aligned} & \int_X \frac{H_n(\mathbf{z}_n, \max(\mathbf{z}_n))}{\max(\mathbf{z}_n)} w(\mathbf{z}_n) \, d\mathbf{z}_n \\ & \leq \int_{\Lambda^{n-1}} \frac{H_n(\mathbf{z}_n, \max(\mathbf{z}_n))}{\max(\mathbf{z}_n)} w(\mathbf{z}_n) \, d\mathbf{z}_n \cdot \int_X w(\mathbf{z}_n) \, d\mathbf{z}_n \end{aligned} \quad (12)$$

Since $H_n(\mathbf{z})$ is a non-negative monotonically decreasing function, the function $H'(\mathbf{z}_n) \triangleq \frac{H_n(\mathbf{z}_n, \max(\mathbf{z}_n))}{\max(\mathbf{z}_n)}$ is also non-negative monotonically decreasing. Let $\mathbb{1}_X(\mathbf{z}_n)$ denote the characteristic function of the set X . Note that $\mathbb{1}_X(\cdot)$ is a non-negative monotonically increasing function. Now the inequality (12) simply reads as

$$\begin{aligned} & \int_{\Lambda^{n-1}} H'(\mathbf{z}_n) \cdot \mathbb{1}_X(\mathbf{z}_n) \cdot w(\mathbf{z}_n) \, d\mathbf{z}_n \\ & \leq \int_{\Lambda^{n-1}} H'(\mathbf{z}_n) \cdot w(\mathbf{z}_n) \, d\mathbf{z}_n \cdot \int_{\Lambda^{n-1}} \mathbb{1}_X(\mathbf{z}_n) \cdot w(\mathbf{z}_n) \, d\mathbf{z}_n. \end{aligned}$$

The last inequality for an arbitrary non-negative monotonically increasing function (in our case it is $\mathbb{1}_X(\cdot)$), an arbitrary non-negative monotonically decreasing function (in our case it is $H'(\cdot)$), and an arbitrary product measure (in our case it is $w(\mathbf{z}_i)$) is known as a special case of Fortuin-Kasteleyn-Ginibre (FKG) inequality [11] (in the uni-variate case this inequality is called Chebyshev Integral inequality). Therefore, the inequality holds and the claim follows. \square

5. K -ITEM VICKREY AUCTION

In this section, we consider the benchmark $\text{MAXV}(\mathbf{b}) = \max_{1 \leq k < n} k \cdot b_{k+1}$ provided by the largest revenue of the k -item Vickrey auction across all possible values of supply k .

First of all, we calculate the lower bound on the competitive ratio for $\text{MAXV}(\cdot)$ using similar approach as in [13]. We employ equal revenue distribution \mathbf{B} of bid vectors, that is i.i.d. with the density function $w(b) = \frac{1}{b^2}$, cumulative density $1 - \frac{1}{b}$, and the support $[1, \infty)$. The key technical problem is to compute the expected value of the benchmark $\text{MAXV}(\mathbf{B})$. Following [13], we compute the probability $\Pr[\text{MAXV}(\mathbf{B}) \geq z]$ for any given z . Since $\text{MAXV}(\mathbf{B})$ is at least $n - 1$, we only need to compute the probability for $z \geq n - 1$. Let a random variable V_i be the i -th largest bid in \mathbf{B} . We further define a set of random variables as

$$F_{n,k} = \max_{i=1,2,\dots,n} (k + i - 1) \cdot V_i.$$

We note that so far we follow the idea of [13], but adjust definitions accordingly to the benchmark $\text{MAXV}(\cdot)$. Intuitively, $F_{n,k}$ captures the value of $\text{MAXV}(\cdot)$ given k additional bid-

ders with bid V_1 . Let \mathcal{H}_i denote the event

$$V_i \geq \frac{z}{k+i-1} \quad \text{and} \quad \bigwedge_{j=i+1, i+2, \dots, n} V_j < \frac{z}{k+j-1}.$$

The probability of \mathcal{H}_i can be written as

$$\Pr[\mathcal{H}_i] = \binom{n}{i} \left(\frac{k+i-1}{z} \right)^i \Pr[F_{n-i, k+i} < z].$$

Since \mathcal{H}_i 's are mutually exclusive and the event $F_{n,k} \geq z$ is the union of \mathcal{H}_i for $i = 1, 2, \dots, n$, we get

$$\begin{aligned} \Pr[F_{n,k} \geq z] &= \sum_i \Pr[\mathcal{H}_i] \\ &= \sum_i \binom{n}{i} \left(\frac{k+i-1}{z} \right)^i \Pr[F_{n-i, k+i} < z]. \end{aligned} \quad (13)$$

This gives a recursive relation for $\Pr[F_{n,k} \geq z]$ and the boundary condition is $\Pr[F_{0,k} \geq z] = 0$.

By induction, we can prove that

$$\Pr[F_{n,k} \geq z] = 1 - \frac{(z-n-k+1)(z+1-k)^{n-1}}{z^n}.$$

Therefore, we have

$$\Pr[\text{MAXV}(\mathbf{B}) \geq z] = \Pr[F_{n,0} \geq z] = 1 - \frac{(z-n+1)(z+1)^{n-1}}{z^n}.$$

Now we can get the expectation by the following calculation.

$$\begin{aligned} \mathbf{E}[\text{MAXV}(\mathbf{B})] &= \int_0^\infty \Pr[\text{MAXV}(\mathbf{B}) \geq z] dz \\ &= n-1 + \int_{n-1}^\infty \left(1 - \frac{(z-n+1)(z+1)^{n-1}}{z^n} \right) dz \\ &= n-1 + \int_{n-1}^\infty \frac{\sum_{i=0}^{n-2} \binom{n}{i} (n-i-1) z^i}{z^n} dz \\ &= n-1 + \sum_{i=0}^{n-2} \binom{n}{i} (n-i-1) \int_{n-1}^\infty \frac{1}{z^{n-i}} dz \\ &= n-1 + \sum_{i=0}^{n-2} \binom{n}{i} \frac{n-i-1}{(n-i-1)(n-1)^{n-i-1}} \\ &= n-1 + \frac{(n^n - (n-1)^n - n(n-1)^{n-1})}{(n-1)^{n-1}} \\ &= n \left(\frac{n}{n-1} \right)^{n-1} - n. \end{aligned}$$

Thus, we get a lower bound on the competitive ratio for $\text{MAXV}(\cdot)$

$$\gamma_n \triangleq \frac{\mathbf{E}[\text{MAXV}(\mathbf{B})]}{n} = \left(\frac{n}{n-1} \right)^{n-1} - 1,$$

where the denominator is n , as the expected revenue of any auction is at most n on the equal revenue distribution. Finally, we note that γ_n approaches $e-1$ as n goes to infinity.

On the other hand, we can show that $\text{MAXV}(\cdot)$ is γ_n -attainable by our characterization theorem. The proof is almost identically to our argument for $\mathcal{F}^{(2)}$ in the previous section. We only replace λ_n by γ_n , $\mathcal{F}^{(2)}(\cdot)$ by $\text{MAXV}(\cdot)$, and $H_n(\mathbf{z}) = \max(0, n+1 - F_n(\mathbf{z}))$ by $H_n(\mathbf{z}) = \max(0, n -$

$F_n(\mathbf{z}))$. In the argument for $\mathcal{F}^{(2)}$, we do not use any explicit formula for $\mathcal{F}^{(2)}$, but employ a few properties such as monotonicity, symmetry, linearity, and the fact that the function value does not change if one increases the largest bid. These properties continue to hold for the benchmark $\text{MAXV}(\cdot)$. The definition of $H_n(\mathbf{z}) = \max(0, n+1 - F_n(\mathbf{z}))$ is the only place where we use the formula for $\mathcal{F}^{(2)}$, which is replaced by $H_n(\mathbf{z}) = \max(0, n - F_n(\mathbf{z}))$ for $\text{MAXV}(\cdot)$. Note that the properties that we use are that $H_n(\mathbf{z})$ is decreasing and does not change if one increases the largest bid, which continue to hold for the modified $H_n(\mathbf{z})$. Therefore, the whole proof of the attainability of $\mathcal{F}^{(2)}$ carries over to the benchmark $\text{MAXV}(\cdot)$.

We summarize our results in the following theorem.

THEOREM 5. *For the benchmark $\text{MAXV}(\cdot)$ and any $n \geq 2$, the optimal competitive ratio of a truthful auction is $\gamma_n = \left(\frac{n}{n-1} \right)^{n-1} - 1$.*

6. CONCLUDING REMARKS

Our paper studies designing optimal competitive digital goods auctions. The proof of the characterization theorem gives an explicit procedure that constructs an auction with the optimal competitive ratio. The construction, however, is rather abstract and may take exponential steps for some benchmarks. It is therefore an intriguing question to see if there is any simple form for the description of optimal auctions (with respect to, e.g., the $\mathcal{F}^{(2)}$ benchmark).

We remark that our characterization applies not only to symmetric benchmarks such as $\mathcal{F}^{(2)}$ but also non-symmetric benchmarks. In particular, Theorem 2 may better estimate the competitive ratio of the monotone-price benchmark $\mathcal{M}^{(2)}$ of Leonardi and Roughgarden [22].

Note that the characterization that we have addressed in Theorem 2 is fundamental to our understanding and discovery of the optimal competitive ratios. The next step is to identify more general settings for which such characterizations exist. A deeper understanding of the result could potentially lead to an answer to the following more general question: For what class of prior free auction design problems the performance of the optimal auction with respect to a given benchmark exactly matches the expected performance of an optimal auction tailored to the worst case distribution under that benchmark? That is, what are the essential features of a problem to guarantee that there is no duality gap between the problem of designing an optimal auction and the problem of computing a worst case distribution? A fundamental question in prior free auction design (and perhaps also in online algorithm design) is to identify the class of problems for which such strong duality holds.

All our discussions in this paper are only for the unlimited supply case. However, our results to some extent carry over to the limited supply setting. In a limited supply setting, let k be the number of units for sale and $n > k$ be the number of bidders. Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ be any non-negative and monotone benchmark. We define

$$f_1(b_1, \dots, b_k) = f(b_1, \dots, b_{k-1}, b_k, b_k, \dots, b_k)$$

and

$$f_2(b_1, \dots, b_k) = f(b_1, \dots, b_{k-1}, b_k, 0, \dots, 0),$$

for each $\mathbf{b} = (b_1, \dots, b_n)$ with $b_1 \geq b_2 \geq \dots \geq b_n$. That is, we 'increase' and 'decrease' in $f(\cdot)$ the last $n-k$ bids to b_k

and 0, respectively. We have $f_1(b_1, \dots, b_n) \geq f(b_1, \dots, b_n) \geq f_2(b_1, \dots, b_n)$ for any \mathbf{b} due to the monotonicity condition. We note that $f_1(\cdot)$ and $f_2(\cdot)$ only depend on the highest k bids and may be viewed as two other benchmarks.

Now we may find the optimal competitive ratios of unlimited supply auctions with k bidders for both benchmarks $f_1(\cdot)$ and $f_2(\cdot)$ with the help of Theorem 2. We note that the competitive ratios for the $f_1(\cdot)$ and $f_2(\cdot)$ benchmarks give upper and lower bounds, respectively, on the competitive ratio of k -unit auctions with respect to the benchmark $f(\cdot)$. Indeed, we may construct the following limited supply auction: given a vector of bids $\mathbf{b} = (b_1, \dots, b_n)$ with $b_1 \geq b_2 \geq \dots \geq b_n$, we admit only the k highest bids and then run the optimal competitive auction with respect to the $f_1(\cdot)$ benchmark. We note that the revenue of such an auction can only be higher than the respective revenue of the unlimited supply auction for $f_1(\cdot)$.

We observe that if $f(\cdot)$ only depends on the k highest bids, then $f(\cdot) = f_1(\cdot) = f_2(\cdot)$ and their respective competitive ratios are the same. This is the case for the benchmark $\mathcal{F}^{(2,k)}$, considered previously in [19, 14] and defined as the optimal omniscient fixed price auction that sells between 2 and k items. Therefore, without much work we obtain optimal competitive k -unit auctions with respect to the benchmark $f(\cdot)$. We note that for other benchmark functions, the question of designing optimal competitive k -unit auctions remains open; we leave this interesting question for the future work.

7. ACKNOWLEDGMENTS

We thank Jason Hartline for helpful discussions and valuable comments.

8. REFERENCES

- [1] G. Aggarwal, A. Fiat, A. Goldberg, J. Hartline, N. Immerlica, and M. Sudan. Derandomization of auctions. In *STOC*, pages 619–625. ACM, 2005.
- [2] S. Alaei, A. Malekian, and A. Srinivasan. On random sampling auctions for digital goods. In *EC*, pages 187–196, 2009.
- [3] I. Ashlagi, M. Braverman, A. Hassidim, and D. Monderer. Monotonicity and implementability. *Econometrica*, 78(5):1749–1772, 2010.
- [4] Z. Bar-yossef, K. Hildrum, and F. Wu. Incentive-compatible online auctions for digital goods. In *SODA*, pages 964–970. ACM/SIAM, 2002.
- [5] S. Bhattacharya, E. Koutsoupias, J. Kulkarni, S. Leonardi, T. Roughgarden, and X. Xu. Near-optimal multi-unit auctions with ordered bidders. In *EC'13*, pages 91–102, 2013.
- [6] S. Bikhchandani, S. Chatterji, R. Lavi, A. Mu'alem, N. Nisan, and A. Sen. Weak monotonicity characterizes deterministic domain strategy implementation. *Econometrica*, 74(4):1109–1132, 2006.
- [7] A. Blum and J. Hartline. Near-optimal online auctions. In *SODA*, pages 1156–1163. SIAM, 2005.
- [8] N. Devanur and J. Hartline. Limited and online supply and the bayesian foundations of prior-free mechanism design. In *EC*, pages 41–50. ACM, 2009.
- [9] U. Feige, A. Flaxman, J. D. Hartline, and R. D. Kleinberg. On the competitive ratio of the random sampling auction. In *WINE*, pages 878–886, 2005.
- [10] A. Fiat, A. V. Goldberg, J. D. Hartline, and A. R. Karlin. Competitive generalized auctions. In *STOC*, pages 72–81, 2002.
- [11] C. M. Fortuin, P. W. Kasteleyn, and J. Ginibre. Correlation inequalities on some partially ordered sets. *Communications in Mathematical Physics*, 22(2):89–103, 1971.
- [12] A. V. Goldberg and J. D. Hartline. Competitiveness via consensus. In *SODA*, pages 215–222. ACM/SIAM, 2003.
- [13] A. V. Goldberg, J. D. Hartline, A. R. Karlin, and M. Saks. A lower bound on the competitive ratio of truthful auctions. In *STACS*, pages 644–655, 2004.
- [14] 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.
- [15] A. V. Goldberg, J. D. Hartline, and A. Wright. Competitive auctions and digital goods. In *SODA*, pages 735–744, 2001.
- [16] N. Gravin and P. Lu. Competitive auctions for markets with positive externalities. In *ICALP (2)*, pages 569–580. Springer, 2013.
- [17] S. Hart and N. Nisan. Approximate revenue maximization with multiple items. In *EC*, pages 656–656. ACM, 2012.
- [18] J. Hartline and Q. Yan. Envy, truth, and profit. In *EC*, pages 243–252. ACM, 2011.
- [19] J. D. Hartline and R. McGrew. From optimal limited to unlimited supply auctions. In *EC*, pages 175–182. ACM, 2005.
- [20] J. D. Hartline and T. Roughgarden. Optimal mechanism design and money burning. In *STOC*, pages 75–84, 2008.
- [21] T. Ichiba and K. Iwama. Averaging techniques for competitive auctions. In *ANALCO*, pages 74–81, 2010.
- [22] S. Leonardi and T. Roughgarden. Prior-free auctions with ordered bidders. In *STOC'12*, pages 427–434, 2012.
- [23] M. Mahdian and A. Saberi. Multi-unit auctions with unknown supply. In *EC*, pages 243–249. ACM, 2006.
- [24] R. B. Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58, 1981.
- [25] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic game theory*. Cambridge University Press, 2007.
- [26] M. E. Saks and L. Yu. Weak monotonicity suffices for truthfulness on convex domains. In *EC*, pages 286–293. ACM, 2005.