Competitive Analysis via Benchmark Decomposition

NING CHEN, Nanyang Technological University
NICK GRAVIN, Microsoft Research.
PINYAN LU, Microsoft Research.

We propose a uniform approach for the design and analysis of prior-free competitive auctions and online auctions. Our philosophy is to view the benchmark function as a variable parameter of the model and study a broad class of functions instead of a individual target benchmark. We consider a multitude of well-studied auction settings, and improve upon a few previous results.

— Multi-unit auctions. Given a β-competitive unlimited supply auction, the best previously known multi-unit auction is 2β-competitive. We design a $(1 + \beta)$-competitive auction reducing the ratio from 4.84 to 3.24. These results carry over to matroid and position auctions.
— General downward-closed environments. We design a 6.5-competitive auction improving upon the ratio of 7.5. Our auction is noticeably simpler than the previous best one.
— Unlimited supply online auctions. Our analysis yields an auction with a competitive ratio of 4.12, which significantly narrows the margin of $[4, 4.84]$ previously known for this problem.

A particularly important tool in our analysis is a simple decomposition lemma, which allows us to bound the competitive ratio against a sum of benchmark functions. We use this lemma in a "divide and conquer" fashion by dividing the target benchmark into the sum of simpler functions.

Categories and Subject Descriptors: J.2 [**Social and Behavioral Sciences**]: Economics; F.2.3 [**Analysis of Algorithms and Problem Complexity**]: Tradeoffs among Complexity Measures

Additional Key Words and Phrases: Competitive analysis; Digital good auctions; Online auctions; Position auction; Benchmark decomposition

1. INTRODUCTION

Revenue maximization in multi-unit auctions is a canonical problem that has attracted a lot of attention from algorithmic game theory community over the last fifteen years. In this framework auctioneer sells an abstract service to n potential customers participating in the auction. The auctioneer has a feasibility constraint on which sets of agents can be served simultaneously. Each bidder i values the service at a single privately known value v_i. Auctioneer runs a single-round auction, where each bidder submits a sealed bid b_i. After soliciting the bids the auctioneer decides on whether each bidder i receives the service and the amount that i pays. The auctioneer's goal is to maximize *revenue*, i.e., the total payment of the bidders. In our work we consider the following well-motivated and intensively studied settings that fall into this general framework.

— Digital goods. The auctioneer sells unlimited number of copies of a single item, so that any set of winners is feasible.

Authors addresses: N. Chen, Division of Mathematical Sciences School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore; email: ningc@ntu.edu.sg; N. Gravin, Microsoft Research, Cambridge, MA, USA; email: ngravin@microsoft.com; Pinyan Lu, Microsoft Research, Shanghai, China pinyanl@microsoft.com.

EC'15, June 15–19, 2015, Portland, OR, USA. ACM 978-1-4503-3410-5/15/06.
http://dx.doi.org/10.1145/2764468.2764491

— Limited supply (a.k.a. multi-unit) auctions. There are only ℓ copies of the item in total, and thus, at most ℓ bidders can be served. There is a reduction [21] to limited supply auctions from position and matroid environments, which are related to FCC spectrum auction and sponsored search.
— General downward-closed permutation environments. Set system of feasible sets is downward-closed if any subset of a feasible set is also feasible. The auctioneer has a probabilistic feasibility constraint, i.e., given by a probability distribution over feasible sets. Bidders are assumed to be symmetric, i.e., bidders' values are revealed in a random order. This is a generalization of multi-unit and position auctions, as well as matching environments, where each feasible set represents vertices on the one side of a bipartite matching.
— Online auctions. The bidders arrive online one by one and the auctioneer makes an irrevocable decision (whether the bidder receives a service or not and at what price) immediately as each new bidder arrives. Online auctions capture important scenarios such as sales on the Internet, where bidders may appear at any time and want to receive service right away.

Our approach. A common theme in the design and analysis of online algorithms and prior-free auctions is the competitive framework. In both cases, an online algorithm, which has to make irrevocable decisions online, or a truthful mechanism, whose outcomes must be aligned with bidders' incentives, are competing against a benchmark corresponding to a desirable outcome. We propose a uniform approach for the design and analysis of prior-free competitive auctions and online auctions. Our philosophy is to view the benchmark function as a variable parameter of the model and study a broad class of functions instead of a individual target benchmark. Our approach is surprisingly simple and relies on a very basic decomposition lemma, which allows us to bound the competitive ratio against a sum of benchmark functions. We use this lemma in a "divide and conquer" fashion by dividing the target benchmark into the sum of simpler benchmark, each of which admit a good approximation. This allows us to improve some of the best known results for the aforementioned settings and sometimes along the way simplify mechanisms and analysis from the prior literature.

Notably, all our significant modifications to the benchmarks are derived *only* for the basic setting of the digital good auction. On the the other hand, the black-box reductions to the digital good auction employed in the prior work usually deal with the same, or analogous benchmarks in the base and derived settings. In contrast to this paradigm, our work demonstrates that it may be better to deliberately use different benchmarks in the two settings. This signifies importance of developing theory and performing competitive analysis for a variety of benchmarks, if nothing else but at least as an intermediate step towards better design and analysis of more complex settings.

1.1. Benchmarks and Competitive Analysis

We assume that all bidders are selfish and aim to maximize their own utility. Buyers are assumed to have quasi-linear utility, i.e., if bidder i gets served, his *utility* is the difference between his value v_i and his payment; otherwise, the bidder pays 0 and his utility is 0. We say that an auction is *truthful* or incentive compatible if it is a dominant strategy for every bidder i to bid his private value, i.e., $b_i = v_i$, no matter how other bidders behave. A randomized auction is (universally) truthful if it is given by a distribution over deterministic truthful auctions.

The objective is to design auctions that maximize revenue of the auctioneer. To evaluate the performance of an auction, we need to define a benchmark function $f : \mathbb{R}^n \to \mathbb{R}$, where $f(\mathbf{b})$ measures our target revenue for the bid vector $\mathbf{b} \in \mathbb{R}^n$. 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

$$\mathbf{E}\left[\mathcal{A}(\mathbf{b})\right] \geq \frac{f(\mathbf{b})}{\lambda}, \quad \forall \mathbf{b} = (b_1, \ldots, b_n) \in \mathbb{R}^n$$

where $\mathbf{E}[\mathcal{A}(\mathbf{b})]$ is the expected revenue of auction $\mathcal{A}$ on the bid vector $\mathbf{b}$.

A benchmark function should be, on the one hand, economically meaningful in providing a revenue target, and on the other hand not too ambitious so that a truthful auction may have small competitive ratio against the benchmark. For the auction with unlimited supply, the most well-studied benchmark is $\mathcal{F}^{(2)}(\mathbf{b}) = \max_{2 \leq k \leq n} \ k \cdot b_{(k)}$, where bids are ordered so that $b_{(1)} \geq b_{(2)} \geq \cdots \geq b_{(n)}$. That is, $\mathcal{F}^{(2)}$ gives the largest possible revenue of a fixed price sale given that there are at least two buyers. The reason for having at least two winners is that otherwise, all but one bidders may have 0 value and then no truthful auction can be competitive against a single bidder with arbitrary large private value. Another meaningful benchmark is $\textsc{maxV}(\mathbf{b}) = \max_{1 \leq k < n} \ k \cdot b_{(k+1)}$. We note that $k \cdot b_{(k+1)}$ is the revenue of the Vickrey auction selling k units. Hence, $\textsc{maxV}$ is the maximum revenue of a k-unit Vickrey auction for all possible values of the supply k.

For the ℓ-unit auctions, one can naturally extend the definition of $\mathcal{F}^{(2)}$ to

$$\mathcal{F}^{(2,\ell)}(\mathbf{b}) = \max_{2 \leq k \leq \ell} k \cdot b_{(k)}.$$

This is the largest possible revenue of a fixed price sale given that there are at least two and at most ℓ buyers (as there are only ℓ copies).

Hartline and Yan [21] gave another interpretation of $\mathcal{F}^{(2)}$ for the unlimited supply setting. Namely, it is the optimal revenue one can extract in an envy-free allocation with at least two winners. The definition was extended in [21] to more general environments such as limited supply and general downward-closed environments. We denote by $\mathrm{EFO}^{(2)}$ the largest revenue that can be obtained in an envy-free allocation for a slightly modified bid vector $\mathbf{b}^2 = (b_2, b_2, \ldots, b_n)$. Interestingly, although $\mathrm{EFO}^{(2)}$ coincides with $\mathcal{F}^{(2)}$ in the digital goods setting, $\mathrm{EFO}^{(2)}$ is not the same as $\mathcal{F}^{(2,\ell)}$ for the limited supply case, where the precise formula for $\mathrm{EFO}^{(2)}$ will be given in Section 2.

For the online setting, we focus on the model of Koutsoupias and Pierrakos [24] of unlimited supply auctions competing against the benchmark $\mathcal{F}^{(2)}$ or against $\textsc{maxV}$. We assume random arrival order of the bidders, as if bidders arrive in an adversarial order, competitive ratio cannot be a constant [24].

1.2. Results and Techniques

Our recent work [7] on digital goods auction proposed a uniform procedure for calculating the *optimal* competitive ratio against any monotone benchmark. In particular, it yielded tight competitive ratios against the $\textsc{maxV}$ and $\mathcal{F}^{(2)}$ benchmarks. Here, we study the design of competitive auctions in other settings. We summarize previous and our new results in the following table. (Since some of the earlier work used auctions for digital goods as a black-box, we update their bounds accordingly with the new tight bounds of [7].)

	Limited supply $\mathrm{EFO}^{(2)}$	Downward-closed $\mathrm{EFO}^{(2)}$	Online $\mathcal{F}^{(2)}$	Online $\textsc{maxV}$
Previous upper bounds	4.84 [21]	7.5 [9]	4.84 [24]	-
Our upper bounds	3.24	6.5	4.12	2.42
Lower bounds	2.42 [14]	2.42 [14]	4 [24]	2

All bounds in the table are for the worst-case scenarios when the number of bidders n can be arbitrarily large. Better bounds are known for every fixed number of bidders, although, as n grows, these bounds quickly converge to the worst-case bounds given in the table.

An important new perspective of [7] was to view benchmark function as a variable parameter of the model. It should be noted that to a limited degree an earlier work [22] also studied a broad class of benchmarks in the digital goods setting. In this paper, we continue to explore this idea in more general settings. Unlike [7] we *explicitly design* auctions with improved competitive ratios.

The following general yet simple observation appears to be very helpful in our analysis. When we seek for an auction with good performance against a specific benchmark $f(\mathbf{b})$, it is often the case that $f(\mathbf{b})$ can be decomposed into the sum of two functions $f(\mathbf{b}) = f_1(\mathbf{b}) + f_2(\mathbf{b})$, such that it is easier to find good competitive auctions separately against $f_1(\cdot)$ and $f_2(\cdot)$. The following lemma gives an upper bound on the competitive ratio against the original benchmark function $f(\cdot)$.

LEMMA 1.1 (DECOMPOSITION LEMMA). *Let $\mathcal{A}_1$ and $\mathcal{A}_2$ be truthful λ_1 and λ_2 competitive auctions against the benchmarks $f_1(\cdot)$ and $f_2(\cdot)$, respectively. Then there is a truthful $\lambda_1 + \lambda_2$ competitive auction against the benchmark $f_1(\cdot) + f_2(\cdot)$.*

PROOF. We construct an auction that runs $\mathcal{A}_1$ with probability $\frac{\lambda_1}{\lambda_1+\lambda_2}$ and runs $\mathcal{A}_2$ with probability $\frac{\lambda_2}{\lambda_1+\lambda_2}$. The constructed auction is (universally) truthful by definition. Its performance for any bid vector $\mathbf{b}$ is at least

$$\frac{\lambda_1}{\lambda_1+\lambda_2}\mathcal{A}_1(\mathbf{b}) + \frac{\lambda_2}{\lambda_1+\lambda_2}\mathcal{A}_2(\mathbf{b}) \geq \frac{\lambda_1}{\lambda_1+\lambda_2}\frac{f_1(\mathbf{b})}{\lambda_1} + \frac{\lambda_2}{\lambda_1+\lambda_2}\frac{f_2(\mathbf{b})}{\lambda_2} = \frac{f_1(\mathbf{b})+f_2(\mathbf{b})}{\lambda_1+\lambda_2}.$$

□

All our results primarily depend on the analysis of non-standard benchmarks *only* for the basic setting of digital good auctions. Namely, we only consider benchmarks of the form $f(\mathbf{b}) = \max(k_2 \cdot b_{(2)}, k_3 \cdot b_{(3)}, \cdots, k_n \cdot b_{(n)})$, where $k_1, \ldots, k_n$ is a set of constants[1]:

limited supply. $f(\mathbf{b}) = l * b_{(l+1)}$,
downward-closed. $f(\mathbf{b}) = b_{(2)}$,
online auctions. $f(\mathbf{b}) = \max(4b_{(2)}, 3b_{(3)}, 4b_{(4)} \cdots, nb_{(n)})$,

whereas our remaining mechanisms are appropriate adjustments to the existing mechanism from the prior literature.

Specifically for the online auctions, Koutsoupias and Pierrakos [24] gave a lower bound of 4 on the competitive ration of any mechanism with just $n = 2$ bidders. They conjectured that competitive ratio of 4 is tight. Our analysis provides matching upper bound of 4 for the cases of $n = 2, 3, 4$ bidders yielding a straightforward mechanism: the auctioneer posts a price for every new arriving bidder equal to the maximal bid seen so far. However, this mechanism does not extend to the case of more than four bidders and, moreover, we believe that in general case the upper bound of 4 is wrong and we conjecture that the right lower bound is in fact 4.12.

1.3. Related Work

The worst-case study of digital goods auctions was initiated by Goldberg et al. [16]. The competitive prior-free framework was formulated by Fiat et al. [12]. Over the past

[1] some preliminary results for this class of benchmarks were derived in [22]

decade a lot of effort has been devoted to improving the analysis and competitive ratios of digital goods auctions, see, e.g., [13], [14], [22], [11], [1] and [23]. In our recent work [7], we showed the optimal bound on the competitive ratio for digital goods auctions.

A few other closely related settings have stemmed from the study of digital goods auctions with the most immediate extension being the limited supply auctions also known as multi-unit auctions. Multi-unit environments have been traditionally studied with respect to the $\mathcal{F}^{(2,\ell)}$ benchmark, which allows a straightforward reduction [15] to the unlimited supply case with a specific number of bidders. Thus optimal bounds of [7] carry over to the multi-unit auctions with respect to the $\mathcal{F}^{(2,\ell)}$ benchmark.

The general downward-closed single-parameter environments include, e.g,. matching, matroid, and position auctions have also received considerable attention in recent years. Hartline and Yan [21] characterized the optimal revenue in the envy-free outcomes and proposed $\text{EFO}^{(2)}$ as a uniform benchmark for all of these environments. They presented a truthful multi-unit auction with a constant competitive ratio and established a no-loss reduction from position and matroid auctions to a simpler multi-unit setting. Devanur et al. [10] improved the competitive ratio to 9.6 and gave a 189-competitive auction for the more general downward-closed environments. Ha and Hartline [19] further improved the competitive ratio to 30.4 for the downward-closed environments. In an unpublished followup paper [18], the authors presented a 11-competitive auction using elegant combination of biased sampling and profit extraction ideas. The best known ratio is due to Devanur at al. [9] (official version of [18]) with a 7.5-competitive auction that builds upon the biased sampling approach in a significantly more intricate manner than in [18] .

As a multi-parameter extension of the digital goods auctions setting, Gravin and Lu [17] studied competitive auction in the presence of positive externalities among the buyers.

Another thread of work considers digital goods auctions is in the online framework. Motivated by internet advertising, Mahdian and Saberi [25] proposed a model where supply is unknown in advance. Devanur and Hartline [8] studied prior-free auctions in this model and by applying random sampling technique derived results in the prior-free setting. There was substantial interest from machine learning community [5], [6],[3] in a closely related online pricing problem. However, this work together with an earlier work [4] on online auctions does not assume random order of arrivals. It also uses machine learning techniques resulting in a worse performance guarantees that depends on h, the ratio between the highest and the lowest bid. Lastly, the setting of Koutsoupias and Pierrakos [24] is closely related to generalized secretary problem (for a survey on secretary problem and online digital goods auction see [2]). They gave a black-box reduction of the online problem to the standard off-line digital goods setting with a factor 2 loss in the competitive ratio. It should be noted that in [24] the choice of the offline competitive mechanism as long as it is constant approximation to the offline optimum does not matter. In contrast, the choice of the corresponding offline auction is important and non trivial part of our mechanism, since our goal is to get mechanism performance as close to the theoretical optimum as possible.

2. LIMITED SUPPLY AUCTIONS

It was pointed out in [15] that there is an equivalence between the unlimited supply auction problem for the $\mathcal{F}^{(2)}$ benchmark and the limited supply auction problem for the $\mathcal{F}^{(2,\ell)}$ benchmark. Namely, any unlimited supply auction with ℓ bidders that is β-competitive against $\mathcal{F}^{(2)}$ can be converted into a β-competitive ℓ-unit auction against $\mathcal{F}^{(2,\ell)}$. This equivalence and the tight results of [7] for unlimited supply auctions with

ℓ bidders against $\mathcal{F}^{(2)}$ benchmark yield tight results for ℓ-unit auction against $\mathcal{F}^{(2,\ell)}$ benchmark with the same competitive ratio λ_ℓ.[2]

A similar equivalence was established in [7] between ℓ-unit auctions competing with any benchmark $f(\cdot)$ that depends only on the first ℓ highest bids and unlimited supply auctions with ℓ bidders. However, beyond such benchmarks [7] does not provide a satisfactory way to compute tight competitive ratios in the limited supply case. For example, economically meaningful benchmark EFO$^{(2)}$ of Hartline and Yan [21] depends not only on the first ℓ highest bids.

Definition 2.1. For a fixed valuation profile $\mathbf{v}$, order all valuations by $v_{(1)} \geq v_{(2)} \geq \cdots \geq v_{(n)}$ and let $g(j) = j \cdot v_{(j)}$ for each $2 \leq j \leq n$. Consider the concave envelope $\widehat{g}(\cdot)$ of the function $g(\cdot)$ on the interval $[2, n]$, i.e., iron $g(\cdot)$. For ℓ-unit auction $\text{EFO}^{(2)}(\mathbf{v}) = \max\limits_{2 \leq i \leq \ell} \widehat{g}(i)$.

There is only a constant gap between EFO$^{(2)}$ and $\mathcal{F}^{(2,\ell)}$ benchmarks for ℓ-unit auctions.

LEMMA 2.2. *For any valuation profile* $\mathbf{v}$*,*

$$\mathcal{F}^{(2,\ell)}(\mathbf{v}) \leq \textbf{\textit{EFO}}^{(2)}(\mathbf{v}) \leq \mathcal{F}^{(2,\ell)}(\mathbf{v}) + (\ell - 2) \cdot v_{(\ell+1)}.$$

PROOF. The first inequality holds because

$$\mathcal{F}^{(2,\ell)}(\mathbf{v}) = \max_{2 \leq i \leq \ell} g(i) \leq \max_{2 \leq i \leq \ell} \widehat{g}(i) = \text{EFO}^{(2)}(\mathbf{v}).$$

We next prove the second inequality. We assume that $g(\cdot)$ is ironed from i to j, where $i < \ell < j$ and $g(i) < g(j)$, otherwise $\text{EFO}^{(2)}(\mathbf{v}) = \mathcal{F}^{(2,\ell)}(\mathbf{v})$ and the second inequality holds true. Similarly, if $\text{EFO}^{(2)}(\mathbf{v}) > \mathcal{F}^{(2,\ell)}(\mathbf{v})$, we can assume that $\text{EFO}^{(2)}(\mathbf{v}) = \widehat{g}(\ell)$. We observe that $v_{(i)}$ the slope of the line from $(0,0)$ to $(i, g(i))$ is greater than or equal to $v_{(j)}$ the slope of the line from $(0,0)$ to $(j, g(j))$, which implies that the latter slope is greater than or equal to the slope of the line between $(i, g(i))$ and $(j, g(j))$. We further note that $v_{(\ell+1)} \geq v_{(j)}$ is greater than or equal to the slope of the line between $(i, g(i))$ and $(j, g(j))$. Finally, since $\ell - 2 \geq \ell - i$ we get that

$$\mathcal{F}^{(2,\ell)}(\mathbf{v}) + (\ell - 2) \cdot v_{(\ell+1)} \geq g(i) + (\ell - i) \cdot v_{(\ell+1)} \geq \widehat{g}(\ell) = \text{EFO}^{(2)}(\mathbf{v}),$$

which concludes the proof. $\square$

One can further estimate $\mathcal{F}^{(2,\ell)}(\mathbf{v}) \geq (\ell-1) \cdot v_{(\ell)} \geq (\ell-2) \cdot v_{(\ell+1)}$; this implies a trivial upper bound of $2\mathcal{F}^{(2,\ell)}(\mathbf{v})$ on $\text{EFO}^{(2)}(\mathbf{v})$. As λ_ℓ is the exact competitive ratio against the $\mathcal{F}^{(2,\ell)}$ benchmark, the competitive ratio against $\text{EFO}^{(2)}(\mathbf{v})$ lies between λ_ℓ and $2\lambda_\ell$. These two bounds were the best currently known [10]. However, these bounds are not tight. In particular, we can improve on the upper bound.

We decompose the upper bound on EFO$^{(2)}$ in Lemma 2.2 into the sum of two benchmarks $f_1(\mathbf{v}) = \mathcal{F}^{(2,\ell)}(\mathbf{v})$ and $f_2(\mathbf{v}) = (\ell - 2) \cdot v_{(\ell+1)}$. The competitive ratio against the first benchmark is λ_ℓ. On the other hand, the revenue of VCG mechanism selling ℓ items is $\ell \cdot v_{(\ell+1)}$, which shows that the competitive ratio against $f_2(\mathbf{v})$ is $\frac{\ell-2}{\ell}$. By combining Lemma 1.1 and Lemma 2.2 we obtain the following claim, which improves the upper bound on the competitive ratio against EFO$^{(2)}$ to $\lambda_\ell + 1$.

[2] $\lambda_\ell = 1 - \sum\limits_{i=2}^{\ell} \frac{i}{i-1} \binom{\ell-1}{i-1} \left(\frac{-1}{\ell}\right)^{i-1}$, which converges to 2.42 when ℓ approaches infinity.

THEOREM 2.3. *For multi-unit auctions with ℓ units for sale, there is a $(\lambda_\ell + \frac{\ell-2}{\ell})$-competitive auction against the EFO$^{(2)}$ benchmark, where λ_ℓ is the optimal competitive ratio of unlimited supply auction with ℓ bidders against the $\mathcal{F}^{(2)}$ benchmark.*

3. DOWNWARD-CLOSED ENVIRONMENTS

In this section, we consider general downward-closed permutation environments. We denote by $\text{EFO}(\mathbf{v})$ the optimal revenue achievable in an envy-free allocation for the vector of values $\mathbf{v}$. Our benchmark of interest is $\text{EFO}^{(2)}(\mathbf{v}) = \text{EFO}(v_2, v_2, \ldots, v_n)$. The basic ingredients of our auction are biased random sampling and the profit extraction auction from [18]. Our auction is slightly different from the one presented in [18] and has a better competitive ratio of 6.51 compared to 11 of [18]. It is much simpler than another auction with competitive ratio 7.5 presented in [9], which has a few more components on the top of random sampling and profit extraction.

The profit extraction (PER) auction receives as a parameter a target valuation profile $\widetilde{\mathbf{v}}$. When $\text{PER}^{\widetilde{\mathbf{v}}}(\mathbf{v})$ is run on the actual valuation profile $\mathbf{v}$, it is able to extract revenue greater than or equal to the value of the envy-free benchmark $\text{EFO}(\widetilde{\mathbf{v}})$ as well as $\text{EFO}^{(2)}(\widetilde{\mathbf{v}})$, if profile $\mathbf{v}$ dominates $\widetilde{\mathbf{v}}$ ($\mathbf{v} \succeq \widetilde{\mathbf{v}}$), i.e., $v_{(i)} \geq \widetilde{v}_{(i)}$ for every bidder i. If $\mathbf{v} \not\succeq \widetilde{\mathbf{v}}$, PER rejects all bidders.

LEMMA 3.1 (HA AND HARTLINE, 2012). *For any downward-closed permutation environment, there is a truthful profit extraction auction PER$^{\widetilde{\mathbf{v}}}(\mathbf{v})$ with a profit of at least EFO$^{(2)}(\widetilde{\mathbf{v}})$, if $v_{(i)} \geq \widetilde{v}_{(i)}$ for each bidder i.*

Our auction is quite simple: with some probability p we run the single-item Vickrey auction; with probability $1 - p$ we run the following σ-biased random-sampling profit-extraction auction (denoted by σ-BSPE).

- Divide all bidders into two groups market M and sample S: Place the *two highest* bids in M. Sample the rest bids independently with probability $\sigma < 1/2$ in S and with probability $1 - \sigma$ in M.
- Let $\widetilde{\mathbf{v}} = \mathbf{v}_S$. Allocate to the winners of $\text{PER}^{\widetilde{\mathbf{v}}}(\mathbf{v}_M)$.

THEOREM 3.2. *For any downward-closed permutation environment, σ-BSPE is a 6.51-competitive truthful auction against the envy-free benchmark EFO$^{(2)}(\mathbf{v})$.*

PROOF. For any random coin flips of σ-BSPE, the allocation rule of $\text{PER}^{\widetilde{\mathbf{v}}}$ is monotone. This implies that σ-BSPE also has a monotone allocation rule. Therefore, since our environment is a single-parameter domain, σ-BSPE allocation with the threshold payment rule makes the auction universally truthful.

We next estimate the expected revenue of σ-BSPE. We follow closely the proof strategy described in [18], the main difference being in the way we deal with the benchmark $\text{EFO}^{(2)}(\mathbf{v})$. We note that if $\mathbf{v}_M \succeq \mathbf{v}_S$, then the total sum of the threshold payments of $\text{PER}^{\widetilde{\mathbf{v}}}(\mathbf{v}_M)$ is at least $\text{EFO}(\mathbf{v}_S)$; we further observe that the threshold payments of σ-BSPE can be only larger than that, as we could only increase payments of the two highest bidders.

CLAIM 3.3. *The probability that $\mathbf{v}_M \succeq \mathbf{v}_S$ is at least $1 - (\frac{\sigma}{1-\sigma})^3$.*

PROOF. Sort all bidders in the original profile $\mathbf{v} : v_{(1)} \geq \cdots \geq v_{(n)}$ (without loss of generality we assume that all inequalities are strict). We simulate our random sampling process by independently flipping a biased coin for each bidder (i) in this order. Each time we count the difference between the number of bidders in M and S. Note that because we always place the highest two bids in M, after the first two steps the difference becomes two. Note that $\mathbf{v}_M \not\succeq \mathbf{v}_S$ if and only if at some step (i) this difference becomes negative. We next estimate the probability that this event never happens.

We consider an infinite random walk on a one-dimensional infinite line; each time we move to the left with probability σ and to the right with probability $1-\sigma$. It is well known that the probability that such a random walk starting at a point x eventually makes one step to the left from x is $\frac{\sigma}{1-\sigma}$. As our random walk starts at point 2, it would take three such steps to move below 0. The probability of this event is $(\frac{\sigma}{1-\sigma})^3$. Therefore, the probability that this never happens after n steps is at least $1-(\frac{\sigma}{1-\sigma})^3$. $\square$

We conclude that the expected revenue of the σ-BSPE is at least

$$\begin{aligned}
\mathbf{E}\left[\sigma\text{-BSPE}\right] &\geq \mathbf{E}\left[\mathrm{EFO}(\mathbf{v}_S) \mid \mathbf{v}_M \succeq \mathbf{v}_S\right] \cdot \mathbf{Pr}\left[\mathbf{v}_M \succeq \mathbf{v}_S\right] \\
&= \mathbf{E}\left[\mathrm{EFO}(\mathbf{v}_S)\right] - \mathbf{E}\left[\mathrm{EFO}(\mathbf{v}_S) \mid \mathbf{v}_M \not\succeq \mathbf{v}_S\right] \cdot \mathbf{Pr}\left[\mathbf{v}_M \not\succeq \mathbf{v}_S\right] \\
&\geq \sigma \cdot \mathrm{EFO}(\mathbf{v}_{-\{1,2\}}) - \mathrm{EFO}(\mathbf{v}_{-\{1,2\}}) \cdot \mathbf{Pr}\left[\mathbf{v}_M \not\succeq \mathbf{v}_S\right] \\
&\geq \left(\sigma - \left(\frac{\sigma}{1-\sigma}\right)^3\right) \mathrm{EFO}(\mathbf{v}_{-\{1,2\}}),
\end{aligned}$$

where $\mathbf{v}_{-\{1,2\}}$ is the bid vector without first two highest bids. The maximum of the function $(\sigma - (\frac{\sigma}{1-\sigma})^3)$ is attained at $\sigma \approx 0.29$ with a value around 0.22. Thus, the competitive ratio of σ-BSPE against the benchmark $\mathrm{EFO}(\mathbf{v}_{-\{1,2\}})$ is 4.51.

On the other hand, by running the single-item Vickrey auction, we extract revenue of at least $\frac{1}{2} \cdot \mathrm{EFO}(v_2, v_2)$. Note that by subadditivity of EFO (shown in [21]) we have $\mathrm{EFO}(v_2, v_2) + \mathrm{EFO}(\mathbf{v}_{-\{1,2\}}) \geq \mathrm{EFO}(v_2, v_2, \mathbf{v}_{-\{1,2\}}) = \mathrm{EFO}^{(2)}(\mathbf{v})$. Therefore, according to Lemma 1.1 one can achieve the competitive ratio of $4.51 + 2 = 6.51$ against the benchmark $\mathrm{EFO}(v_2, v_2) + \mathrm{EFO}(\mathbf{v}_{-\{1,2\}})$. Thus, we obtain a 6.51-competitive auction against $\mathrm{EFO}^{(2)}(\mathbf{v})$, which runs 0.22-BSPE with probability $4.51/6.51$ and the single-item Vickrey auction with probability $2/6.51$. $\square$

Remark 3.4. In fact, our analysis also implies the same competitive ratio of 6.51 for a slightly stronger benchmark $\mathrm{EFO}(2v_2, v_2, \dots, v_n)$[3]. Indeed, single-item Vickrey auction in the proof of Theorem 3.2 generates revenue equal to $\frac{1}{2} \cdot \mathrm{EFO}(2v_2, v_2)$.

4. ONLINE AUCTIONS

Let $\{\mathcal{M}_{\text{off}}^n\}_{n=2}^{\infty}$ be a sequence of β-competitive offline digital goods auctions against a benchmark $f(\cdot)$ for each number of bidders n. To simplify notation, we refer $\{\mathcal{M}_{\text{off}}^n\}_{n=2}^{\infty}$ as $\mathcal{M}_{\text{off}}$ auction omitting the number of bidders when it could be inferred from the context.

THEOREM 4.1 (KOUTSOUPIAS, PIERRAKOS [24]). *Let $\mathcal{M}_{\mathit{off}}$ be a β-competitive offline auctions against the $\mathcal{F}^{(2)}$ benchmark. The online sampling auction is a 2β-competitive against the $\mathcal{F}^{(2)}$ benchmark with bidders arriving in a random order.*

The online sampling auction by [24] uses a black box reduction from an offline digital-goods auction $\mathcal{M}_{\text{off}}$ to construct an online competitive auction $\mathcal{M}_{\text{on}}$. Their auction, upon the arrival of each bidder k, observes first $k-1$ bids $\mathbf{b}^{[k-1]} \triangleq (b_1, \dots, b_{k-1})$ and runs $\mathcal{M}_{\text{off}}^k(\mathbf{b}^{[k-1]})$ for bidder k.

[3]this benchmark was proposed by Hartline [20] as an upper bound on the revenue of the optimal auction for any i.i.d. prior in any downward-closed environment.

In particular, [24] used the offline auction of [22] with a competitive ratio of 3.24, they obtained an upper bound of 6.48. In a recent paper [7], it was shown that the optimal competitive ratio of $\mathcal{M}_{\text{off}}$ is in fact 2.42, which gives an upper bound of 4.84 for the online problem. There is also a lower bound of 4 in [24] for online auctions with only 2 bidders.

COROLLARY 4.2. *The optimal competitive ratio of online auctions is between* 4 *and* 4.84*.*

We next propose another simple black-box reduction from offline to online auctions with a better competitive guarantee. Any online auction can be thought of as a sequence of offline auctions run for a set of bidders already present at each time. The main idea of our design is to tailor each of our offline auctions to a different from $\mathcal{F}^{(2)}$ benchmark so that their combination has good performance with respect to $\mathcal{F}^{(2)}$.

THEOREM 4.3. *Let* $f(\mathbf{b}) = \max(4b_2, 3b_3, 4b_4, \ldots, kb_k)$ *and* $\mathcal{M}_{\mathit{off}}$ *be a* β*-competitive auction against the* $f(\cdot)$ *benchmark. Then there is a* β*-competitive online auction against* $\mathcal{F}^{(2)}$*, where bidders arrive in a random order.*

PROOF. Any truthful online auction $\mathcal{M}_{\text{on}}$ can be viewed as a weighted combination of offline auctions $\{A^n\}_{n=2}^{\infty}$ running on $n = 1, 2 \ldots$ bidders

$$\mathcal{M}_{\text{on}}(\mathbf{b}) = \sum_{n=1} \frac{1}{n} \cdot A^n(\mathbf{b}^{[n]}).$$

Indeed, each time when $\mathcal{M}_{\text{on}}$ observes first $n-1$ bids $\mathbf{b}^{[n-1]}$ and offers a price to bidder n, it could have seen any combination of $n-1$ bids among $\mathbf{b}^{[n]}$ equally likely, since bidders arrive uniformly at random. Therefore, $\mathcal{M}_{\text{on}}$ derives $\frac{1}{n}$ of the revenue of offline auction $A^n(\mathbf{b}^{[n]})$. We denote by $\mathcal{M}_{\text{on}}^k$ the online auction run only up to k rounds, i.e.,

$$\mathcal{M}_{\text{on}}^k(\mathbf{b}) = \sum_{n=1}^{k} \frac{1}{n} \cdot A^n(\mathbf{b}^{[n]}).$$

We are going to construct our online auction $\mathcal{M}_{\text{on}}$ inductively at each time increasing the number of bidders by one. Namely, we assume that for all $n = k-1$ bidders our $\mathcal{M}_{\text{on}}$ auction is β-competitive. Next we specify an offline auction A^k which together with $\mathcal{M}_{\text{on}}^{k-1}$ is β-competitive for k bidders.

For $n = 1$ bidder $\mathcal{F}^{(2)}$ is 0, so $\mathcal{M}_{\text{on}}$ is competitive regardless of A^1. By induction hypothesis, we know that for any fixed bid vector $\mathbf{b}^{[k-1]}$,

$$\beta \cdot \mathcal{M}_{\text{on}}^{k-1}(\mathbf{b}^{[k-1]}) \geq \mathcal{F}^{(2)}(\mathbf{b}^{[k-1]}).$$

Since the first $k-1$ bids are uniformly selected from $\mathbf{b}^{[k]}$, we have

$$\beta \cdot \mathcal{M}_{\text{on}}^{k-1}(\mathbf{b}^{[k]}) = \frac{1}{k} \sum_{i=1}^{k} \beta \cdot \mathcal{M}_{\text{on}}^{k-1}(\mathbf{b}_{\text{-}i}^{[k]}) \geq \frac{1}{k} \sum_{i=1}^{k} \mathcal{F}^{(2)}(\mathbf{b}_{\text{-}i}^{[k]}).$$

Let us sort the bids in $\mathbf{b}^{[k]}$ by $b_1 \geq b_2 \geq \cdots \geq b_k$. We compare the revenue of $\mathcal{M}_{\text{on}}^{k-1}$ with each b_ℓ, for $2 \leq \ell \leq k$. We have $\mathcal{F}^{(2)}(\mathbf{b}_{\text{-}i}^{[k]}) \geq \ell \cdot b_\ell$ for every $i > \ell$. If $\ell > 2$ and $i \leq \ell$, then $\mathcal{F}^{(2)}(\mathbf{b}_{\text{-}i}^{[k]}) \geq (\ell - 1) \cdot b_\ell$. Unfortunately, for $\ell = 2$, we cannot write $\mathcal{F}^{(2)}(\mathbf{b}_{\text{-}1}^{[k]}) \geq b_2$ or

$\mathcal{F}^{(2)}(\mathbf{b}^{[k]}_{-2}) \geq b_2$. Therefore,

$$\begin{aligned}\frac{1}{k}\sum_{i=1}^{k}\mathcal{F}^{(2)}(\mathbf{b}^{[k]}_{-i}) \;&\geq\; \max\left(\frac{2k-4}{k}b_2, \frac{3k-3}{k}b_3, \ldots, \frac{\ell(k-\ell)+(\ell-1)\ell}{k}b_\ell, \ldots, \frac{k^2-k}{k}b_k\right)\\ &\geq\; \mathcal{F}^{(2)}(\mathbf{b}^{[k]}) - \frac{1}{k}\max(4b_2, 3b_3, 4b_4, \ldots, kb_k).\end{aligned}$$

We want the offline auction A^k to have good performance against $f(\mathbf{b}) = \max(4b_2, 3b_3, 4b_4, \ldots, kb_k)$. We know that there is a β-competitive auction $\mathcal{M}_{\text{off}}$ with respect to this benchmark $f(\cdot)$. Thus, there is a β-competitive auction for k bidders in the online setting. This completes the proof. $\square$

Note that $f(\mathbf{b}) \leq \mathcal{F}^{(2)}(\mathbf{b}) + 2b_2$. According to Lemma 1.1, we can run a mixture of the optimal auction against $\mathcal{F}^{(2)}$ and single-item Vickery auction against $2b_2$ to achieve a $(\lambda+2)$-competitive auction with respect to $f(\cdot)$, which is already an improvement over the result of [24]. However, we can actually derive the optimal ratio using the same approach as that for $\mathcal{F}^{(2)}$, which yields an even better competitive ratio for the online auction problem.

THEOREM 4.4. *The optimal competitive ratio of (offline) digital good auction with respect to the benchmark* $f(\mathbf{b}) = \max(4b_2, 3b_3, 4b_4, \ldots, nb_n)$ *is at most* 4.12.[4]

PROOF. By the same argument as in [7] for $\mathcal{F}^{(2)}$, the matching lower bound for the optimal competitive ratio is achieved by the equal revenue distribution with the support $\mathbb{R}^n_{\geq 1}$. For $n \leq 4$, $f(\mathbf{b}) = 4b_2$. In the following, we always assume $n > 4$. We first observe that $f(\mathbf{b}) = \max(4b_2, \mathbf{EFO}^{(2)}(\mathbf{b}))$.

We recall that equal revenue distribution $\mathcal{D}^n$ over the bid vectors is i.i.d. with the density function $w(b) = \frac{1}{b^2}$ and cumulative density $1 - \frac{1}{b}$ supported on $[1, \infty)$. Let $\mathbf{B}$ be a random vector drawn from $\mathcal{D}^n$. The key technical problem for us is to compute the expected value of the benchmark $f(\mathbf{B})$. Following [14], we compute the probability $\mathbf{Pr}[f(\mathbf{B}) \geq z]$ for any given z. Since $f(\mathbf{B})$ is at least n, we may only consider $z \geq n$. Let a random variable V_i be the i-th largest bid in $\mathbf{B}$. We also define a set of random variables

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

Let $\mathcal{H}_i$ denote the event

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

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

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

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

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

[4]the actual ratio for a fixed n is $1 - \sum_{i=2}^{n}\left(\frac{-1}{n}\right)^{i-1}\frac{i}{i-1}\binom{n-1}{i-1} + \frac{3n}{2(n-2)}\left(\left(1-\frac{2}{n}\right)^{n-1} + 1 - \frac{2}{n}\right)$.

This gives a recursive relation for $\mathbf{Pr}[F_{n,k} \geq z]$ and the boundary condition is $\mathbf{Pr}[F_{0,k} \geq z] = 0$. This recursion has been solved in [14]:

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

Let $\mathcal{H}_2'$ denote the event $V_2 \geq \frac{z}{4}$ and $\bigwedge_{j=3,4,\ldots,n} V_j < \frac{z}{j}$. Then

$$\mathbf{Pr}\left[\mathcal{H}_2'\right] = \binom{n}{2}\left(\frac{4}{z}\right)^2 \mathbf{Pr}\left[F_{n-2,2} < z\right].$$

This implies that

$$\begin{aligned}
&\mathbf{Pr}\left[f(\mathbf{B}) \geq z\right] \\
&= \mathbf{Pr}\left[\mathcal{H}_2'\right] + \sum_{j=3,4,\ldots,n} \mathbf{Pr}\left[\mathcal{H}_j\right] \\
&= \mathbf{Pr}\left[\mathcal{H}_2'\right] + \mathbf{Pr}\left[F_{n,0} \geq z\right] - \mathbf{Pr}\left[\mathcal{H}_1\right] - \mathbf{Pr}\left[\mathcal{H}_2\right] \\
&= \binom{n}{2}\left(\frac{4}{z}\right)^2 \mathbf{Pr}\left[F_{n-2,2} < z\right] + \frac{n}{z} - \frac{n}{z}\mathbf{Pr}\left[F_{n-1,1} < z\right] - \binom{n}{2}\left(\frac{2}{z}\right)^2 \mathbf{Pr}\left[F_{n-2,2} < z\right] \\
&= \frac{n}{z} - \frac{n}{z}\left(\frac{z-1}{z}\right)^{n-1}\left(\frac{z-n}{z-1}\right) + \frac{6n(n-1)}{z^2}\left(\frac{z-2}{z}\right)^{n-2}\left(\frac{z-n}{z-2}\right)
\end{aligned}$$

Therefore, we have

$$\begin{aligned}
\mathbf{E}\left[f(\mathbf{B})\right] &= \int_0^\infty \mathbf{Pr}\left[f(\mathbf{B}) \geq z\right] dz \\
&= n + \int_n^\infty \left(\frac{n}{z} - \frac{n}{z}\left(\frac{z-1}{z}\right)^{n-1}\left(\frac{z-n}{z-1}\right) + \frac{6n(n-1)}{z^2}\left(\frac{z-2}{z}\right)^{n-2}\left(\frac{z-n}{z-2}\right)\right) dz \\
&= n - n\sum_{i=2}^{n}\left(\frac{-1}{n}\right)^{i-1}\frac{i}{i-1}\binom{n-1}{i-1} + 6n(n-1)\int_n^\infty \frac{1}{z^2}\left(\frac{z-2}{z}\right)^{n-2}\left(\frac{z-n}{z-2}\right) dz.
\end{aligned}$$

The integration part is

$$\int_n^\infty \frac{1}{z^2}\left(\frac{z-2}{z}\right)^{n-2}\left(\frac{z-n}{z-2}\right)dz$$
$$=\int_n^\infty \frac{1}{z^2}(1-\frac{n}{z})\left(\frac{z-2}{z}\right)^{n-3}dz$$
$$=\int_n^\infty \frac{1}{z^2}(1-\frac{n}{z})\left(\sum_{i=0}^{n-3}\binom{n-3}{i}\frac{(-2)^i}{z^i}\right)dz$$
$$=\sum_{i=0}^{n-3}\binom{n-3}{i}(-2)^i\int_n^\infty\left(\frac{1}{z^{i+2}}-\frac{n}{z^{i+3}}\right)dz$$
$$=\sum_{i=0}^{n-3}\binom{n-3}{i}(-2)^i\left(\frac{1}{(i+1)n^{i+1}}-\frac{1}{(i+2)n^{i+1}}\right)$$
$$=\sum_{i=0}^{n-3}\binom{n-3}{i}(-2)^i\frac{1}{(i+1)(i+2)n^{i+1}}$$
$$=\frac{n}{4(n-1)(n-2)}\sum_{i=0}^{n-3}\binom{n-1}{i+2}\left(\frac{-2}{n}\right)^{i+2}$$
$$=\frac{n}{4(n-1)(n-2)}\left(\left(1-\frac{2}{n}\right)^{n-1}-1-\binom{n-1}{1}\frac{-2}{n}\right)$$
$$=\frac{n}{4(n-1)(n-2)}\left(\left(1-\frac{2}{n}\right)^{n-1}+1-\frac{2}{n}\right)$$

Therefore, we have

$$\mathbf{E}\left[f(\mathbf{B})\right]=n-n\sum_{i=2}^{n}\left(\frac{-1}{n}\right)^{i-1}\frac{i}{i-1}\binom{n-1}{i-1}+\frac{3n^2}{2(n-2)}\left(\left(1-\frac{2}{n}\right)^{n-1}+1-\frac{2}{n}\right).$$

And the competitive ratio is

$$1-\sum_{i=2}^{n}\left(\frac{-1}{n}\right)^{i-1}\frac{i}{i-1}\binom{n-1}{i-1}+\frac{3n}{2(n-2)}\left(\left(1-\frac{2}{n}\right)^{n-1}+1-\frac{2}{n}\right).$$

□

As $f(\mathbf{b})=4b_2$ for $n=2,3$, and 4 bidders we get competitive ratio of 4, which exactly matches the lower bound. Therefore, our online auction is optimal for the case of 2, 3, and 4 bidders.

4.1. The benchmark MAXV

The results of [24] carry over for another standard benchmark, namely, the maximum Vickery MAXV. As the exact competitive ratio of the optimal offline auction against MAXV was shown in [7] to be $e-1$ and since $2\text{MAXV}(\mathbf{v})=\mathcal{F}^{(2,\ell)}(\mathbf{v})$ for $n=2$ bidders, the approach of [24] implies the following claim.

THEOREM 4.5 (KOUTSOUPIAS ET AL.[24]). *The competitive ratio of the online sampling auction of [24] is at most $2(e-1)$ against the* MAXV *benchmark. The competitive ratio of any online auction against* MAXV *is at least* 2.

Interestingly, if we run $\mathcal{M}_{\text{on}}$ exactly as proposed in [24], i.e., as a sequence of $\mathcal{M}_{\text{off}}$ tailored to $\mathcal{F}^{(2)}$, then $\mathcal{M}_{\text{on}}$ appears to be specifically well suited for the MAXV benchmark. This observation once again highlights how useful is the idea of thinking about the problem with respect to different benchmarks.

THEOREM 4.6. *Let $\mathcal{M}_{\mathit{off}}$ be a β-competitive auction against the $\mathcal{F}^{(2)}$ benchmark. The online sampling auctions composed of a sequence of offline auctions $\mathcal{M}_{\mathit{off}}$ against $\mathcal{F}^{(2)}$ is β-competitive against* MAXV.

PROOF. Similar to the proof of Theorem 4.3, we proceed by induction on the number of bidders. We have

$$\beta \cdot \mathcal{M}_{\text{on}}^{k-1}(\mathbf{b}^{[k]}) \geq \frac{1}{k}\sum_{i=1}^{k} \text{MAXV}(\mathbf{b}_{\text{-}i}^{[k]}).$$

We sort the bids in $\mathbf{b}^{[k]} : b_1 \geq \cdots \geq b_k$. For a fixed ℓ, we want to estimate how the revenue of $\mathcal{M}_{\text{on}}^{k-1}$ is compared to b_ℓ. For each $i > \ell$ we have $\mathcal{F}^{(2)}(\mathbf{b}_{\text{-}i}^{[k]}) \geq (\ell-1)\cdot b_\ell$; and for $i \leq \ell$, we have $\mathcal{F}^{(2)}(\mathbf{b}_{\text{-}i}^{[k]}) \geq (\ell-2)\cdot b_\ell$. Hence,

$$\begin{aligned}
\frac{1}{k}\sum_{i=1}^{k}\mathcal{F}^{(2)}(\mathbf{b}_{\text{-}i}^{[k]}) \;&\geq\; \max\left(\frac{k-2}{k}b_2,\ldots,\frac{(\ell-1)(k-\ell)+(\ell-2)\ell}{k}b_\ell,\ldots,\frac{(k-2)k}{k}b_k\right)\\
&=\; \max\left(\frac{k-2}{k}b_2,\ldots,\frac{(\ell-1)k-\ell}{k}b_\ell,\ldots,\frac{(k-1)k-k}{k}b_k\right)\\
&\geq\; \text{MAXV}(\mathbf{b}) - \frac{1}{k}\mathcal{F}^{(2)}(\mathbf{b}).
\end{aligned}$$

Thus, the online sampling auction by running a β-competitive auction against $\mathcal{F}^{(2)}$ benchmark is β-competitive against MAXV. □

COROLLARY 4.7. *The competitive ratio of online auctions against* MAXV *is between* 2 *and* 2.42*.*

REFERENCES

S. Alaei, A. Malekian, and A. Srinivasan. On random sampling auctions for digital goods. In *EC*, pages 187–196, 2009.

Moshe Babaioff, Nicole Immorlica, David Kempe, and Robert Kleinberg. Online auctions and generalized secretary problems. *SIGecom Exchanges*, 7(2), 2008.

Maria-Florina Balcan, Avrim Blum, Jason D. Hartline, and Yishay Mansour. Reducing mechanism design to algorithm design via machine learning. *J. Comput. Syst. Sci.*, 74(8):1245–1270, 2008.

Ziv Bar-Yossef, Kirsten Hildrum, and Felix Wu. Incentive-compatible online auctions for digital goods. In *SODA*, pages 964–970, 2002.

A. Blum and J. Hartline. Near-optimal online auctions. In *SODA*, pages 1156–1163. SIAM, 2005.

Avrim Blum, Vijay Kumar, Atri Rudra, and Felix Wu. Online learning in online auctions. *Theor. Comput. Sci.*, 324(2-3):137–146, 2004.

Ning Chen, Nick Gravin, and Pinyan Lu. Optimal competitive auctions. In *STOC*, pages 253–262, 2014.

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.

Nikhil R. Devanur, Bach Q. Ha, and Jason D. Hartline. Prior-free auctions for budgeted agents. In *EC*, pages 287–304, 2013.

Nikhil R. Devanur, Jason D. Hartline, and Qiqi Yan. Envy freedom and prior-free mechanism design. *CoRR*, abs/1212.3741, 2012.

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.

A. Fiat, A. V. Goldberg, J. D. Hartline, and A. R. Karlin. Competitive generalized auctions. In *STOC*, pages 72–81, 2002.

A. V. Goldberg and J. D. Hartline. Competitiveness via consensus. In *SODA*, pages 215–222. ACM/SIAM, 2003.

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.

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.

A. V. Goldberg, J. D. Hartline, and A. Wright. Competitive auctions and digital goods. In *SODA*, pages 735–744, 2001.

N. Gravin and P. Lu. Competitive auctions for markets with positive externalities. In *ICALP (2)*, pages 569–580. Springer, 2013.

Bach Q. Ha and Jason D. Hartline. The biased sampling profit extraction auction. *CoRR*, abs/1206.4955, 2012.

Bach Q. Ha and Jason D. Hartline. Mechanism design via consensus estimates, cross checking, and profit extraction. In *SODA*, pages 887–895, 2012.

J. Hartline. Personal communications, 2014.

J. Hartline and Q. Yan. Envy, truth, and profit. In *EC*, pages 243–252. ACM, 2011.

J. D. Hartline and R. McGrew. From optimal limited to unlimited supply auctions. In *EC*, pages 175–182. ACM, 2005.

T. Ichiba and K. Iwama. Averaging techniques for competitive auctions. In *ANALCO*, pages 74–81, 2010.

Elias Koutsoupias and George Pierrakos. On the competitive ratio of online sampling auctions. *ACM Trans. Economics and Comput.*, 1(2):10, 2013.

M. Mahdian and A. Saberi. Multi-unit auctions with unknown supply. In *EC*, pages 243–249. ACM, 2006.