On the Competitive Ratio of Online Sampling Auctions

ELIAS KOUTSOUPIAS, University of Athens
GEORGE PIERRAKOS, University of California, Berkeley

We study online profit-maximizing auctions for digital goods with adversarial bid selection and uniformly random arrivals; in this sense, our model lies at the intersection of prior-free mechanism design and secretary problems. Our goal is to design auctions that are constant competitive with $\mathcal{F}^{(2)}$. We give a generic reduction that transforms any offline auction to an online one with only a loss of a factor of 2 in the competitive ratio. We also present some natural auctions, both randomized and deterministic, and study their competitive ratio. Our analysis reveals some interesting connections of one of these auctions with RSOP.

Categories and Subject Descriptors: F.2.0 [**Analysis of Algorithms and Problem Complexity**]: General; G.2.1 [**Discrete Mathematics**]: Combinatorics

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Mechanism design, digital goods auctions, prior-free mechanisms, online algorithms

ACM Reference Format:
Koutsoupias, E. and Pierrakos, G. 2013. On the competitive ratio of online sampling auctions. *ACM Trans. Econ. Comp.* 1, 2, Article 10 (May 2013), 10 pages.
DOI:http://dx.doi.org/10.1145/2465769.2465775

10

1. INTRODUCTION

The design of auctions that maximize the auctioneer's profit is a well-studied question in mechanism design. Even though most of the economics literature assumes a prior on the distribution of bidders' values and aims at maximizing the *expected* profit [Myerson 1981], the question of designing a profitable auction without any prior assumptions has also received a lot of attention during the past decade [Aggarwal et al. 2005; Dhangwatnotai et al. 2010; Feige et al. 2005; Goldberg et al. 2002; Hartline and Roughgarden 2008]. This is called *prior-free mechanism design* [Goldberg et al. 2002] and it adopts the following worst-case approach: bids are no longer coming from a distribution, but are rather picked by an adversary, and the goal is to design an auction that performs favorably compared to some well-behaved benchmark.

Most of the work in prior-free mechanism design assumes that the bids are known in advance; we consider instead the online setting where bidders arrive one at a time with a random order. In this setting, the design of a profitable, truthful auction reduces to making the "right" take-it-or-leave-it offer to every bidder as she arrives at the auction,

E. Koutsoupias was supported in part by IST-2008-215270 (FRONTS) and ESF-NSRF research program Thales (AGT). G. Pierrakos was supported by NSF grant CCF-0964033 and by a Google University Research Award.
Authors' addresses: E. Koutsoupias, Department of Computer Science, University of Oxford; G. Pierrakos (corresponding author), School of Electrical Engineering and Computer Sciences, University of California at Berkeley; email: georgios@cs.berkeley.edu.

DOI:http://dx.doi.org/10.1145/2465769.2465775

using the previous bids as the only information. We call such auctions *online sampling auctions*.

This model bears a lot of similarities to the *secretary model*: the adversary picks the values of the elements, which are then presented in (uniformly) random order, and we are called to design an algorithm that maximizes the probability of picking the largest element. There is an extensive literature about online auctions and generalized secretary problems (for a survey, see Babaioff et al. [2008]). In the generalized secretary problem, the difficulty for designing an algorithm comes from the restriction to select among a feasible subset of the input; for example, a subset of at most k elements in the k-choice secretary problem, or an independent set in the matroid secretary problem. In contrast, in our work there is no such restriction since all possible subsets are feasible as we consider digital goods; the difficulty arises for setting the appropriate take-it-or-leave-it offer. Combining the two models remains an interesting open problem. For example, the k-choice secretary problem corresponds to online sampling auctions with limited supply. Another difference between the generalized secretary problems and our model is on the objective: in the secretary problem the objective is the social welfare (total value of selected elements), while in our model the objective is the extracted profit.

1.1. The Model

We study auctions of digital goods where bidders arrive online. Formally, we have n bidders with valuations $v_1, \ldots, v_n$ (where we assume $v_1 \geq \ldots \geq v_n$) and n identical items for sale. Bidders arrive in a random order, specified by the function $\pi : [n] \to [n]$, which is a permutation on $[n] = \{1, \ldots, n\}$; we assume a uniform distribution over all different permutations of the n bids and adversarial (worst-case) choice of the values of the bids.

As each bidder arrives, we make her a take-it-or-leave-it offer for a copy of the item at some price p. We want to make the offer before the bidder declares her bid (or equivalently we do not want our offer to depend on her declared bid) so that our auction is *truthful* (i.e., it is in the bidder's best interest to bid her true value v_i); hence, from now on we shall use $b_1, \ldots, b_n$ to refer to both bids and actual values of the players. Formally, we want to make the j-th bidder b_{π_j} an offer $p_j = p(b_{\pi_1}, \ldots, b_{\pi_{j-1}})$; the bidder will accept the offer if $b_{\pi_j} \geq p_j$ and will pay p_j.

Our goal is to maximize the expected profit of our auction, defined as $\mathrm{E}\left[\sum_{j=1}^{n} p_j \cdot \mathrm{I}(b_{\pi_j} \geq p_j)\right]$, where $\mathrm{I}(\cdot)$ is an indicator function. We are going to consider both deterministic and randomized pricing rules $p(b_{\pi_1}, \ldots, b_{\pi_{j-1}})$; therefore the expectation is over all possible orderings of the input bids and—in the case of random pricing—over the randomization in our mechanism.

We are going to use the competitive framework proposed in Goldberg et al. [2002] and compare the expected profit of our auctions to the profit of the *optimal single-price auction* that sells at least two items, namely $\mathcal{F}^{(2)}(b_1, \ldots, b_n) = \max_{i \geq 2} i \cdot b_i$. Notice that although the benchmark $\mathcal{F}^{(2)}$ uses a single price, we impose no such restriction on our auctions. We say that an online auction is ρ-competitive if its expected profit is at least $\mathcal{F}^{(2)}/\rho$. Our goal is to design constant-competitive auctions (i.e., auctions where ρ is a constant).

1.2. Related Work

The work closest in spirit to ours is Hajiaghayi et al. [2004]. This article studies limited-supply online auctions where an auctioneer has k items to sell and bidders arrive and depart dynamically; the analysis assumes worst-case input bids and random arrivals and the main result is an online auction that is constant competitive for both efficiency and revenue. The profit benchmark considered for $k > 1$ items is essentially the same as the one here, namely the optimal single-price sale profit that sells at least two items $\mathcal{F}^{(2)}$. The setting considered in that paper is strictly more general than the one considered here because of the limited supply; moreover that paper also address the issue of possible arrival times misreports. For that more complicated problem they give an auction that is 6,338-competitive. Our auctions achieve much better competitive ratios (below 10), and are arguably simpler to analyze for the particular problem we are interested in.

One of the simplest offline competitive auctions, and arguably the most studied [Alaei et al. 2009; Feige et al. 2005; Goldberg et al. 2002], is the *Random Sampling Optimal Price* auction (RSOP). In RSOP the bidders are uniformly partitioned into two parts, and the optimal single price of each part (i.e., $\arg\max i \cdot b_i$) is offered to the bidders of the other part. RSOP is conjectured to be 4-competitive; to date the best upper bound is 4.68 [Alaei et al. 2009].

Online auctions for digital goods have also been studied before in Bar-Yossef et al. [2002], Blum et al. [2004], Blum and Hartline [2005], and Balcan et al. [2008]. Their model is different from ours in that they do not assume random arrivals. Most of the algorithms presented in these papers are based on techniques from machine learning and their performance depends on h, the ratio of the highest to the lowest bid. Our auctions are arguably more natural and in most cases achieve better competitive ratios; however, in our model, auctions heavily rely on learning the actual values of past bids and not just whether a bidder accepted or rejected the offer (as opposed to some of the auctions in Blum and Hartline [2005]).

1.3. Our Results

Our main result in Section 2.1 is a general black-box reduction that transforms any ρ-competitive offline auction to an online auction while blowing up the competitive ratio by at most a factor of two. Using the best known offline algorithm, this yields a competitive ratio of 6.48 for the problem in hand; we also provide a lower bound of 4 for all randomized online sampling auctions. Next, in Section 2.2, we propose a natural family of deterministic online auctions, the best-price-so-far auctions, and we show that one representative of this family achieves a constant-competitive ratio. We conjecture that its actual competitive ratio is 4, same as the conjectured competitive ratio of RSOP, and we highlight some connections between the two auctions.

2. ONLINE SAMPLING AUCTIONS

2.1. Randomized Competitive Online Sampling Auctions

We start by noting that it is relatively straightforward to implement an auction that achieves constant competitive ratio; indeed the following simple application of the existing literature achieves a competitive ratio equal to twice that of RSOP (i.e., at most 9.36): We start by noticing that RSOP partitions the bidders in two sets of size $(k, n-k)$ where $k \sim \text{Binomial}(n, 1/2)$; let this k be the first k bidders that arrive in our setting, and suppose we offer the optimal price of those k bidders to the remaining $n-k$ bidders. This of course results in the loss of the profit that RSOP extracts from the first k bidders when offering them the optimal price of the remaining $n-k$ bidders; however, because of the random order assumption about the input, it follows that this extracts at least half of RSOP's revenue. Since RSOP is at most 4.68-competitive, it follows that the preceding auction is at most 9.36-competitive. (In fact we can achieve an even better competitive ratio of 8 by using RSPE instead of RSOP, a slightly different auction that relies on the same random partitioning approach; see Goldberg et al. [2002] for more details.) The obvious drawback of this approach is that the number of bidders n must be known in advance; in this section we provide a generic reduction from offline to online auctions, yielding algorithms that do not assume such prior knowledge of n, and with a competitive ratio at most twice as large as that of their offline counterpart.

We first notice that any truthful (offline) auction $\mathcal{A}$ for digital goods has the following format: every bidder i is given a take-it-or-leave-it offer p_i which is a function $f(b_{-i})$ of the bids of the other players; if the bidder accepts she pays p_i, otherwise she pays nothing (this follows from Myerson's characterization [Myerson 1981]). Then we notice that every such truthful offline auction gives rise to the following auction for the online setting, called the *online version* of auction $\mathcal{A}$: simply set the price offered to the j-th arriving bidder to be $p_j = f(b_{\pi_1}, \dots, b_{\pi_{j-1}})$, for the same function f. Intuitively this means that we run the offline auction on the whole set of revealed bids, but actually charge only the bidder that has just arrived; because we restrict our attention to truthful offline auctions, we know that the price offered to p_j will not depend on b_j and so we can offer the j-th bidder a price before she even reveals her bid. Our theorem now says that the resulting online auction has at most twice the competitive ratio of the offline auction.

THEOREM 2.1. *The online version of an offline auction with competitive ratio ρ has competitive ratio at most 2ρ. More precisely, if b_k is the price of the optimal single-price auction, then the competitive ratio of the online auction is at most $\rho \cdot k/(k-1)$.*

PROOF. Consider the first t bids of the online auction. The online auction runs the offline auction on them. The expected profit of the offline auction from this set of t bids is at least $\frac{1}{\rho}\mathcal{F}^{(2)}(b_{\pi_1}, \dots, b_{\pi_t})$; by the random-order assumption about the input, the expected profit from every bid is equal and, in particular, the expected gain from b_{π_t} is at least $\frac{1}{t}\frac{1}{\rho}\mathcal{F}^{(2)}(b_{\pi_1}, \dots, b_{\pi_t})$.

Let $\mathcal{F}^{(2)}(b_1, \dots, b_n) = kb_k$; with probability $\binom{t}{m}\binom{n-t}{k-m}/\binom{n}{k}$ the first t bids have exactly m of the highest k bids which contribute to the optimum. Also, for $m \geq 2$, $\mathcal{F}^{(2)}(b_{\pi_1}, \dots, b_{\pi_t}) \geq mb_k$.[1] So, it follows that when $m \geq 2$, with the previous probability the expected gain of the online auction from b_{π_t} is at least $\frac{1}{t}\frac{1}{\rho}mb_k$.

So, the expected profit of the online auction is at least

$$\sum_{t=2}^{n}\sum_{m=2}^{\min\{t,k\}} \frac{\binom{t}{m}\binom{n-t}{k-m}}{\binom{n}{k}}\frac{1}{t}\frac{1}{\rho}mb_k = \frac{1}{\rho}b_k\binom{n}{k}^{-1}\sum_{t=2}^{n}\sum_{m=2}^{k}\binom{t-1}{m-1}\binom{n-t}{k-m}.$$

The rest of the proof is a technical result that simplifies

$$\binom{n}{k}^{-1}\sum_{t=2}^{n}\sum_{m=2}^{k}\binom{t-1}{m-1}\binom{n-t}{k-m}$$

[1]Notice that when $m = 1$, there is no decent lower bound for $\mathcal{F}^{(2)}$; this is the reason that the online auction has larger competitive ratio than the offline auction.

to $k-1$. Indeed we have that

$$\begin{aligned}
\binom{n}{k}^{-1}\sum_{t=2}^{n}\sum_{m=2}^{k}\binom{t-1}{m-1}\binom{n-t}{k-m} &= \binom{n}{k}^{-1}\sum_{t=1}^{n-1}\sum_{m=1}^{k-1}\binom{t}{m}\binom{n-1-t}{k-1-m},\\
&= \binom{n}{k}^{-1}\sum_{t=1}^{n-1}\left(\binom{n-1}{k-1}-\binom{n-1-t}{k-1}\right),\\
&= \binom{n}{k}^{-1}\left((n-1)\binom{n-1}{k-1}-\sum_{j=k-1}^{n-2}\binom{j}{k-1}\right),\\
&= \binom{n}{k}^{-1}\left((n-1)\binom{n-1}{k-1}-\binom{n-1}{k}\right),\\
&= k-1,
\end{aligned}$$

where in the second equality we used the Chu-Vandermonde identity and in the second-to-last equality we used the identity $\sum_{j=k}^{n}\binom{j}{k}=\binom{n+1}{k+1}$. The theorem now follows, since the expected profit of the online auction is at least

$$(k-1)\frac{b_k}{\rho}=\frac{k-1}{k}\frac{kb_k}{\rho}.$$

□

The competitiveness of online sampling auctions follows by combining the preceding theorem with known guarantees for offline digital goods auctions.

COROLLARY 2.2. *The competitive ratio of online sampling auctions is between 4 and 6.48.*

PROOF. The upper bound is given by the online version of the (offline) auction presented in Hartline and McGrew [2005] which achieves a competitive ratio of 3.24. The lower bound follows from the next lemma, which shows that no randomized online algorithm can have a competitive ratio less than 4. □

LEMMA 2.3. *In the random-order model, no randomized online algorithm has competitive ratio less than 4 against $\mathcal{F}^{(2)}$.*

PROOF. We will show the lemma for two bids. This extends directly to many bids, by padding the 2 bids with other bids of 0 value.

A randomized algorithm for two bids is defined by a cdf $P_1(x)$, which is the probability to offer price at most x to the first bid, and by a cdf $P_2(x|y)$, which is the probability to offer price at most x to the second bid, given that the history is y (y encapsulates both the first price and the first bid). It is not clear how to use Yao's lemma in this case. However, a simple variant which uses Yao's lemma for the second bid suffices: The adversary selects the two bids independently from the equal-revenue distribution with cdf $F_a(x)=1-a/x$, for some a which depends on $P_1(x)$. In other words, the auction designer selects $P_1(x)$, then the adversary selects a that defines a cdf for the two (independent) bids. In the next paragraph we show that for every $P_1(x)$, the adversary can select an a for which the expected profit from the first bid is very small. In particular, for every $\epsilon > 0$ and $P_1(x)$, there is a value a such that the expected online gain from the first bid is at most ϵa. The expected online gain from the second bid is at most a; this is the expected gain from the equal revenue distribution $F_a(x)$ and it is independent of the selected strategy of the algorithm. The total online gain from both bids is at most $(1+\epsilon)a$. On the other hand, the expected value of $\mathcal{F}^{(2)}$ for the equal-revenue distribution $F_a(x)$ is well-known to be exactly $4a$ (see, for example, Goldberg et al. [2002]). The lower bound is therefore $4/(1+\epsilon)$, for every $\epsilon > 0$, which proves the lemma.

It remains to show that there is a such that the expected gain from the first bid is at most ϵa. Indeed, for given a, the expected gain from the first bid is given by

$$\begin{aligned}\int_a^\infty \int_0^b xP_1'(x)\,dx F_a'(b)\,db &= \int_a^\infty \frac{a}{b^2}\int_0^b xP_1'(x)\,dx\,db\\ &= \int_0^\infty \int_{\max(a,x)}^\infty \frac{a}{b^2} xP_1'(x)\,db\,dx\\ &= \int_0^\infty xP_1'(x)\frac{a}{\max(a,x)}\,dx\\ &= \int_0^a xP_1'(x)\,dx + a\int_a^\infty P_1'(x)\,dx\\ &= a - \int_0^a P_1(x)\,dx.\end{aligned}$$

The limit of the expression $(a - \int_0^a P_1(x)\,dx)/a$ as a tends to infinity is the same with the limit of $1 - P_1(a)$ (l'Hôpital's rule). Since the last tends to 0, for every ϵ there is a sufficiently high a for which the expected online gain from the first bid is at most ϵa. The proof of the lemma is complete. □

In a first attempt to bridge the gap between the lower and the upper bound we studied the competitive ratio that can be achieved by the online version of the *Sampling Cost Sharing* auction (SCS); this auction partitions bidders uniformly into two parts and extracts the optimal single-price sale profit of each side from the other (if possible, otherwise it extracts no profit) [Goldberg et al. 2002]. Given that the competitive ratio of SCS is no more than $\rho(k) = \left(\frac{1}{2} - \binom{k-1}{\lfloor k/2 \rfloor}2^{-k}\right)^{-1}$ (as proved in Goldberg et al. [2002]), where k is the number of winners, it is tempting to conclude that the online version of SCS will have competitive ratio at most $\frac{k}{k-1}\rho(k)$; unfortunately this is not true because Theorem 2.1 gives guarantees only in terms of $\rho = \max_{k\geq 2}\rho(k)$. The proof of the theorem fails when we replace ρ with $\rho(k)$, because the optimal value of k may not be the same for all prefixes of the input sequence[2].

2.2. A Deterministic Online Sampling Auction: BPSF$_r$

The two online auctions considered in the previous section are randomized, like their offline counterparts. In this section we shift our focus to deterministic online sampling auctions. For the offline setting, we know from Goldberg et al. [2002] that every symmetric and deterministic truthful auction has unbounded competitive ratio against $\mathcal{F}^{(2)}$.

There exist asymmetric deterministic auctions with constant competitive ratio, but most of them are derived by derandomization [Aggarwal et al. 2005]. In the

[2]An earlier version of this article included the incorrect claim that the competitive ratio of SCS is at most 4, for the special case in which the optimal single-price auction for the whole set of bids sells the item to at least 5 buyers. We thank Alkmini Sgouritsa for pointing out the error.

online setting, where order matters, we can hope to design a constant-competitive and deterministic (truthful) auction that is also natural. To this end we define the best-price-so-far auction.

Definition 2.4. Let BPSF_r be the auction which offers as price the bid among the highest r of the previous bids, which maximizes the single-price sale profit of past requests.

We are going to focus our attention on two representatives of this family, BPSF_1 and BPSF_∞, henceforth denoted by BPSF. BPSF_1 is an interesting auction which offers as price the maximum revealed bid. BPSF is an auction that offers the j-th bidder the price $p_j = p(b_{\pi_1}, \ldots, b_{\pi_{j-1}}) = b_{\pi_{i^*}}$, where $i^* = \arg\max_{i \leq j-1} i \cdot b_{\pi_i}$.

THEOREM 2.5. *The expected profit of $BPSF_1$ is at least $\sum_{i=2}^{n} \frac{1}{i} b_i$; in fact, it is exactly equal to this quantity when all bids are distinct. Furthermore, if b_k is the price of the optimal single-price auction, then the competitive ratio of $BPSF_1$ is $\frac{k}{H_k - 1}$, where $H_k = 1 + 1/2 + \cdots + 1/k$ is the k-th harmonic number, and this is tight.*

PROOF. Assume first that the bids are distinct. Notice that b_j is going to be offered as price exactly when b_j appears before $b_1,\ldots,b_{j-1}$. Every such bid is accepted if there is a higher bid after b_j appears. Thus b_j is going to be accepted at some point when $j \geq 2$; notice that each bid can be accepted as a price at most once. The probability that b_j appears before $b_1,\ldots,b_{j-1}$ is exactly $1/j$. It follows that the expected profit of BPSF_1 is $\sum_{i=2}^{n} \frac{1}{i} b_i$. When the bids are not distinct, BPSF_1 can only have higher expected profit, as b_j can be also accepted when it is followed by a bid b_l of equal value and lower-order statistics, that is, when $l < j$.

For the second fact, let b_k be the optimal single price; then the online profit is at least $\sum_{i=2}^{k} \frac{1}{i} b_i \geq \sum_{i=2}^{k} \frac{1}{i} b_k = (H_k - 1) b_k$. Since the optimal profit is $k b_k$, the claim follows.

Finally, it is easy to verify that the preceding bound is tight for any set of n bids with $b_1 > \ldots > b_n$ and $b_n \geq b_1 - \epsilon$, for sufficiently small ϵ. □

Even though BPSF_1 is not constant competitive, we note that, when the number of winners is small (in particular $k \leq 5$), its competitive ratio is at most 4, thus matching our lower bound. This is interesting since having a small number of winners is in general considered to be the "hard" case for digital goods auction (see Alaei et al. [2009] for an example).

It seems plausible that the competitive ratio of BPSF_r decreases with r and that BPSF has the best competitive ratio among these algorithms. BPSF is a very natural online auction and corresponds to the online version of the (offline) *Deterministic Optimal Price* (DOP) auction that offers bidder j the optimal single price of the other bidders. DOP is not competitive [Goldberg et al. 2002], but BPSF on the contrary is constant competitive as our next theorem shows; this is not very surprising because DOP fails to be competitive because it is both symmetric and deterministic, while BPSF violates either one or the other property (depending on how one interprets the random order of the input).

Before stating and proving the theorem about BPSF's comptetiveness, we would like to point out the relationship between BPSF and RSOP (defined in Section 1.2): in both auctions the price offered to a bidder is chosen to be the optimal price for a subset of other bidders of cardinality k, which are chosen uniformly at random among all other $n-1$ bidders. The difference between BPSF and RSOP is that, in BPSF k is a random variable chosen uniformly at random from $\{0, \ldots, n-1\}$, while in RSOP k is chosen from the Binomial$(n-1, 1/2)$ distribution. Specifically, using Profit(S, b_i) to denote the profit we get by offering the optimal single price of S to bidder b_i, we can compute the profit of both auctions with

$$\sum_{b_i} \sum_{S \subseteq \{b_1,\ldots,b_n\} \setminus \{b_i\}} \text{Profit}(S, b_i)\Pr[S] = \sum_{S \subseteq \{b_1,\ldots,b_n\}} \Pr[S] \sum_{b_i \notin S} \text{Profit}(S, b_i),$$

where, as discussed earlier, the probability of S is determined as

$$\Pr[S] = f(|S|)\binom{n-1}{|S|}^{-1},$$

where $f(\cdot)$ is the density function of the Binomial$(n-1, 1/2)$ distribution for RSOP and of the Uniform$(\{0, \ldots, n-1\})$ distribution for BPSF. Therefore

For RSOP: $\Pr[S] = \frac{1}{2^{n-1}}$

For BPSF: $\Pr[S] = \frac{1}{n}\binom{n-1}{|S|}^{-1}$.

Keeping this relation between RSOP and BPSF in mind, we can now prove that BPSF is indeed constant competitive.

THEOREM 2.6. *BPSF has constant-competitive ratio.*

PROOF. We give a proof that is similar to the proof in Feige et al. [2005] which establishes that RSOP is 15-competitive. For simplicity, we make no attempt to optimize the parameters.

We need to set up some notation first. Fix a permutation $b_{\pi_1}, \ldots, b_{\pi_n}$ of the bids. Let k_t^* denote the index of the bid that is the best single price for the first t bids, that is, the optimal single-price value for the first t bids is $k_t^* b_{k_t^*}$. Let k^* be such that $\mathcal{F}^{(2)} = k^* b_{k^*}$. Let $s_{i,t} = |\{b_1, \ldots, b_i\} \cap \{b_{\pi_1}, \ldots, b_{\pi_t}\}|$. Fix two small constants λ and μ and consider the three events:

$\mathcal{R}$. denotes the event that the maximum bid b_1 appears in the second half of the sequence.
$\mathcal{B}$. denotes the event $s_{k^*,n/4} \geq \lambda k^*$.
$\mathcal{E}$. denotes the event: for every $i = 1, \ldots, n$: $s_{i,n/2} \leq (1-\mu)i$.

We want to argue that event $\mathcal{E} \cap \mathcal{B}$ occurs with positive constant probability. This is guaranteed if $\Pr[\mathcal{E}] + \Pr[\mathcal{B}] > 1$. Unfortunately, this does not hold in general. The problem is that with probability 1/2 event $\mathcal{R}$ does not happen (i.e., the maximum bid b_1 appears in the first half of the sequence) and this immediately forces the probability of event $\mathcal{E}$ below 1/2.

To resolve this, we will assume that event $\mathcal{R}$ happens; this assumption will double the competitive ratio. Under this assumption, the probability of event $\mathcal{E}|\mathcal{R}$ approaches 1, as μ approaches 0; one can use Chernoff bounds to get an explicit bound[3]. Also, the probability of $\mathcal{B}|\mathcal{R}$ can be easily bounded from below by a constant; again, one can use Chernoff bounds (or even Markov inequality) to get an explicit bound. Putting the two together, we get that $Pr[\mathcal{E} \cap \mathcal{B}|\mathcal{R}]$ is bounded below by a positive constant.

Suppose then that events $\mathcal{B}$ and $\mathcal{E}$ happen. Then we can lower bound the expected gain from bid at position $t+1$ when $t \in \{n/4+1, \ldots, n/2\}$ as follows: the price offered

[3]A very similar event with $1-\mu = 3/4$ was used in the proof of Feige et al. [2005], for which it was shown that the probability is at least 0.9. Intuitively, the same bound holds for our event $\mathcal{E}|\mathcal{R}$, but it does not appear to exist an obvious way to make a rigorous connection between the two bounds.

by BPSF to $b_{\pi_{t+1}}$ is $b_{k_t^*}$ and the probability for accepting it is equal to $\frac{k_t^* - s_{k_t^*,t}}{n-t} b_{k_t^*}$. Event $\mathcal{E}$ guarantees that this can be bounded from below.

$$\begin{aligned} \frac{k_t^* - s_{k_t^*,t}}{n-t} b_{k_t^*} &\geq \frac{k_t^* - s_{k_t^*,n/2}}{n-t} b_{k_t^*} \\ &\geq \frac{\mu}{(n-t)} k_t^* b_{k_t^*} \\ &\geq \frac{\mu}{(1-\mu)(n-t)} s_{k_t^*,t} b_{k_t^*} \\ &\geq \frac{4\mu}{3(1-\mu)n} s_{k_t^*,t} b_{k_t^*} \end{aligned}$$

On the other hand, the optimality of k_t^* and event $\mathcal{B}$ guarantee that

$$s_{k_t^*,t} b_{k_t^*} \geq s_{k^*,t} b_{k^*} \geq s_{k^*,n/4} b_{k^*} \geq \lambda k^* b_{k^*} = \lambda \mathcal{F}^{(2)}.$$

Putting everything together, under the assumption that both events $\mathcal{B}$ and $\mathcal{E}$ happen, we get that for $t \in \{n/4+1, \ldots, n/2\}$, the expected profit from $b_{\pi_{t+1}}$ is at least

$$\frac{4\mu}{3(1-\mu)n} \lambda \mathcal{F}^{(2)}.$$

The expected profit from all $t \in \{n/4+1, \ldots, n/2\}$ is at least $\frac{\mu\lambda}{3(1-\mu)} \mathcal{F}^{(2)}$. This shows that BPSF has constant-competitive ratio. □

The previous theorem shows that the competitive ratio of BPSF is a large constant. We conjecture that the competitive ratio of BPSF is in fact 4, the same as the conjectured competitive ratio of RSOP.

CONJECTURE 2.7. *The competitive ratio of BPSF is 4.*

ACKNOWLEDGMENTS

We would like to thank Alkmini Sgouritsa for pointing out an error in a previous version of this article and an anonymous reviewer who suggested the constant upper bound on the competitive ratio of BPSF.

REFERENCES

Gagan Aggarwal, Amos Fiat, Andrew V. Goldberg, Jason D. Hartline, Nicole Immorlica, and Madhu Sudan. 2005. Derandomization of auctions. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC'05)*. 619–625.

Saeed Alaei, Azarakhsh Malekian, and Aravind Srinivasan. 2009. On random sampling auctions for digital goods. In *Proceedings of the ACM Conference on Electronic Commerce*. 187–196.

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

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

Ziv Bar-Yossef, Kirsten Hildrum, and Felix Wu. 2002. Incentive-compatible online auctions for digital goods. In *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'02)*. 964–970.

Avrim Blum and Jason D. Hartline. 2005. Near-optimal online auctions. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'05)*. 1156–1163.

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

Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. 2010. Revenue maximization with a single sample. In *Proceedings of the ACM Conference on Electronic Commerce*. 129–138.

Uriel Feige, Abraham Flaxman, Jason D. Hartline, and Robert D. Kleinberg. 2005. On the competitive ratio of the random sampling auction. In *Proceedings of the 1st International Workshop on Internet and Network Economics (WINE'05)*. 878–886.

Andrew V. Goldberg, Jason D. Hartline, Anna R. Karlin, Andrew Wright, and Michael Saks. 2002. Competitive auctions. *Games Econ. Behav. 55*, 2.

Mohammad Taghi Hajiaghayi, Robert D. Kleinberg, and David C. Parkes. 2004. Adaptive limited-supply online auctions. In *Proceedings of the ACM Conference on Electronic Commerce*. 71–80.

Jason D. Hartline and Robert McGrew. 2005. From optimal limited to unlimited supply auctions. In *Proceedings of the ACM Conference on Electronic Commerce.* 175–182.

Jason D. Hartline and Tim Roughgarden. 2008. Optimal mechanism design and money burning. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC'08).* 75–84.

Roger B. Myerson. 1981. Optimal auction design. *Math. Oper. Res. 6*, 58–73.

Received August 2011; revised December 2012; accepted December 2012