Balloon Popping With Applications to Ascending Auctions

Nicole Immorlica* Anna R. Karlin† Mohammad Mahdian‡ Kunal Talwar§

Abstract

We study the power of *ascending auctions* in a scenario in which a seller is selling a collection of identical items to anonymous unit-demand bidders. We show that even with full knowledge of the set of bidders' private valuations for the items, if the bidders are ex-ante identical, no ascending auction can extract more than a constant times the revenue of the best fixed price scheme.

This problem is equivalent to the problem of coming up with an optimal strategy for blowing up indistinguishable balloons with known capacities in order to maximize the amount of contained air. We show that the algorithm which simply inflates all balloons to a fixed volume is close to optimal in this setting.

*Microsoft Research, Redmond, WA. email: nickle@microsoft.com

†University of Washington, Seattle, WA. email: karlin@cs.washington.edu

‡Yahoo Research, Santa Clara, CA. email: mahdian@yahoo-inc.com

§Microsoft Research, Mountain View, CA. email: kunal@microsoft.com

1 Introduction

Ascending price auctions based on a mechanism due to Moulin [**?**] have been a major building block in the design of revenue-competitive mechanisms in the computer science literature []. Such mechanisms are used to design objects called *profit extractors*, which given a target revenue (often computed using random sampling) and a set of bidders, try to extract this amount from the bidders in an incentive-compatible manner.

A natural question (which was the initial motivation for the present work) is whether the Moulin mechanism at the heart of this paradigm is the best choice; in other words, are there other mechanisms that can raise considerably higher revenues? We address this question by studying ascending auctions that treat bidders anonymously. As we will argue in Section 2, these conditions are both reasonable and necessary. We start by formulating the problem as a balloon popping puzzle. The connection between this puzzle and ascending auctions as well as other implications of our results for online auction design will be explained in the next section.

The balloon popping problem. Suppose that n indistinguishable balloons, yet to be blown up, are laid out on a table in front of you. Your challenge is to blow up the balloons to a maximum total volume. Before you start, you are told that the capacities of the balloons are $1, \frac{1}{2}, \frac{1}{3}, \ldots, \frac{1}{n}$. Of course, if you knew which balloon was which, you'd be able to blow each balloon up to its capacity for a total volume of H_n, the n-th Harmonic number. Unfortunately, you are told that the balloons have been laid out on the table in a random order, and if you blow up a balloon beyond its capacity, it will pop. What is the best strategy for blowing up the balloons in order to maximize the total volume? What total volume is achievable?

To get a feel for the problem, consider the following strategies:

Strategy 1: Blow all the balloons up to $1/k$ for some k.

Strategy 2: Initialize V to $\{1, 1/2, \ldots, 1/n\}$, the set of possible balloon capacities still under consideration. For each balloon successively, blow it up to the maximum value v in V and, if the balloon pops at volume $1/i$, set $V = V - \{1/i\}$, otherwise set $V = V - \{v\}$.

In Strategy 1, all the balloons of capacity $1/i$, $1 \leq i \leq k$ will end up at a volume of $1/k$ and all other balloons will pop, since their capacity will be exceeded. Thus, the final total volume of the remaining balloons will be 1. However, in Strategy 2, you benefit by allowing different balloons to contribute different amounts to the total volume. In fact, it is not hard to see that the balloon with capacity $1/k$ gets filled to capacity precisely if it appears after all the balloons of larger capacity, which happens with probability $1/k$. Thus, the expected final total volume is $\sum_{1 \leq k \leq n} \frac{1}{k^2}$ which is $\pi^2/6 \approx 1.64$ in the limit as n goes to infinity.

Unfortunately, Strategy 2 is not optimal either. However, it is not too far off. Our main result is the following:

Theorem 1 *Suppose the capacities of the balloons are $v_1 \geq v_2 \geq \ldots \geq v_n$. No balloon popping strategy can achieve an expected total volume more than $\alpha \max_i iv_i$, where $\alpha \leq 4.331$.*

2 Motivation

The balloon popping problem is motivated by our desire to understand how much revenue can be raised by an *ascending auctions* with *anonymous bidders* for selling *digital goods*. Digital goods auctions, also known as unlimited supply auctions, have been the subject of extensive research in the computer science literature [][1]. Most of the work in this literature has focused on designing mechanisms whose revenue is competitive with respect to the maximum revenue that can be raised by an omniscient auctioneer who is restricted to offer the good at the *same* price to all bidders. However, it is not clear that this is a reasonable restriction, and whether without this restriction, the auctioneer can achieve a considerably higher revenue.

One must be careful in formulating this question, as for any set of utilities, it is possible to craft a mechanism that extracts the full surplus if the agents happen to have the predicted set of utilities. Therefore, we cannot hope to find a mechanism that achieves the highest revenue among *all* incentive-compatible mechanisms for *every* set of utilities. In the economic literature, this problem is often overcome by assuming a distribution on the set of utilities, and looking for a mechanism that achieves the highest *expected* revenue among all mechanisms for this particular distribution (see, for example, [14, 6, 2, 7]). In this paper, we take a different approach: we restrict the set of mechanisms we are considering by two properties, namely treating bidders anonymously and being an ascending mechanism (which as we will argue are both reasonable and necessary), and show that on *every* set of utilities, a fixed price scheme achieves a revenue that is at least a constant fraction of the revenue generated by *any* mechanism satisfying these properties.[2]

The first property is *anonymity*. Bidders are anonymous if an auctioneer is unable to distinguish between them before they start bidding (i.e., each bidder is equally likely to have each value a priori). Our assumption that bidders are anonymous is natural for many auctions on the Internet, and is motivated by applications in which it is difficult, prohibitively costly, or legally restricted to verify the identities of the bidders or discriminate among them based on their identities. Also, note that without the anonymity assumption, for any set of utilities, there is an auction that extracts the full surplus by charging each bidder her utility. Clearly, a fixed price scheme cannot achieve a revenue within a constant factor of the full surplus.

We will also restrict our attention to the class of *ascending auctions*, which are mechanisms in which the published prices can only rise over time (such as the English auction). Ascending auctions are particularly interesting and widely studied for several reasons. First, they represent a very intuitive way to sell items: increase the prices for the overly eager bidders or for the over-demanded items. Second, there is empirical evidence that they are appealing to bidders and increase the "trust" of the bidders in the auctioneer because it is easy for the bidders to understand what is happening and why. In addition, the bidders reveal less information than they do with a truthful mechanism. Finally, they are widely used in many real-life settings ranging from the FCC spectrum auctions and auctions for US Treasury Bills to most online auctions like eBay, eBid, and Yahoo! Auctions.

However, in the case of interdependent values, one might imagine that a seller could use the value dependencies to deduce optimal prices for the buyers, and indeed, Cremer and McLean [7] prove that there is an auction mechanism that extracts the full buyers' surplus when values are interdependent for a range of priors, including some where the buyers are anonymous (i.e., values are drawn according to a *symmetric* joint distribution). Cast in this light, our results imply that in the same settings considered

[1] As noted by other papers on the subject (see, for example, []), multi-unit auctions and auctions with constant marginal cost of producing an additional unit of supply can be easily reduced to digital goods auctions. Also, even though there is formal reduction, most of the approaches applied for designing digital goods auctions have been generalized to settings with budget constraints [**?**, **?**].

[2] From a mechanism design point of view, the fixed price mechanism can be implemented either directly, if the mechanism designer has information about the distribution of the utilities, or using competitive auctions of [] which can be implemented as ascending auctions.

by Cremer and McLean, ascending auctions can not extract significantly more revenue than the ex-post optimal uniform pricing, and in particular, there are instances where such mechanisms can only extract an $O(1/\log n)$ fraction of the total buyers' surplus.

In contrast to the competitive auction framework, profit maximization problems in the economics literature are typically approached by assuming that the privately known values v_i are drawn according to a publicly known distribution. When buyers have independent and identically distributed values, a sequence of work [14, 6, 2] shows that for a general class of cost functions (including the fixed cost case considered here), the maximum profit is extracted by a uniform price.

However, in the case of interdependent values, one might imagine that a seller could use the value dependencies to deduce optimal prices for the buyers, and indeed, Cremer and McLean [7] prove that there is an auction mechanism that extracts the full buyers' surplus when values are interdependent for a range of priors, including some where the buyers are anonymous (i.e., values are drawn according to a *symmetric* joint distribution). Cast in this light, our results imply that in the same settings considered by Cremer and McLean, ascending auctions can not extract significantly more revenue than the ex-post optimal uniform pricing, and in particular, there are instances where such mechanisms can only extract an $O(1/\log n)$ fraction of the total buyers' surplus.

Our focus here is on a setting of a monopolistic auctioneer who wishes to sell a set of identical goods to bidders with unit demand. Each bidder has a privately known value for the good. In an ascending auction, as in a balloon popping mechanism, the auctioneer maintains at all times a current price vector whose i'th component represents the current price being offered to the i'th remaining bidder. (The current price associated with a bidder obviously corresponds to the current volume of a balloon.) The auctioneer is allowed to arbitrarily raise the prices for any bidder. In doing so, however, he risks the possibility that a price will be raised above a bidder's private value for the item, denoted by v_i for the i'th bidder, in which case that bidder will withdraw from the auction. At some point of the auctioneer's choosing, he closes the auction and sells to the active bidders at the current prices.

A simple strategy for the auctioneer is to just charge a fixed price to all bidders. The auctioneer may hope, however, that he will be able to price discriminate – sell identical goods to different buyers at different prices. With discriminatory pricing, a seller can often achieve a much higher revenue than that achievable by uniform pricing.

We endow the ascending auctions we study with a great deal of information – the precise set of bidders' values v_i. It is not hard to imagine that highly accurate distributional information could be collected via market research. However, because of our anonymity assumption, the bidders are presumed to have a random permutation of these values. Note however that this only makes our result stronger; the result still applies for the case when such accurate information is not available.

Our bound on the maximum total volume any balloon popping strategy can achieve can be viewed as a negative result about the power of ascending auctions with respect to profit maximization. Even with *full knowledge* of the set of bidders' values, no ascending auction can extract a profit more than that of the best fixed price profit by more than a small constant factor. (Of course, the best fixed price profit is trivially achievable by an ascending auction when the set of bidder values is known – simply choose v_i such that iv_i is maximized and charge all bidders v_i.)

There is a fairly vast literature on ascending auctions, focused largely on their power with respect to the goal of maximizing efficiency. Notably, Blumrosen and Nisan [5] recently undertook a systematic study of this question (with respect to the goal of maximizing efficiency). The current paper begins the study of the power of ascending auctions with respect to the goal of profit maximization.

There are a few other perspectives on this result also worth noting.

2.1 Competitive auction design

In the computer science literature, Goldberg, Hartline, and Wright [10] initiated the study of *competitive auction design* in which the goal is to guarantee a revenue approximately equal to that of an omniscient, but constrained, auctioneer. The best-known algorithms for profit maximization in the competitive auction framework work by reducing the profit maximization problem to the *profit extraction problem* [9, 11, 8]. Loosely speaking, the profit extraction problem is to:

> Design a truthful auction that, given a target profit R, on each input, extracts a profit of R whenever "possible", where "possible" means that some natural profit metric of the input is at least R.

The most well-known and widely-used profit extractor to date is based on the *Moulin mechanism* with equal cost shares. It extracts a revenue of R whenever $R \leq \max_p |\{v_i : v_i \geq p\}| \cdot p$, and is implementable by a simple ascending auction. In fact, *all* known profit extractors in the literature can be implemented as ascending auctions (though no precise characterization or equivalence is known). Our result shows that the Moulin mechanism with equal cost shares is essentially the best (up to a constant factor) possible ascending auction mechanism for profit extraction for selling identical items to unit-demand bidders, even in the presence of perfect market research.

2.2 Online auction design

Several recent papers [1, 3, 4, 12] have studied the design of revenue maximizing auctions for settings in which bidders arrive one at a time *online*, as happens, for example, in the popular online travel agency, `priceline.com`. (See [15] for a short introduction). In these models, in between bidder arrivals the auctioneer computes an offer price for the next bidder using previously submitted bids. When a new bidder arrives, he is allocated the good if his bid exceeds the current offer price, and pays the current offer. Such an auction is obviously truthful (the bidder has no incentive to lie about his value). For such a model, there are auction designs which are guaranteed to achieve a constant fraction of the optimal fixed price revenue (e.g., [3]).

Perhaps the most interesting aspect of our proof is a reduction of the balloon popping problem to an online version of the balloon popping problem that corresponds precisely to the model just described. This reduction shows that precise distributional knowledge cannot improve existing online auction results, for bidders that arrive in a random order, by more than a small constant factor.

3 Model

We draw an analogy between ascending auctions and algorithms for inflating balloons, called *balloon popping mechanisms*. A balloon popping mechanism takes as input a set of identical balloons and their capacities $\{v_i\}$. We label the balloons' capacities in decreasing order, i.e. $v_i \geq v_{i+1}$ for $1 \leq i \leq n$. We assume that the balloons themselves are ordered according to a random permutation π unknown to the mechanism, so that the capacity of balloon i is v_{π_i}.

As the balloons are identical, the mechanism can not distinguish capacities of individual balloons. Still, we wish to design a mechanism which inflates balloons in a way that maximizes the amount of contained air. Let s be an n-dimensional vector where the i'th component is either P, indicating that balloon i has popped, or a nonnegative real number s_i, indicating that balloon i's current size is s_i. We initialize s to the all 0 vector.

Definition 2 *An* offline balloon popping mechanism *for capacities $\{v_i\}$ is defined by a function $Blow(\mathbf{s})$, which outputs either Stop or a pair (i, b), which indicates that b units of air, for some positive number b, should be blown into balloon i. The vector $\mathbf{s}$ gets updated appropriately: if the output of Blow is Stop, then $\mathbf{s}$ is unchanged and the process ends. Otherwise, entry i of $\mathbf{s}$ is replaced by P if balloon i pops (which happens if $s_i + b > v_{\pi_i}$) or by $s_i + b$ otherwise. Let S be the final set of indices such that $s_i \neq P$. The* payoff *or* revenue *of the balloon popping mechanism is $\sum_{i \in S} s_i$ at the first time the output is Stop.*

Notice that the mechanism as defined is memoryless. It is easy to see, however, that there is no loss of generality in describing the mechanism this way.

Corresponding to online auctions, we also define online balloon popping mechanisms. These online mechanisms process balloons sequentially in a fixed order (say, $1, 2, \dots, n$) and are granted the additional knowledge of a balloon's capacity once they have completed processing that balloon *regardless of whether or not the balloon popped*. (This corresponds precisely to an agent submitting his bid in an online auction.)

Definition 3 *An* online balloon popping mechanism *is defined by a function $Blow(\{v_{\pi_1}, \dots, v_{\pi_{i-1}}\})$ that outputs a non-negative number b, indicating that b units of air should be blown into balloon i. If $b \leq v_{\pi_i}$, let $s_i = b$. Otherwise, let $s_i = 0$. The* payoff *or* revenue *of the online balloon popping mechanism is $\sum_i s_i$.*

Note that with the above definition, it is not a priori clear whether online or the offline balloon popping mechanisms are stronger, since the online mechanisms are given the additional knowledge of the values of the balloons that are processed and not popped, while the offline mechanisms are given the flexibility to get back to an already-inflated balloon and inflate it further.

4 The main result

We now recall the main theorem:

Theorem 4 *Suppose the capacities of the balloons are $v_1 \geq v_2 \geq \dots \geq v_n$. No balloon popping strategy can achieve an expected total volume more than $\alpha \max_i i v_i$, where $\alpha \leq 4.331$.*

The theorem is proven in three steps. First, we observe that we can upper bound α by assuming that $(v_1, \dots, v_n) = (1, \frac{1}{2}, \dots, \frac{1}{n})$. We call this latter vector of capacities the *harmonic capacities*. Then we show that for any capacity vector, the expected maximum volume achievable by an offline balloon popping mechanism is upper bounded by that of the optimal online balloon popping mechanism. Finally, we obtain an expression for the maximum expected volume that can be achieved by the optimal online balloon popping mechanism and upper bound that expression using Chernoff bounds for the harmonic capacities.

We begin with step 1:

Lemma 5 *Let $v_1 \geq v_2 \geq \dots \geq v_n$ be a set of balloon capacities. Without loss of generality, assume that $\max_i i v_i = 1$. Then the maximum expected volume achievable by a balloon popping mechanism on $(v_1, v_2, \dots, v_n)$ is at most the maximum achievable on harmonic capacities.*

Proof. Since $\max_i(iv_i) = 1$, we have that $v_i \leq 1/i$ for all i. First observe that there is never any point in blowing up any balloon to a value which lies strictly in between v_j and v_{j+1} for some j. Consider any balloon popping mechanism tailored to the set $\{v_i\}$. Simulate it on the harmonic capacities as follows: whenever a balloon is blown up to v_i, blow up the corresponding balloon in the harmonic case to $1/i$. Since the probability of success on any particular blow is the same, and the volume obtained is greater for the harmonic capacities, the claim is proved. ■

4.1 Reduction to online balloon popping

Next, we prove that the revenue of an offline balloon popping mechanism can be bounded by that of the optimal online balloon popping mechanism.

First, we show that the revenue of the optimal online balloon popping mechanism can be characterized by a simple formula.

Lemma 6 *Let T be a random subset of $\{v_1, \ldots, v_n\}$ and let $\{y_1, y_2, \ldots, y_k\}$ be the order statistics of the set T (so y_1 is the largest element of T, y_2 is the second largest element of T and so on). Define the random variable $g(T)$ to be*

$$g(T) = \frac{\max_j(jy_j)}{|T|}. \tag{1}$$

Then the revenue of the optimal online balloon popping mechanism is

$$\sum_{k=1}^{n} E_T[g(T)]. \tag{2}$$

where the expectation in the k-th term in the summation is over a random subset T of $\{v_1, \ldots, v_n\}$ of size k.

Proof. Let OPT be an optimal balloon popping mechanism. Consider the $(n - k + 1)$'th step of OPT, i.e., when there are exactly k balloons remaining. The set of capacities of these balloons is denoted by $T = \{y_1 \geq y_2 \geq \ldots y_k\}$. Clearly, T is a random k-subset of $\{v_1, \ldots, v_n\}$. At this step, OPT inflates balloon $n - k + 1$ to the volume y_j for some j. Therefore, the probability over the random permutation π that the this balloon survives is exactly $j/|T|$, and the contribution to the total payoff given that it survives is y_j. Thus, the expected revenue of OPT at this step is $\frac{jy_j}{|T|}$, which is at most $g(T)$. Therefore, by the linearity of expectation, the expected revenue of the optimal mechanism is at most $\sum_{k=1}^{n} E_T[g(T)]$, where the expectation is over a random subset T of $\{v_1, \ldots, v_n\}$ of size k. On the other hand, it is easy to see that the online mechanism that inflates the $(n - k + 1)$'th balloon to the volume y_j for a j that maximizes jy_j has an expected revenue equal to $\sum_{k=1}^{n} E_T[g(T)]$. ■

Now consider the optimal offline balloon popping mechanism. Initially, the set of values of the balloons is the set $\{v_1, \ldots, v_n\}$, and in each step, if a balloon pops, we know its value. Therefore, at any moment we know the set of values of the remaining balloons. Let $x_1, x_2, \ldots$ denote the set of remaining values in decreasing order, and $S = \{x_1, x_2, \ldots\}$. We can restrict the algorithm to always inflate the balloons to one of the levels in S. For balloon j, if it is already inflated to level x_{ℓ_j}, we say that ℓ_j is the current lower bound for this balloon. (Note the slight abuse of notation here – ℓ_j is actually the *index* of the element in S whose value is a lower bound on balloon j. Nonetheless, we will refer to both ℓ_j and x_{ℓ_j} as the lower bounds for balloon j.) Let ℓ denote the vector $(\ell_1, \ldots, \ell_{|S|})$ of current lower bounds, sorted in decreasing order. The set S and the lower bounds summarize all the information that the algorithm has at this moment. Importantly, each permutation that is consistent with these lower bounds on the capacities is equally likely.

For example, if $S = \{1/3, 1/7, 1/10\}$ and there are two balloons (wlog the first two) that have been inflated to 1/7 and one balloon that has been inflated to 1/10, then the lower bound vector (ℓ_1, ℓ_2, ℓ_3) is $(2, 2, 3)$. In this case, the permutations $(1/3, 1/7, 1/10)$ or $(1/7, 1/3, 1/10)$ are equally likely.

Any balloon popping algorithm can be broken down into minimal steps in which a balloon with current lower bound x_{ℓ_k} is blown up to $x_{\ell_k - 1}$. In addition wlog we can maintain the invariant that the balloons are ordered in order of decreasing lower bounds: on a step in which the optimal algorithm chooses to blow up a balloon with current lower bound x_{ℓ_k} to x_{ℓ_k-1}, we will pick the leftmost balloon of lower bound x_{ℓ_k} to blow into. Since the distribution of capacities for balloons with the same lower bound is identical, we lose nothing by doing this. This selection ensures that the balloons remain sorted in decreasing order of lower bounds.

Let $\mathrm{Rev}(S, \ell)$ denote the maximum expected revenue starting from a position indicated by (S, ℓ). We also define a function $f(S, \ell)$ (that we will prove is an upper bound for $\mathrm{Rev}(S, \ell)$) as follows.

Fix a permutation $(x_{\pi_1}, x_{\pi_2}, \ldots, x_{\pi_{|S|}})$ of the elements of S that is consistent with the lower bounds ℓ (i.e., $x_{\pi_i} \geq x_{\ell_i}$ for all i). Define $S_j(\pi)$ to be the subset of balloons $\{j, \ldots, |S|\}$ in the permutation π whose capacity is at least x_{ℓ_j}, i.e.,

$$S_j(\pi) = \{x_{\pi_i} | i \geq j \text{ and } x_{\pi_i} \geq x_{\ell_j}\}).$$

We define

$$h(\pi, \ell) = \sum_{j=1}^{|S|} g(S_j(\pi)).$$

where g is the function defined in Equation 1. Finally, we let

$$f(S, \ell) = \mathrm{Exp}_\pi[h(\pi, \ell)],$$

where the expectation is over the random choice of a permutation π from the set of all permutations of S satisfying $x_{\pi_j} \geq x_{\ell_j}$ for all j. Intuitively, $f(S, \ell)$ is the revenue the optimal online mechanism would get if it had the same information regarding lower bounds as the offline mechanism. We prove the following theorem.

Theorem 7 *For every set S and every non-increasing vector of lower bounds ℓ, we have*

$$\mathrm{Rev}(S, \ell) \leq f(S, \ell). \tag{3}$$

Proof. We prove this statement by induction on $\sum_j \ell_j$. Consider the first move of the optimal algorithm at (S, ℓ). This move is either to stop and extract the revenue of $\sum_j x_{\ell_j}$, or to blow up the k'th balloon to $x_{\ell_k - 1}$, for some k (which either pops this balloon or raises its lower bound from ℓ_k to $\ell_k - 1$). In the former case, the inequality follows from the fact that for any set T, $g(T)$ is at least the minimum value in T, and $S_j(\pi)$ always contains at least one element of value x_{ℓ_j} or more. Thus, in every permutation π, $h(\pi, \ell) \geq \sum_j x_{\ell_j}$.

Now, assume that the best move of the optimal algorithm is to blow the k'th balloon to $x_{\ell_k - 1}$ (hence changing its lower bound from ℓ_k to $\ell_k - 1$). In this case, the k'th balloon pops with probability p_{fail}, and survives with probability $p_{success}$. Therefore,

$$\mathrm{Rev}(S, \ell) \leq p_{fail} \mathrm{Rev}(S \setminus \{x_{\ell_k}\}, \ell^*_{-k}) + p_{success} \mathrm{Rev}(S, \ell - 1_k),$$

where 1_k denotes a vector with a 1 in the k'th entry and zero everywhere else, and ℓ^*_{-k} denotes the vector ℓ with the k'th element removed and other elements adjusted so that if an element j had a lower bound of x_{ℓ_k} in ℓ, it now has a lower bound of x_{ℓ_k-1} (since x_{ℓ_k} is no longer among the set of remaining values), and all other elements have the same lower bound. By the induction hypothesis, we have

$$\begin{aligned}\text{Rev}(S,\ell) &\leq p_{fail} f(S \setminus \{x_{\ell_k}\}, \ell^*_{-k}) + p_{success} f(S, \ell - 1_k) \\ &= p_{fail}\text{Exp}_{\pi'}[h(\pi', \ell^*_{-k})] + p_{success}\text{Exp}_{\pi'}[h(\pi', \ell - 1_k)],\end{aligned}$$

where the first expectation is over the set of permutations π' of $S \setminus \{x_{\ell_k}\}$ that satisfy the lower bounds in ℓ^*_{-k}, and the second expectation is over permutations π' of S satisfying the lower bounds in $\ell - 1_k$. Equivalently, the random choice of π' in the first expectation can be done by picking a random permutation π of S satisfying the lower bounds in ℓ and satisfying $x_{\pi_k} = x_{\ell_k}$, and letting $\pi' = \pi_{-k}$. Similarly, the random choice of π' in the second expectation can be done by picking a random permutation π of S satisfying the lower bounds in ℓ and satisfying $x_{\pi_k} > x_{\ell_k}$, and letting $\pi' = \pi$. Also, notice that p_{fail} and $p_{success}$ are precisely the probabilities of the events $x_{\pi_k} = x_{\ell_k}$ and $x_{\pi_k} > x_{\ell_k}$. By this observation, the above inequality can be written as follows:

$$\text{Rev}(S,\ell) \quad \leq \quad \text{Exp}_{\pi}[h(\pi', \ell')], \tag{4}$$

where the expectation is over the random choice of a permutation π of S satisfying the lower bounds in ℓ, and

$$(\pi', \ell') = \begin{cases} (\pi_{-k}, \ell^*_{-k}) & \text{if } x_{\pi_k} = x_{\ell_k} \\ (\pi, \ell - 1_k) & \text{if } x_{\pi_k} > x_{\ell_k}. \end{cases}$$

By Equation 4, the definition of f, and linearity of expectation, to prove inequality 3 it will suffice to show:

$$\text{Exp}_{\pi}[h(\pi,\ell) - h(\pi', \ell')] \geq 0. \tag{5}$$

By the definition of h, for a fixed permutation π, we have

$$h(\pi,\ell) - h(\pi',\ell') = \begin{cases} g(S_k(\pi)) & \text{if } x_{\pi_k} = x_{\ell_k} \\ g(S_k(\pi)) - g(\{x_{\pi_i} | i \geq k \text{ and } x_{\pi_i} \geq x_{\ell_k - 1}\}) & \text{if } x_{\pi_k} > x_{\ell_k}. \end{cases} \tag{6}$$

Equation (6) follows from the fact that for any permutation π such that $x_{\pi_k} = x_{\ell_k}$, $g(S_j(\pi)) = g(S_j(\pi'))$ for any $1 \leq j < k$, and $g(S_j(\pi)) = g(S_{j-1}(\pi'))$ for any $j > k$, whereas for any permutation π such that $x_{\pi_k} > x_{\ell_k}$, $g(S_j(\pi)) = g(S_j(\pi'))$ for any $1 \leq j < k$, and $g(S_j(\pi)) = g(S_j(\pi'))$ for any $j > k$.

We now partition the set of permutations π of S satisfying the lower bounds in ℓ into several groups as follows: each group corresponds to one permutation π^* satisfying the lower bounds in ℓ, with $x_{\pi^*_k} = x_{\ell_k}$. Let $T_{\pi^*} = \{x_{\pi^*_i} | i \geq k \text{ and } x_{\pi^*_i} \geq x_{\ell_k}\}$. For each element $x_{\pi^*_r}$ of T_{π^*}, we add a permutation to π^*'s group which is obtained by swapping π^*_r and π^*_k (note that by the definition of T_{π^*}, this permutation satisfies the lower bounds ℓ). This specifies the set of all permutations in π^*'s group.

Observe that each permutation of S satisfying lower bounds ℓ appears in exactly in one of the groups defined above. This follows from the assumption that the mechanism keeps the lower bounds in non-increasing order, and therefore at the moment that the mechanism inflates the k'th balloon to $x_{\ell_k - 1}$, every balloon before this balloon must have a lower bound of $x_{\ell_k - 1}$ or higher. Therefore, for every permutation π satisfying the lower bounds ℓ, there is exactly one $i \geq k$ such that $x_{\pi_i} = x_{\ell_k}$. Swapping π_i and π_k, we obtain the unique permutation π^* that has π in its group.

Therefore, in order to show that the expectation in Equation 5 is non-negative, it is enough to show that for every group, the expectation taken over permutations in that group satisfying the lower bounds is non-negative. Consider a group represented by the permutation π^*, and let $t = |T_{\pi^*}|$. By Equation 6 and the definition of π^*'s group, the sum of the terms $h(\pi, \ell) - h(\pi', \ell')$ for permutations π in this group is at most:

$$g(T_{\pi^*}) + (t-1)(g(T_{\pi^*}) - g(T_{\pi^*} \setminus \{x_{\ell_k}\}))$$

By the definition of g, we have $g(T_{\pi^*} \setminus \{x_{\ell_k}\}) \le \frac{t}{t-1} g(T_{\pi^*})$. Therefore, the above summation is at least

$$g(T_{\pi^*}) \left(1 + (t-1)(1 - \frac{t}{t-1})\right) = 0.$$

This completes the proof of the induction step. ■

Corollary 8 *The expected payoff of the optimal offline balloon popping mechanism is at most that of the optimal online balloon popping mechanism.*

Proof. Setting S to $\{v_1, \ldots, v_n\}$ and ℓ to the vector all of whose entries are equal to n (meaning that the initial lower bounds on the capacities of all balloons are $1/n$), we see the maximum revenue in the offline balloons problem can be bounded by that of the optimal online balloons algorithm, Equation 2 above. ■

4.2 Bounding the optimal payoff

Finally, we show a constant upper bound on the expression we obtained in Lemma 6.

Theorem 9 *The revenue of the optimal online balloon popping mechanism is at most 4.331.*

Proof. We need to bound the expectation of the random variable

$$g(T) = \frac{\max_j (j y_j)}{|T|}$$

for a random subset $T = \{y_1, \ldots, y_k\}$ of the set $\mathbf{H_n} = \{1, \frac{1}{2}, \ldots, \frac{1}{n}\}$, such that $y_1 \ge y_2 \ge \ldots \ge y_k$. Equivalently, if we let n_i denote $|\{j \in T : y_j \ge 1/i\}|$, we have

$$g(T) = \max_i \frac{n_i}{i|T|}.$$

It is immediate that

$$g(T) \le 2 \max_s \frac{n_{2^s}}{2^s|T|}.$$

We wish to show that if T is random k-subset of H_n, then the expected value of $|T|g(T)$ is close to $\frac{k}{n} \log \frac{n}{k}$. To handle the max operator above, we shall use the inequality $E[\max(X,Y)] \le E[X] + E[Y]$ repeatedly, to derive:

$$E[|T|g(T)] \le 2(\sum_{s \le s^*} E[\frac{n_{2^s}}{2^s}] + E[\max_{s > s^*} \frac{n_{2^s}}{2^s}]).$$

We shall use this inequality for a suitable s^*.

Toward this end, let B_i^r be the event that n_i/i is at least $rk/(n-k)$. We first bound the probability of B_i^r. Let the set T be drawn at random by picking elements $z_1, \ldots, z_k$ one by one uniformly at random from the set of remaining elements, and let X_j be the indicator variable $\mathbf{1}(z_j \geq \frac{1}{i})$. Clearly for each j, X_j is one with probability at most $\frac{i}{n-k}$. Thus the random variables $X_1, \ldots, X_k$ are stochastically dominated by random variables $Y_1, \ldots, Y_k$, where each Y_i is one with probability exactly $i/(n-k)$. Thus the random variable $X = \sum_j X_j$ is stochastically dominated by $Y = \sum_j Y_j$. Note that $E[Y] = ik/(n-k)$.

Let $r > 1$ be arbitrary. Then by Chernoff bounds [13], the probability that $\sum Y_j \geq rik/(n-k)$ is at most $(e^{r-1}/r^r)^{(ik/(n-k))}$. Since Y stochastically dominates $X = n_i$, we have that $Pr[B_i^r] < (e^{r-1}/r^r)^{(ik/(n-k))}$. Next we simply use the union bound:

$$\begin{aligned}
Pr[\exists s > s^* \text{ s.t. } B_{2^s}^r] &\leq \sum_{s>s^*} Pr[B_{2^s}^r] \\
&\leq \sum_{s>s^*} (e^{r-1}/r^r)^{(k/(n-k))2^s}
\end{aligned}$$

We shall choose s^* to be $\lceil \log(n/k) \rceil > \log(\frac{n-k}{k})$. For this choice of s^*, the summation becomes $\sum_{j \geq 1} (e^{r-1}/r^r)^{2^j}$. Thus we have

$$Pr[\exists s > s^* \text{ s.t. } B_{2^s}^r] \leq \min(1, \sum_{j \geq 1} (e^{r-1}/r^r)^{2^j})$$

The expected value of $\max_{s>s^*} \frac{n_{2^s}}{2^s}$ is thus at most

$$\begin{aligned}
E[\max_{s>s^*} \frac{n_{2^s}}{2^s}] &\leq \int_{r \geq 0} (\frac{k}{n-k}) Pr[\max_{s>s^*} \frac{n_{2^s}}{2^s} \geq rk/(n-k)] dr \\
&\leq (\frac{k}{n-k}) \int_{r \geq 0} \min(1, \sum_{j \geq 1} (e^{r-1}/r^r)^{2^j}) dr \\
&\leq (\frac{2.31k}{n-k})
\end{aligned}$$

where the last inequality is obtained by numerically estimating the integral.

Also note that for any s, the expected value of $\frac{n_{2s}}{2^s}$ is at most k/n. Thus

$$\begin{aligned}
E[|T| g(T)] &\leq 2(\sum_{s \leq s^*} E[\frac{n_{2^s}}{2^s}] + E[\max_{s>s^*} \frac{n_{2^s}}{2^s}]) \\
&\leq 2(s^*(k/n) + \frac{2.31k}{n-k}) \\
&\leq (2k/n)(\lceil \log(n/k) \rceil + 2.464)
\end{aligned}$$

whenever $k < \frac{n}{16}$. Moreover, note that for any k, $|T| g(T) \leq |\mathbf{H_n}| g(\mathbf{H_n}) = 1$.

By Lemma 6, the total expected revenue is thus bounded by

$$\sum_{k=1}^{n} E_T[g(T)] \leq \sum_{k=1}^{\frac{n}{16}} (\frac{2}{n}) (\lceil \log(n/k) \rceil + 2.464) + \sum_{k=\frac{n}{16}+1}^{n} \frac{1}{k}$$

$$
\begin{aligned}
&\leq \sum_{k=1}^{\frac{n}{16}} (\frac{2}{n}) \left(\lceil \log(n/k) \rceil + 2.464\right) + \ln 16 \\
&\leq \frac{2.464}{8} + \sum_{k=1}^{\frac{n}{16}} (\frac{2}{n}) \lceil \log(n/k) \rceil + \ln 16 \\
&\leq 3.081 + \frac{2}{n} \sum_{j=0}^{\lceil \log(\frac{n}{16}) \rceil - 1} \sum_{l=2^j}^{2^{j+1}} \lceil \log(n/l) \rceil \\
&\leq 3.081 + \frac{2}{n} \sum_{j=0}^{\lceil \log(\frac{n}{16}) \rceil - 1} 2^j (\lceil \log n \rceil - j) \\
&\leq 3.081 + \frac{2}{n} \sum_{t=\lceil \log n \rceil - \lceil \log(\frac{n}{16}) \rceil}^{\log n} t 2^{(\log n - t)} \\
&\leq 3.081 + 2 \sum_{t=4}^{\log n} t 2^{-t} \\
&\leq 3.081 + 1.25 = 4.331
\end{aligned}
$$

Hence the claim. ■

5 Conclusions

In this paper, we have shown that no balloon popping mechanism for blowing up indistinguishable, randomly permuted balloons with varying capacities can obtain a total expected volume more than $4.331 \max_i iv_i$ for any capacities. This upper bound is not tight. For the harmonic capacities, we know that the lower bound of $\pi^2/6$ that we presented is not tight either—it is easy to raise the expected payoff beyond $\pi^2/6$ algorithm by a small constant[3] On the other hand, numerical calculations using Lemma 6 and Theorem 7 suggests that the right upper bound is probably not more than 2 (for example, for $n = 1000$, sampling 10^5 random permutations yields a numerical estimate of ≈ 1.89 for the optimal online revenue). It would be interesting to pin down the precise constant for the harmonic capacities.

Note that by a simple application of the Yao minimax principle, we can conclude that no randomized balloon popping mechanism can obtain a total expected volume more than $4.331 \max_i iv_i$ for any capacities either (on an adversarially chosen ordering of the bidders).

The balloon popping problem is motivated by our desire to understand the limitations of ascending auctions in terms of auctioneer profit maximization. It would be very interesting to extend these kinds of results to more complex ascending auctions, such as ascending combinatorial auctions.

Acknowledgments. We would like to thank Ryan O'Donnell for suggesting the $\pi^2/6$ algorithm.

[3]For example, consider the following simple modification to Strategy 2 from the introduction: If, at the time that a balloon is successfully blown up to volume 1, the balloon of capacity 1/2 has popped but the balloons of capacity 1/3 and 1/4 have not – call this event A – then blow up successive balloons to a capacity of 1/4 until two balloons have successfully been blown up to this volume. Then revert back to Strategy 2. If event A does not occur, simply continue with Strategy 2. Conditioned on event A, the payoff from the two balloons of capacity 1/3 and 1/4 is 1/2. On the other hand, if we continued with Strategy 2, our payoff from these two balloons would be $1/3 + 1/8 < 1/2$. Since event A happens with constant probability, this modification to strategy 2 increases the total payoff by a constant.

References

[1] Z. Bar-Yossef, K. Hildrum, and F. Wu. Incentive-compatible online auctions for digital goods. In *Proceedings of 13th ACM Symposium on Discrete Algorithms (SODA)*, 2002.

[2] J. Bergin and L. Zhou. Monotonic assignment rules and common pricing. *Mathematics of Operations Research*, 31, 2006.

[3] A. Blum and J. Hartline. Near-optimal online auctions. In *Proceedings of 16th ACM Symposium on Discrete Algorithms (SODA)*, 2005.

[4] A. Blum, V. Kumar, A. Rudra, and F. Wu. Online learning in online auctions. In *Proceedings of 14th ACM Symposium on Discrete Algorithms (SODA)*, 2003.

[5] L. Blumrosen and N. Nisan. On the computational power of iterative auctions II: Ascending auctions. In *Stanford Institute on Theoretical Economics Workshop on Bounded Rationality in the Design of Markets and Organizations*, 2004.

[6] J. Bulow and J. Roberts. The simple economics of optimal auctions. *Journal of Political Economy*, 97:1060–1090, 1990.

[7] J. Cremer and R.P. McLean. Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica*, 53(2):345–361, 1985.

[8] K. Deshmukh, A. Goldberg, J. Hartline, and A. Karlin. Truthful and competitive double auctions. In *Proceedings of 10th Annual European Symposium on Algorithms (ESA)*, 2002.

[9] A. Goldberg and J. Hartline. Competitiveness via consensus. In *Proceedings of 14th ACM Symposium on Discrete Algorithms (SODA)*, 2003.

[10] A. Goldberg, J. Hartline, and A. Wright. Competitive auctions and digital goods. In *Proceedings of 12th ACM Symposium on Discrete Algorithms (SODA)*, 2001.

[11] J. Hartline. *Optimization in the Private Value Model: Competitive Analysis Applied to Auction Design*. PhD thesis, University of Washington, 2003.

[12] R. Kleinberg and T. Leighton. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *Proceedings of 14th ACM Symposium on Discrete Algorithms (SODA)*, 2003.

[13] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.

[14] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.

[15] D. Parkes. Online mechanisms. In *Algorithmic Game Theory*, chapter 16. Cambridge University Press, 2007. Editors: N. Nisan and T. Roughgarden and E. Tardos and V. Vazirani. In press.