Detail-free, Posted-Price Mechanisms for Limited Supply Online Auctions

Moshe Babaioff *   Shaddin Dughmi†   Aleksandrs Slivkins*

February 2010

Abstract

We consider online posted-price mechanisms with limited supply. A seller has k items for sale and is facing n potential buyers ("agents") that are arriving sequentially. Each agent is interested in buying one item. Each agent's value for an item is an IID sample from some fixed distribution with support $[0, 1]$. The seller offers a take-it-or-leave-it price to each arriving agent (possibly different for different agents), and aims to maximize his expected revenue.

We focus on mechanisms that do not use any information about the distribution; such mechanisms are called *detail-free*. They are desirable because knowing the distribution is unrealistic in many practical scenarios. We study how the revenue of such mechanisms compares to the revenue of the optimal offline mechanism that knows the distribution ("offline benchmark").

We present a detail-free online posted-price mechanism whose revenue is within $O((k \log n)^{2/3})$, in additive terms, of the offline benchmark, for every distribution that is regular. In fact, this guarantee holds without *any* assumptions if the benchmark is relaxed to fixed-price mechanisms.

The upper bound can be improved to $O(\sqrt{k} \log n)$ for $k < \frac{n}{2e}$ under a stronger, yet quite common, assumption on the distribution: monotone hazard rate. A strong intuition from prior work suggests that one should not hope for a sufficiently general upper bound that is better than $O(\sqrt{k})$.

1 Introduction

Consider an airline (or a travel agency) that is interested in selling k seats on a plane between New York and London, at a date that is right before the 2012 London Olympics. The seller is interested in maximizing his revenue from selling these flight tickets, and is offering the tickets on a website such as Expedia. Potential buyers ("agents") arrive one after another, each with the goal of purchasing a ticket if the price is smaller than the agent's valuation. The seller expects n such agents to arrive. Whenever an agent arrives the seller presents to him a take-it-or-leave-it price, and the agent makes a purchasing decision according to that price. The seller can update the price taking into account the observed history and the number of remaining items and agents.

We adopt a Bayesian view that the valuations of the buyers are IID samples from a fixed distribution, called *demand distribution*. A standard assumption in a Bayesian setting is that the demand distribution is known to the seller, who can design a specific mechanism tailored to this knowledge. (For example, the Myerson optimal auction for one item sets a reserve price that is a function of the distribution). However, in some settings this assumption is very strong, and should be avoided if possible. For example, when the seller enters a new market, she might not know the demand distribution, and learning it through market research

*Microsoft Research Silicon Valley, Mountain View CA. {moshe,slivkins} @ microsoft.com
†Department of Computer Science, Stanford University. shaddin@cs.stanford.edu.

might be costly. Likewise, when the market has experienced a significant recent change, the new demand function might not be easily derived from the old data.

Ideally we would like to design mechanisms that perform well for any demand distribution, and yet do not rely on knowing it. Such mechanisms are called *detail-free*, in the sense that the specification of the mechanism does not depend on the details of the "environment", in the spirit of Wilson's Doctrine [34]. Learning about the demand distribution is an integral part of the problem that a detail-free mechanism faces. The performance of such mechanisms is compared to a benchmark that *does* depend on the specific demand distribution, as in [24, 21, 5].

In this paper we take this approach and design detail-free, online posted-price mechanisms with revenue that is close to the revenue of the optimal offline mechanism (that can depend on the demand distribution and is not restricted to be posted price) for two families of distributions. The guarantee we provide is either for any demand distribution that is regular, or any demand distribution that satisfy the stronger condition of Monotone Hazard Rate. Both conditions are mild and standard, and even the stronger one is satisfied by most common distributions, such as the normal, uniform, and exponential distributions.

Posted price mechanisms are appealing for many reasons.[1] First, they are commonly used in practice. Second, these mechanisms are trivially truthful (in dominant strategies) and moreover also group strategy-proof (a notion of collusion resistance when side payments are no allowed). Third, an agent only needs to evaluate his offer, which might be easier than exactly computing his private value. Fourth, agents do not reveal their entire private information to the seller: rather, they only reveal whether their private value is larger than the posted price. Further, detail-free posted-price mechanisms are particularly useful in practice as the seller is not required to estimate the demand distribution.

Next we discuss our model more formally. We consider the following limited supply auction model. A seller has k items he can sell to a set of n agents (potential buyers), aiming to maximize his expected revenue. The agents arrive sequentially to the market and the seller interacts with each agent before observing future agents (in an online manner). We make the simplifying assumption that each agent interacts with the seller only once, and the timing of the interaction cannot be influenced by the agent (this assumption is also made in other papers that consider our problem for special supply amounts [28, 5, 11] and is also standard in the literature on auctions based on secretary algorithms [25, 6]). Each agent i ($1 \leq i \leq n$) is interested in buying one item, and has a private value v_i for an item. The private values are independently drawn from the same *demand distribution* F. The demand distribution F is *unknown* to the seller, but it is known that F has support in $[0, 1]$.[2]

Whenever agent i arrives to the market the seller offers him a price p_i for an item. The agent buys the item if and only if $v_i \geq p_i$, and in case he buys the item he pays p_i (so the mechanism is trivially truthful). The seller never learns the exact value of v_i, he only observes the agent's binary decision to buy the item of not. We are interested in designing such online posted-price mechanisms with high revenue compared to a natural benchmark, with minimal assumption on the demand distribution. Our benchmark is the revenue of the offline optimal mechanism that is allowed to use the demand distribution. Note that the offline mechanism that is optimal is well characterized, it is the Myerson Auction [32] (which *does* depend on knowledge of the demand distribution). The defined benchmark is a very strong benchmark, it has the following advantages over our mechanism: it is allowed to use the demand distribution, it is not constrained to posted prices and is not constrained to run online.

[1] Similar argument have been made in prior work, e.g. [18].

[2] Assuming that $\texttt{support}(F) \subset [0, 1]$ is without loss of generality (by normalizing) as long as the seller knows an upper bound on the support.

1.1 Our Contribution

We present detail-free, online posted-price mechanisms with revenue that is close to the revenue of the optimal (revenue maximizing) offline mechanism, for large families of natural distributions. All our mechanisms are deterministic and (trivially) run in polynomial time. Our main result follows.

Theorem 1.1. *There exists a detail-free, online posted-price mechanism such that for any regular demand distribution its expected revenue is within an additive term of $O((k \log n)^{2/3})$ from the expected revenue of the optimal offline mechanism.*

We emphasize that Theorem 1.1 holds for a mechanism that does *not* know the demand distribution. The mechanism is trivially truthful as it is a posted price mechanism.

The proof of Theorem 1.1 consists of two stages. The first stage (immediate from Yan [35]) is to observe that for any regular demand distribution the revenue of the best fixed-price mechanism[3] is close to the revenue of the optimal offline mechanism. The second stage, which is our main technical contribution, is to show that our posted price mechanism achieves revenue that is close to the revenue of the best fixed-price mechanism. Surprisingly, this holds without *any* assumptions on the demand distribution.

Theorem 1.2. *There exists a detail-free, online posted-price mechanism whose expected revenue is within an additive term of $O(k \log n)^{2/3}$ from the expected revenue of the best fixed-price mechanism. This result holds for every demand distribution.*

The mechanism in Theorem 1.2 builds on a technique in the design and analysis of a multi-armed bandit algorithm in [2]. This technique, and even the intuition behind it, are not directly applicable. The main conceptual challenge is to re-invent this technique in the limited supply setting, and recover the key property that it does not separate "exploration" and "exploitation". (This property results in a much more efficient "exploration" of suboptimal prices; see Section 4.1 for further discussion.) As such, our work contributes to the literature on multi-armed bandits, which may be of independent interest.

The bounds in Theorem 1.1 and Theorem 1.2 can be improved to $O(c_F \cdot \sqrt{k} \log n)$ for $k < \frac{n}{2e}$, where c_F is a constant that depends on the distribution F, under a stronger assumption on the demand distribution: Monotone Hazard Rate (MHR). This assumption is quite common in the literature; it is satisfied by many natural distributions, e.g. uniform, exponential and normal distributions. The fact that the upper bound dependence on k is of the rate $\sqrt{k}$ is particularly interesting due to the matching lower bound from prior work [28, 11] for the case $k = \Omega(n)$. [4] These bounds provide a strong intuition that, informally, one should not hope for a sufficiently general upper bound that is better than $O(\sqrt{k})$. It should be noted that for some distributions F the constant c_F can be very large.

The bounds in Theorem 1.1 and Theorem 1.2 are uninformative when $k = O(\log^2 n)$. We next provide another detail-free, online posted-price mechanism that gives meaningful bounds in the case that k is very small (but bigger than some constant). Assuming MHR, we show that its expected revenue is within $O(k^{3/4} \operatorname{poly} \log(k))$ of the maximal expected revenue of any offline mechanism.

1.2 Related Work

Special cases. Several special cases of our setting have been studied in [28, 5, 11].

First, Kleinberg and Leighton [28] consider the unlimited supply case. Among other results, they study IID valuations, i.e. our setting with $k = n$. They provide matching upper and lower bounds on regret,

[3] A fixed-price mechanism is a posted-price mechanism that offers the same price to all agents, as long as it has items to sell. The "best" fixed-price mechanism is one with the maximal expected revenue.

[4] The $\Omega(\sqrt{n})$ lower bounds in [11, 28] hold even if the demand distributions are constrained to satisfy some non-degeneracy and smoothness conditions; the conditions in the two papers are incomparable. The result in [28] only applies to the case $k = n$.

of order $\sqrt{n}$. [5] The upper bound assumes a version of regularity, and depends on a distribution-specific constant. This is similar to our $O(c_F\sqrt{k})$ result for MHR demand distributions. (Our result assumes $\frac{k}{n} < \frac{1}{2e}$ and hence does not subsume theirs.) Absent any assumptions, the upper bound analysis [28] results in regret $O(n^{2/3})$, which is subsumed by Theorem 1.2.

On the other extreme, Babaioff et al. [5] consider the case that the seller has only one item to sell ($k = 1$). They provide a super-constant multiplicative lower bound for unrestricted demand distribution (with respect to the online optimal mechanism), and a constant-factor approximation assuming MHR. Note that we also use MHR to derive bounds that apply to the case of a very small k.

Finally, Besbes and Zeevi [11] consider a technically different, continuous-time version. It can be specialized to the discrete time, and then it is (essentially) equivalent to our setting with $k = \Omega(n)$. They derive a $\Omega(\sqrt{n})$ lower bound on regret. Further, they provide a number of upper bounds on regret with respect to the fixed-price benchmark, assuming that the demand distribution $F(\cdot)$ and its inverse $F^{-1}(\cdot)$ are Lipschitz-continuous. Without any extra assumptions, they achieve regret $O(n^{3/4})$, using a mechanism with separate exploration and exploitation phases. They improve it to $O(n^{2/3})$ for demand distributions that are parameterized (in a way that is known to the algorithm), and to $O(\sqrt{n})$ if furthermore they depend on a single parameter. Both results rely on knowing the parametrization: the mechanisms continuously update the estimates of the parameter(s) and revise the current price according to these estimates.

The upper bounds in [11] should be contrasted with our $O(k^{2/3})$ upper bound that applies to an arbitrary k and makes no assumptions on the demand distribution, and the $O(c_F\sqrt{k})$ improvement for MHR demand distributions.

Online mechanisms. The study of online mechanisms was initiated by Lavi and Nisan [30], which unlike us consider the case that each agent is interested in multiple items, and provide a logarithmic multiplicative approximation. Below we survey only the most relevant papers in this line of work, in addition to the special cases of our setting that we have already discussed.

Several papers [9, 13, 28, 12] consider online mechanisms with unlimited supply and adversarial valuations (as opposed to limited supply and IID valuations in our setting). The mechanism in the initial paper [9] requires the agents to submit bids and so is not posted-price. The subsequent work [13, 28, 12] provides various improvements. In particular, Blum et al. [13] (among other results) design a simple *posted-price* mechanism which achieves multiplicative approximation $1+\epsilon$, for any $\epsilon > 0$, with an additive term that depends on ϵ. [6] Blum and Hartline [12] use a more elaborate posted-price mechanism to improve the additive term. Kleinberg and Leighton [28] show that the simple mechanism in [13] achieves regret $O(n^{2/3})$; moreover, they provide a nearly matching lower bound of $\Omega(n^{2/3})$.

Papers [23, 19] study online mechanisms for limited supply and IID valuations (same as us), but their mechanisms are not posted-price. Hajiaghayi et al. [23] consider an online auction model where players arrive and depart online, and may misreport the time period during which they participate in the auction. This makes designing strategy-proof mechanisms more challenging, and as a result their mechanisms achieve a constant multiplicative approximation rather than additive regret. Moreover, their revenue benchmark is incomparable with ours.[7] Devanur and Hartline [19] study several variants of the limited-supply mechanism design problem: supply is known or unknown, online or offline. Most related to our paper is their mechanism for limited, known, online supply. This mechanism is based on random sampling and achieves constant (multiplicative) approximation, but is not posted-price. Our mechanism is posted-price and achieves low (additive) regret.

[5] Throughout this section, we omit the log factors in regret bounds.

[6] This result considers valuations in the range $[1, H]$, and the additive term also depends on H.

[7] The revenue benchmark in [23] is in terms of the realized valuations whereas ours is in expectation over the prior. On the other hand, their benchmark requires selling at least two items, which is not necessarily optimal.

Other work. Absent the supply constraint, our problem (and a number of related formulations) fit into the *multi-armed bandit* (MAB) framework.[8] MAB has a rich literature in Operations Research, Computer Science and Economics. A proper discussion of this literature is beyond the scope of this paper; a reader can refer to [16, 10] for background. Most relevant to our specific setting is the work on (prior-free) MAB with stochastic payoffs, e.g. [29, 2], and MAB with Lipschitz-continuous stochastic payoffs, e.g. [1, 26, 4, 27, 15]. The posted-price mechanisms in [13, 28, 12] described above are based on a well-known MAB algorithm [3] for adversarial payoffs. The connection between online learning and online mechanisms has been explored in a number of other papers, including [33, 20, 8, 7].

Recently, [18, 17, 35] studied the problem of designing an offline, sequential posted-price mechanisms in Bayesian settings, where the distributions of valuations are not necessarily identical, yet are known to the seller. Chawla et al. [18] provide constant multiplicative approximations. Yan [35] obtains a multiplicative bound that is optimal for large k, and Chakraborty et al. [17] obtain a PTAS for all k.

2 Preliminaries

A (monopolist) seller has k items to sell. Potential buyers (agents) arrive online (sequentially): agent t arrives in round t. We assume that the seller knows that n buyers will arrive. The seller sells his supply using an *online posted-price* mechanism. Such a mechanism is essentially an online algorithm which in each round outputs a price and observes whether there was a sale.

Throughout, we assume that agents' valuations are drawn independently from a distribution F with support in $[0,1]$, called *demand distribution*. We use $p \in [0,1]$ to denote a price. We let $F(p)$ denote the c.d.f, $f(p)$ denote the p.d.f and $S(p) = 1 - F(p)$ denote the survival rate at price p. Let $R(p) = p\,S(p)$ denote the revenue at price p. The demand distribution is called *regular* if $R(F^{-1}(\alpha))$ is a concave function of cumulative probability α. We call it *strictly regular* if furthermore $R(\cdot)$ has a unique maximizer.[9] We say F is a *Monotone Hazard Rate (MHR)* distribution if the hazard rate $H(p) = f(p)/S(p)$ is monotone non-decreasing. All MHR distributions are regular.

A mechanism is called *detail-free* if it does not use the knowledge of the demand distribution. We are interested in designing detail-free, online posted-price mechanisms with good performance for *every* demand distribution in some (large) family of distributions. We compare our mechanisms to two benchmarks: the maximal revenue of an offline mechanism (the *offline benchmark*), and the maximal revenue of a fixed price mechanism (the *fixed-price benchmark*). An offline mechanism with a maximal revenue was given by the seminal paper of Myerson [32] (it is not an online posted price mechanism).

A *fixed-price mechanism* with n agents, k items and price p, denoted $\mathcal{A}^n_k(p)$, is the mechanism that makes a fixed offer price p to every agent so long as fewer than k items have been sold, and stops afterwards (equivalently, from that point always sets the price to ∞). Note that $\mathcal{A}^n_n(p)$, the mechanism with no supply constraint, sells $S(p)n$ items in expectation.

Let $\texttt{Rev}(\mathcal{A})$ be the total expected revenue achieved by mechanism $\mathcal{A}$. We define the *regret* of $\mathcal{A}$ with respect to the fixed-price benchmark as follows.

$$\texttt{Regret}(\mathcal{A}) \triangleq \max_p \texttt{Rev}[\mathcal{A}^n_k(p)] - \texttt{Rev}(\mathcal{A}).$$

Thus, regret is the additive loss in expected revenue compared to the best fixed-price mechanism.

[8]To void a possible confution, we note that the supply constraint in our setting may appear similar to the budget constraint in line of work on *budgeted MAB* (see [14, 22] for details and further references). However, the "budget" in budgeted MAB is essentially the duration of the experimentation phase (n), rather than the number of rounds with positive reward (k).

[9]This maximizer is called the *Myerson Reserve Price* and denoted p_{r}. The revenue function $R(p)$ is non-decreasing for $p \leq p_{\mathrm{r}}$, and non-increasing when $p \geq p_{\mathrm{r}}$.

3 Benchmarks Comparison

We start by observing that for regular demand distributions, the fixed-price benchmark is close to the offline benchmark; this result is immediate from Yan [35].

Lemma 3.1 (Yan [35])**.** *Assume that the demand distribution is regular. Then there exists a fixed-price mechanism whose expected revenue is at least the optimal offline revenue minus $O(\sqrt{k \log k})$.*

Remark. We provide a self-contained proof in Appendix B.

Remark. Lemma 3.1 implies that any mechanism with regret $O(R)$, $R = \Omega(\sqrt{k \log k})$ with respect to the fixed-price benchmark has the same asymptotic regret $O(R)$ with respect to the offline benchmark, as long as the demand distribution is regular, and in particular if it is MHR. Therefore, the rest of the paper can focus on the fixed-price benchmark. In particular, our main result, Theorem 1.1 for regular distributions, follows from Theorem 1.2 that addresses the fixed-price benchmark.

If we forego an additive factor of $O(\sqrt{k \log k})$ then the expected revenue of a fixed-price mechanism is particularly easy to characterize.

Claim 3.2. *Let $\mathcal{A}$ be the fixed-price mechanism with price p. Then*

$$\texttt{Rev}(\mathcal{A}) \geq p \min(k, n\, S(p))) - O(p\sqrt{k \log k}). \tag{1}$$

Proof. Let X_t be the indicator variable of sale in round t. Denote $X = \sum_{t=1}^{n} X_t$ and let $\mu = \mathbb{E}[X]$. Then by Chernoff Bounds (Theorem A.1(a))

$$\Pr[X \geq \mu - O(\sqrt{\mu \log k})] \geq 1 - \tfrac{1}{k}.$$

Therefore with probability at least $1 - \frac{1}{k}$ it holds that

$$\begin{aligned} \#\text{sales} &= \min(k, X) \\ &\geq \min(k, \mu - O(\sqrt{\mu \log k})) \\ &\geq \min(k, \mu) - O(\sqrt{k \log k}), \end{aligned}$$

which implies the claim since $\mu = n\, S(p)$. □

We use this fact in Section 4. Moreover, we can now characterize a near-optimal fixed price. This characterization provides intuition for the rest of the paper, and it is used directly in Section 6.

Lemma 3.3. *The bound in Lemma 3.1 is satisfied for the fixed price $p^* = \max(r,\, S^{-1}(\frac{k}{n}))$, where $r = \operatorname{argmax}_p p\, S(p)$ is the Myerson reserve price.*

4 The main technical result

This section is devoted to the main technical result: Theorem 1.2 for the fixed-price benchmark. This result is very general, as it makes no assumptions on the demand distribution. Let us restate it here for convenience:

Theorem 4.1. *There exists a detail-free, online posted-price mechanism whose regret with respect to the fixed-price benchmark is at most $O(k \log n)^{2/3}$.*

Remarks. Theorem 4.1 provides a non-trivial bound as long as $k > \Omega(p^{-3})(\log^2 n)$, where p is the price such that $S(p) = \frac{k}{n}$. This is because by Claim 3.2 the expected revenue from the fixed-price mechanism with this price is at least $kp - O(\sqrt{k \log k})$.

4.1 High-level discussion

Absent the supply constraint, our problem fits into the *multi-armed bandit* (MAB) framework: in each round, an algorithm chooses among a fixed set of alternatives ("arms") and observes a payoff, and the objective is to maximize the total payoff over a given time horizon. Our setting corresponds to MAB with *stochastic payoffs*: in each round, the payoff is an independent sample from some unknown distribution that depends on the chosen "arm" (price).

This connection is exploited in [28] for the special case with unlimited supply ($k = n$). The authors use a standard algorithm for MAB with stochastic payoffs, called `UCB1` [2]. Specifically, they focus on the prices $\{i\delta : i \in \mathbb{N}\}$, for some parameter δ, and run `UCB1` with these prices as "arms". The regret bound from [2] applies directly (although some additional work is required to convert it into the final result).

Unfortunately, neither the analysis nor the intuition behind `UCB1` and similar MAB algorithms is directly applicable to the setting with limited supply. Informally, the goal of an MAB algorithm is to converge to an arm with the highest average payoff, whenever such "best arm" exists. This, of course, is a wrong approach if the supply is limited: if selling at the "best price" quickly exhausts the inventory, then a higher price is more profitable.

Our main conceptual challenge is to recover, for the limited supply setting, the appealing feature of `UCB1` that it does not separate "exploration" and "exploitation": it explores arms according to a schedule that continuously adapts to the observed payoffs, rather than is fixed according to some pre-defined parameters. This way `UCB1` ensures that (very) suboptimal arms are chosen (very) rarely even while they are being "explored", which immediately translates into advantageous bounds on regret.

Following `UCB1`, we would like to assign each arm a numerical score, called *index*, so that an arm with the highest index is picked. In `UCB1`, the index is, essentially, the best Upper Confidence Bound (UCB) on the expected payoff of this arm that is available at a given time. The fact that the UCB depends on both the average payoff and the number of times this arm has been played so far provides the desired exploration-exploitation combination.

The main technical hurdles are as follows. First, we need to define a statistic that, unlike the average, reflects the limited inventory, and also takes into account the number of samples. Then we need to deduce that a price with the highest index is, in some useful sense, comparable to the price that is optimal given the inventory size. Finally, we need to find way to "charge" each suboptimal price for each time that it is chosen, in a way that the total regret is bounded by the sum of these charges and the sum can be "tamed" in the proof. The analysis of `UCB1` overcomes these hurdles via simple (but very elegant) tricks which, unfortunately, do not extend to the limited supply setting. We address these challenges one by one in the rest of this section.

4.2 Our mechanism

Let us define our mechanism, called `CappedUCB`. The mechanism is initialized with a set $\mathcal{P}$ of "active prices". In each round t, some price $p \in \mathcal{P}$ is chosen. For each $p \in \mathcal{P}$, let $N_t(p)$ be the number of times price p has been chosen up to round t, and let $N(p)$ be the total number of times this price is chosen. Let $k_t(p)$ be the number of items sold at price p up to time t. Let $\widehat{S}_t(p) \triangleq k_t(p)/N_t(p)$ be the average survival rate for price p up to time t.

A *confidence radius* is some number $r_t(p)$ such that

$$|S(p) - \widehat{S}_t(p)| \leq r_t(p) \quad (\forall\, p \in \mathcal{P}, t \leq n). \tag{2}$$

holds with high probability, namely with probability at least $1 - n^{-2}$. We will use it in round t of the mechanism, so it is essential for $r_t(p)$ to be defined in terms of quantities that are observable at time t, such as $N_t(p)$ and $\widehat{S}_t(p)$. A standard confidence radius used in the literature is (essentially) $r_t(p) = \sqrt{\frac{O(\log n)}{N_t(p)}}$.

We will use a somewhat non-standard confidence radius from [27] (see Lemma A.2) which performs better for prices with low survival rate:

$$r_t(p) \triangleq \frac{O(\log n)}{N_t(p)} + \sqrt{\frac{O(\log n)\, \widehat{S}_t(p)}{N_t(p)}}. \tag{3}$$

Note that $\widehat{S}_t(p) + r_t(p)$ is an Upper Confidence Bound (UCB) for $S(p)$.

By Claim 3.2, the expected revenue from the fixed-price auction $\mathcal{A}_k^n(p)$ can be approximated by

$$\nu(p) \triangleq p \cdot \min(k,\, n\, S(p)).$$

In each round t, we define the *index* $I_t(p)$ as an UCB on $\nu(p)$ that is available at this time:

$$I_t(p) \triangleq p \cdot \min\left(k,\, n\, \left(\widehat{S}_t(p) + r_t(p)\right)\right). \tag{4}$$

In each round, the mechanism picks a price with the maximal index, breaking ties arbitrarily. Once k items have been sold the mechanism always sets the price to ∞ and never sells any additional item.

We will use the active prices given by

$$\mathcal{P} = \{\delta(1+\delta)^i \in [0,1] :\, i \in \mathbb{N}\}, \tag{5}$$

for some parameter $\delta \in (0,1)$.

4.3 Analysis of the mechanism

Our goal is to bound from above the *regret* of `CappedUCB`, which is the difference between the optimal expected revenue of a fixed-price mechanism and the expected revenue of `CappedUCB`. We prove that `CappedUCB` achieves regret $O(k \log n)^{2/3}$ for a suitable choice of parameter δ in (5).

Lemma 4.2. *Mechanism* `CappedUCB` *with parameter* $\delta = k^{-1/3}\,(\log n)^{2/3}$ *achieves regret* $O(k \log n)^{2/3}$.

Since the bound in Lemma 4.2 is trivial for $k < \log^2 n$, we will assume that $k \geq \log^2 n$ from now on. Let `RRev` denote the realized revenue of `CappedUCB` (revenue that is realized in a given execution).

Note that `CappedUCB` "exits" (sets the price to ∞) after it sells k items. For a thought experiment, consider a version of this mechanism that does not "exit" and continues running as if it has unlimited supply of items; let us call this version `CappedUCB`$'$. Then the revenue of `CappedUCB` is exactly equal to the revenue obtained by `CappedUCB`$'$ from selling the first k items. Thus from here on we focus on analyzing the latter.

Let X_t be the indicator variable of the random event that there is a sale in round t. Then

$$\texttt{RRev} = \textstyle\sum_{t=1}^{N} p_t\, X_t, \qquad \text{where } N = \max\{N \leq n :\, \sum_{t=1}^{N} X_t \leq k\}. \tag{6}$$

Let $X \triangleq \sum_{t=1}^{n} X_t$ be the total number of sales if the inventory were unlimited.

High-probability events. We tame the randomness inherent in the sales X_t by setting up three high-probability events, as described below. (The probability bounds are derived via appropriate tail inequalities, see Appendix A.) In the rest of the analysis, we will argue deterministically under the assumption that these three events hold. It suffices because the expected loss in `RRev` from the low-probability failure events will be negligible.

First, for each t, X_t is a 0-1 random variable with expectation $S(p_t)$, where p_t depends on $X_1, \ldots, X_{t-1}$. Let $S = \sum_{t=1}^n S(p_t)$. Using the appropriate tail bound (Lemma A.3 with $\alpha_t \equiv 1$) we obtain that

$$|X - S| < O(\sqrt{S \log n} + \log n) \tag{7}$$

holds with probability at least $1 - n^{-2}$.

Second, taking Lemma A.3 with $\alpha_t = p_t$ we obtain that

$$|\textstyle\sum_{t=1}^n p_t(X_t - S(p_t))| < O(\sqrt{S \log n} + \log n) \tag{8}$$

holds with probability at least $1 - n^{-2}$.

Third, let us prove that (3) is a confidence radius, i.e. that (2) holds with high probability. Indeed, for each price $p \in \mathcal{P}$, let $\{Z_{i,p}\}_{i \le n}$ be a family of independent 0-1 random variables with expectation $S(p)$. Without loss of generality, let us pretend that the i-th time that price p is selected by the mechanism, sale happens if and only if $Z_{i,p} = 1$. Then by Lemma A.2 after the i-th play of price p the bound (2) holds with probability at least $1 - n^{-4}$. Taking the Union Bound over all choices of i and all choices of p, we obtain that (2) holds with probability at least $1 - n^{-2}$ as long as $|\mathcal{P}| \le n$ (which is the case for us).

From now on, we will assume that (2), (8) and (7) hold.

Single-round analysis. Let us analyze what happens in a particular round t of the mechanism. For each price $p \in \mathcal{P}$ and each round t we have

$$\nu(p) \le I_t(p) \le p \cdot \min\left(k,\, n\, \left(S(p) + 2\, r_t(p)\right)\right). \tag{9}$$

Let p_t be the price chosen in round t. Let $p^*_{\text{act}} \in \operatorname{argmax}_{p \in \mathcal{P}} \nu(p)$ be the best active price (according to $\nu(\cdot)$), and let $\nu^*_{\text{act}} \triangleq \nu(p^*_{\text{act}})$. Then

$$\begin{cases} I_t(p_t) \ge I_t(p^*_{\text{act}}) \ge \nu(p^*_{\text{act}}) \triangleq \nu^*_{\text{act}} \\ I_t(p_t) \le p_t \cdot \min\left(k,\, n\, \left(S(p_t) + 2\, r_t(p_t)\right)\right). \end{cases}$$

Combining these two inequalities, we obtain the key inequality:

$$\tfrac{1}{n}\, \nu^*_{\text{act}} \le p_t \cdot \min\left(\tfrac{k}{n},\, S(p_t) + 2\, r_t(p_t)\right). \tag{10}$$

There are several consequences. First, it follows that

$$\begin{cases} p & \ge \frac{1}{k}\, \nu^*_{\text{act}}, \quad \forall p \in \mathcal{P}_{\text{sel}} \\ \Delta(p_t) & \triangleq \frac{1}{n}\, \nu^*_{\text{act}} - p_t\, S(p_t) \le 2\, p_t\, r_t(p_t). \end{cases} \tag{11}$$

Here $\mathcal{P}_{\text{sel}} \triangleq \{p \in \mathcal{P} : N(p) \ge 1\}$ is the set of active prices that have been selected at least once. The notation $\Delta(p) \triangleq \frac{1}{n}\, \nu^*_{\text{act}} - p\, S(p)$ corresponds to the "badness" of price p in a single round, if the bounded inventory size is not taken into account. We will use this notation throughout the analysis: eventually, we will bound regret in terms of $\sum_{p \in \mathcal{P}} \Delta(p)\, N(p)$.

Another consequence of (10) is that

$$\forall p \in \mathcal{P}_{\text{sel}}: \quad \Delta(p) > 0 \Rightarrow S(p) < \tfrac{k}{n}. \tag{12}$$

Indeed, if $\Delta(p) > 0$ for some $p \in \mathcal{P}_{\text{sel}}$, then $pk \ge \nu^*_{\text{act}} > np\, S(p)$, which implies (12).

Note that we have not yet used the definition (3) of the confidence radius. For a given price $p \in \mathcal{P}_{\text{sel}}$, let t be the last round in which this price has been selected by the mechanism. Then using (11) to bound $\Delta(p)$, Lemma A.2 to bound the confidence radius $r_t(p)$, and (12) to bound the survival rate, we obtain:

$$\Delta(p) \leq O(p) \times \max\left(\tfrac{\log n}{N(p)},\; \sqrt{\tfrac{k}{n}\,\tfrac{\log n}{N(p)}}\right). \tag{13}$$

Now we can bound $N(p)$ in terms of $\Delta(p)$:

$$\begin{aligned} N(p) &\leq O(\log n) \times \max\left(\tfrac{p}{\Delta(p)},\; \tfrac{k}{n}\,\tfrac{p^2}{\Delta^2(p)}\right) \\ N(p)\,\Delta(p) &\leq O(\log n)\left(1 + \tfrac{k}{n}\,\tfrac{1}{\Delta(p)}\right). \end{aligned} \tag{14}$$

Analyzing the total revenue. For brevity, let

$$\beta(S) = O(\sqrt{S \log n} + \log n)$$

to be the right-hand side in (7) and (8).

Claim 4.3. $\texttt{RRev} \geq \min(\nu^*_{act},\; \sum_{t=1}^n p_t\, S(p_t)) - \beta(k)$.

Proof. Recall that $p_t \geq \frac{1}{k}\nu^*_{\text{act}}$ by (11). It follows that $\texttt{RRev} \geq \nu^*_{\text{act}}$ whenever $\sum_{t=1}^n X_t > k$. Therefore, if $\texttt{RRev} < \nu^*_{\text{act}}$ then $\sum_{t=1}^n X_t \leq k$ and so $\texttt{RRev} = \sum_{t=1}^n p_t\, X_t$. Thus,

$$\begin{aligned} \texttt{RRev} &\geq \min\left(\nu^*_{\text{act}},\; \textstyle\sum_{t=1}^n p_t\, X_t\right) \\ &\geq \min\left(\nu^*_{\text{act}},\; \textstyle\sum_{t=1}^n p_t\, S(p_t) - \beta(S)\right) \end{aligned}$$

So the claim holds when $S \leq k$.

On the other hand, if $S > k$ then

$$\begin{aligned} X &\geq S - \beta(S) \geq k - \beta(k) \\ \texttt{RRev} &\geq \min(k, X)\,(\tfrac{1}{k}\,\nu^*_{\text{act}}) \geq \nu^*_{\text{act}} - \beta(k), \end{aligned}$$

completing the proof. □

In light of Claim 4.3, we can focus on $\sum_{t=1}^n p_t\, S(p)$, effectively ignoring the capacity constraint.

$$\begin{aligned} \textstyle\sum_{t=1}^n p_t\, S(p_t) &= \textstyle\sum_{t=1}^n \tfrac{1}{n}\,\nu^*_{\text{act}} - \Delta(p_t) \\ &= \nu^*_{\text{act}} - \textstyle\sum_{t=1}^n \Delta(p_t) \\ &= \nu^*_{\text{act}} - \textstyle\sum_{p \in \mathcal{P}} \Delta(p)\, N(p). \end{aligned} \tag{15}$$

Fix $\epsilon > 0$ and let $\mathcal{P}_\epsilon \triangleq \{p \in \mathcal{P}_{\text{sel}} :\; \Delta(p) \geq \epsilon\}$. Then plugging in (14), we obtain

$$\begin{aligned} &\textstyle\sum_{p\in\mathcal{P}} \Delta(p)\, N(p) \\ &\quad\leq \textstyle\sum_{p \in \mathcal{P}\setminus\mathcal{P}_\epsilon} \Delta(p)\, N(p) + \sum_{p\in\mathcal{P}_\epsilon} \Delta(p)\, N(p) \\ &\quad\leq \epsilon n + O(\log n) \textstyle\sum_{p\in\mathcal{P}_\epsilon} \left(1 + \tfrac{k}{n}\tfrac{1}{\Delta(p)}\right) \\ &\quad\leq \epsilon n + O(\log n)\left(|\mathcal{P}_\epsilon| + \tfrac{k}{n}\textstyle\sum_{p\in\mathcal{P}_\epsilon} \tfrac{1}{\Delta(p)}\right). \end{aligned} \tag{16}$$

Combining this with Claim 4.3 yields a claim that summarizes our findings so far.

Claim 4.4. *For any set $\mathcal{P}$ of active prices and any parameter $\epsilon > 0$ it holds that*

$$\nu^*_{act} - \mathbb{E}[\texttt{RRev}] \leq \epsilon n + O(\log n) \left(|\mathcal{P}_\epsilon| + \tfrac{k}{n} \sum_{p \in \mathcal{P}_\epsilon} \tfrac{1}{\Delta(p)} \right) + \beta(k).$$

Now let us use the fact that the active prices are given by (5) for some $\delta \in (0,1)$. Recall that $\nu^* \triangleq \max_p \nu(p)$. Let $p^* \in \operatorname{argmax}_p \nu(p)$ denote the best fixed price with respect to $\nu(\cdot)$, ties broken arbitrarily. If $p^* \leq \delta$ then $\nu^* \leq \delta k$. Else, letting $p_0 = \max\{p \in \mathcal{P} :\, p \leq p^*\}$ we have $p_0/p \geq \frac{1}{1+\delta} \geq 1 - \delta$, and so

$$\nu^*_{\text{act}} \geq \nu(p_0) \geq \tfrac{p_0}{p^*}\, \nu(p^*) \geq \nu^*(1-\delta) \geq \nu^* - \delta k.$$

It follows that for any $\epsilon > 0$ and $\delta \in (0,1)$ we have:

$$\texttt{Regret} \leq O(\log n) \left(|\mathcal{P}_\epsilon| + \tfrac{k}{n} \sum_{p \in \mathcal{P}_\epsilon} \tfrac{1}{\Delta(p)} \right) + \epsilon n + \delta k + \beta(k). \tag{17}$$

Plugging in $\Delta(p) \geq \epsilon$ for each $p \in \mathcal{P}_\epsilon$ in (17), we obtain:

$$\texttt{Regret} \leq O(|\mathcal{P}_\epsilon| \log n) \left(1 + \tfrac{1}{\epsilon}\, \tfrac{k}{n}\right) + \epsilon n + \delta k + \beta(k).$$

Note that $|\mathcal{P}| \leq \frac{1}{\delta} \log n$. To simplify the computation, we will assume upfront that $\delta \geq \frac{1}{n}$ and $\epsilon = \delta\, \frac{k}{n}$. Then

$$\texttt{Regret} \leq O\left(\delta k + \tfrac{1}{\delta^2} (\log n)^2 + \sqrt{k \log n} \right). \tag{18}$$

Finally, it remains to pick δ to minimize (18). Let us simply take δ such that the first two summands are equal: $\delta = k^{-1/3}\, (\log n)^{2/3}$. Then the two summands are equal to $O(k \log n)^{2/3}$. This completes the proof of Lemma 4.2.

5 Improved regret bounds

We show that the mechanism from Section 4 satisfies an improved regret bound, $O(\sqrt{k} \log n)$, for monotone hazard rate (MHR) demand distributions. Unlike the main result, this bound depends on a distribution-specific constant.

Remark. Prior work [28] and [11] provides (essentially) $\Omega(\sqrt{n})$ lower bounds on regret, for $k = n$ and $k = \Omega(n)$ respectively. Both lower bounds hold even if a number of non-degeneracy and smoothness conditions are enforced. Therefore there is a strong intuition that any sufficiently general upper bound of the form $k^\gamma \operatorname{polylog}(n)$ must have $\gamma \geq \frac{1}{2}$.

The distribution-specific constant involves the function

$$g(s) \triangleq s\, S^{-1}(s) : [S(1), 1] \to [0,1].$$

Recall that MHR implies regularity: $g''(\cdot) \leq 0$.

Theorem 5.1. *Assume $\frac{k}{n} < \frac{1}{2e}$. Consider any demand distribution that is non-degenerate and satisfies MHR. For any such distribution, the mechanism* `CappedUCB` *with parameter $\delta = k^{-1/2} \log(n)$ achieves regret*

$$O(\sqrt{k} \log n)(1 + 1/g'(\tfrac{k}{n})).$$

Remark. It holds that $g'(\frac{k}{n}) > 0$. This easily follows from regularity and the fact that any maximizer of $g(\cdot)$ is at least $\frac{1}{e}$ (Claim D.2). Moreover, $g'(\frac{k}{n}) \geq \Omega(\inf |g''(\cdot)|)$.

The remainder of this section is devoted to proving Theorem 5.1. We will use MHR to obtain a lower bound on $\Delta(p)$, which results in savings in (17), which in turn implies the improved regret bound.

Recall the notation from Section 4.3. Let $C = g'(\frac{k}{n})$. Note that by regularity $g'(s) \geq C$ for any $s \leq \frac{k}{n}$. Let $p^* = S^{-1}(\frac{k}{n})$ and $p \in \mathcal{P}_\epsilon$. Note that by (12) it holds that $S(p) < \frac{k}{n}$ and consequently $p > p^*$.

First, we claim that $S(p) < \frac{p^*}{p}\frac{k}{n}$. Indeed, this is because $p\,S(p) = g(S(p)) < g(\frac{k}{n}) = p^*\,\frac{k}{n}$.

Second, we bound $\Delta(p)$ from below:

$$\begin{aligned}
\tfrac{1}{n}\,\nu^*_{\text{act}} &\geq (1-\delta)\,\tfrac{\nu^*}{n} \geq (1-\delta)\,g(\tfrac{k}{n}) \\
\Delta(p) &\geq (1-\delta)\,g(\tfrac{k}{n}) - g(p) \\
&\geq [g(\tfrac{k}{n}) - g(p)] - \delta g(\tfrac{k}{n}) \\
&\geq C(\tfrac{k}{n} - S(p)) - \delta\tfrac{k}{n}\,p^* \\
&\geq C\,\tfrac{k}{n}\,(1 - \tfrac{p^*}{p}) - \delta\tfrac{k}{n}\,p^* \\
&\geq C\,\tfrac{k}{n}\,(1 - \tfrac{p^*}{p}(1 + \tfrac{\delta}{C})).
\end{aligned}$$

Since $\mathcal{P}$ is given by (5), it holds that

$$\mathcal{P}_\epsilon \subset \{p^*\alpha\,(1+\delta)^i :\, i \in \mathbb{N}\}$$

for some $\alpha \geq 1$. Define

$$\mathcal{P}' \triangleq \{p \in \mathcal{P}_\epsilon :\, p = p^*\alpha\,(1+\delta)^i \text{ with } i \geq \tfrac{2}{C}\}.$$

Then for any $p \in \mathcal{P}'$ it holds that

$$\begin{aligned}
p/p^* &= \alpha(1+\delta)^i \geq 1 + i\delta \\
\Delta(p) &\geq C\,\tfrac{k}{n}\,(1 - \tfrac{1+\delta/C}{1+i\delta}) \geq \tfrac{C}{2}\,\tfrac{k}{n}\,\tfrac{i\delta}{1+i\delta}.
\end{aligned}$$

Therefore, noting that $|\mathcal{P}'| \leq |\mathcal{P}| \leq O(\frac{1}{\delta}\,\log\frac{1}{\delta})$,

$$\begin{aligned}
\tfrac{k}{n}\sum_{p\in\mathcal{P}'}\tfrac{1}{\Delta(p)} &\leq \tfrac{2}{C}\,\sum_{p\in\mathcal{P}'}(1 + \tfrac{1}{i\delta}) \\
&\leq \tfrac{2}{C}\,(|\mathcal{P}'| + \tfrac{1}{\delta}\,\log|\mathcal{P}'|) \\
&\leq O(\tfrac{1}{C}\tfrac{1}{\delta}\,\log\tfrac{1}{\delta}) \\
\sum_{p\in\mathcal{P}_\epsilon\setminus\mathcal{P}'}\tfrac{1}{\Delta(p)} &\leq \tfrac{1}{\epsilon}\,|\mathcal{P}\setminus\mathcal{P}'| \leq \tfrac{1}{\epsilon}\,(\tfrac{2}{C} + 1).
\end{aligned}$$

Plugging this into (17) with $\epsilon = \delta\,\frac{k}{n}$, we have

$$\begin{aligned}
\tfrac{k}{n}\sum_{p\in\mathcal{P}_\epsilon}\tfrac{1}{\Delta(p)} &\leq O(\tfrac{1}{\delta}\,\log\tfrac{1}{\delta})(1 + \tfrac{1}{C}) \\
\texttt{Regret} &\leq O(\delta k + \tfrac{1}{\delta}(1 + \tfrac{1}{C})(\log n)^2 + \sqrt{k\log n}) \\
&\leq O(\sqrt{k}\log n)(1 + \tfrac{1}{C})
\end{aligned}$$

for $\delta = k^{-1/2}\,\log n$. Thus, we proved Theorem 5.1.

Mechanism 1 Descending prices

Parameter: Approximation parameters $\delta, \epsilon \in [0,1]$

Let $\alpha = \left(\frac{k}{n}\right)^{1-\delta}$, $\gamma = \min(\alpha, 1/e)$.
$\ell \leftarrow 0,\ \ell_{\max} \leftarrow 0,\ R_{\max} \leftarrow 0$.
repeat
  $\ell \leftarrow \ell + 1,\ p_\ell \leftarrow (1+\delta)^{-\ell}$
  Offer price p_ℓ to $m = \lceil \delta \frac{n}{\log_{1+\delta}(1/\epsilon)} \rceil$ agents.
  Let S_ℓ be the fraction of them who accept.
  Let $R_\ell = p_\ell S_\ell$ be the average per agent revenue.
  If $S_\ell \geq (1+\delta)^{-1}\gamma$ and $R_\ell \geq R_{\max}$,
    then $R_{\max} \leftarrow R_\ell,\ \ell_{\max} \leftarrow \ell$
until $p_\ell \leq \epsilon$ or $S_\ell \geq (1+\delta)\alpha$ or $R_\ell \leq (1+\delta)^{-2} R_{\max}$
Offer price $\widetilde{p} = p_\ell$ so long as unsold items remain.

6 Selling very few items

In this section we target a case when very few items are available for sale (roughly, $k < O(\log^2 n)$), so that the bound in Theorem 1.1 becomes trivial. We provide a different mechanism whose regret does not depend on n, under the mild assumption of monotone hazard rate.

We rely on the characterization in Lemma 3.1: we look for the price $p^* = \max(r,\, S^{-1}(\frac{k}{n}))$, where $r = \operatorname{argmax}_p p\, S(p)$ is the Myerson reserve price. The mechanism proceeds as follows. It considers prices $p_\ell = (1-\delta)^\ell$, $\ell \in \mathbb{N}$ sequentially in the descending order. For each ℓ, it offers the price p_ℓ to a fixed number of agents. The loop stops once the mechanism detects that, essentially, the "best" p_ℓ has been reached: either $S(p_\ell)$ is close to $\frac{k}{n}$, or we are near a maximum of $p\, S(p)$. Parameters are chosen so as to minimize regret, see Mechanism 1.

Theorem 6.1. *For some parameters ϵ and δ, Mechanism 1 achieves regret $O\left(k^{3/4} \operatorname{poly}\log(k)\right)$ with respect to the offline benchmark, for any demand distribution that satisfies the Monotone Hazard Rate condition.*

The rest of this section is devoted to proving Theorem 6.1 for parameters $\epsilon = k^{-1/4}$ and $\delta = (\frac{1}{k}\, \log k)^{1/4}$. We will assume that the demand distribution is MHR, without further notice. We derive Theorem 6.1 from the following multiplicative bound.

Lemma 6.2. *Assume $p^* \geq \epsilon$. Set*

$$\delta = \sqrt[4]{\tfrac{1}{k}\, \log k \log \tfrac{1}{\epsilon} \log\log \tfrac{1}{\epsilon}}. \tag{19}$$

Then the expected revenue of Mechanism 1 is at least $1 - O(\delta)$ fraction of the offline benchmark.

Proof of Theorem 6.1. If $p^* \leq \epsilon$ then the maximum possible loss in revenue is ϵk. Else, using Lemma 6.2, the loss in revenue is at most $O(\delta k)$, where δ is as defined in Lemma 6.2. Thus, in general, the additive regret compared to the optimal offline revenue is at most $\max(\epsilon k, O(k\delta))$. This is (roughly) minimized by setting $\epsilon = k^{-1/4}$. □

6.1 Proof of Lemma 6.2

We use a stronger, multiplicative version of Lemma 3.1 (which is also immediate from Yan [35]). More precisely, we use a somewhat stronger result, Corollary C.3, which is proved in Appendix C using the machinery from Yan [35]. It appears difficult to circumvent the multiplicative bound in Corollary C.3 and prove the additive bound in Theorem 6.1 directly. Also, we take advantage of several properties of MHR distributions, detailed in Appendix D.

We say the exploration phase is δ-*approximate* if

$$S(p_\ell) \geq \gamma \;\Rightarrow\; \tfrac{1}{1+\delta} \leq S_\ell / S(p_\ell) \leq 1+\delta.$$

Claim 6.3. *The exploration phase is δ-approximate with probability at least* $1 - 2\left(\log_{1+\delta} \frac{1}{\epsilon}\right) e^{-\delta^2 \gamma m/4}$.

Proof. This follows directly by applying Chernoff bounds (both the upper and lower tail form) to the event that some S_ℓ violates the condition, then applying the union bound over all choices of ℓ. □

Claim 6.4. *When the exploration phase is δ-approximate, we have* $(1-7\delta) S^{-1}\left(\frac{k}{n}\right) \leq \widetilde{p} \leq p^*$.

Proof. It is easy to see that none of the stopping conditions of the exploration phase can be triggered until the price goes below p^*. Therefore $\widetilde{p} \leq p^*$. For the other inequality observe that, by Claim D.3 it holds that

$$S^{-1}(\alpha) \geq (1-\delta)\, S^{-1}(\tfrac{k}{n}).$$

Therefore it suffices to show that $\widetilde{p} \geq (1-6\delta) S^{-1}(\alpha)$.

Assume for a contradiction that the stopping condition is not triggered in some phase ℓ such that

$$p_{\ell+1} < (1+\delta)^{-6}\, S^{-1}(\alpha).$$

Therefore, at round ℓ we have

$$p_\ell = (1+\delta) p_{\ell+1} < (1+\delta)^{-5} S^{-1}(\alpha) \tag{20}$$

Examining the stopping conditions, and using our assumption that none where triggered at round ℓ, we deduce that:

$$S_\ell < (1+\delta)\alpha \tag{21}$$

$$R_{\max}/R_\ell < (1+\delta)^2, \tag{22}$$

Combining (20) and (21), we get

$$R_\ell = p_\ell S_\ell < (1+\delta)^{-4} \alpha S^{-1}(\alpha) \tag{23}$$

Note that, since we chose round ℓ such that $p_\ell \ll S^{-1}(\alpha)$, the mechanism already encountered some round $t < \ell$ such that p_t is"close" to $S^{-1}(\alpha)$ – in particular

$$(1+\delta)^{-1} S^{-1}(\alpha) \leq p_t \leq S^{-1}(\alpha) \tag{24}$$

and therefore also $S(p_t) \geq \alpha$. Since we assume the exploration phase is δ-approximate, the estimated survival rate at round t satisfies

$$S_t \geq (1+\delta)^{-1} S(p_t) \geq (1+\delta)^{-1} \alpha \tag{25}$$

Combining (25) and (24), we get that the estimated revenue R_t at round t satisfies

$$R_t = p_t S_t \geq (1+\delta)^{-2} \alpha S^{-1}(\alpha) \tag{26}$$

The value of $R_{\max}$ in round ℓ is at least R_t. Combining (26) with (23), this shows that at round ℓ we have $\frac{R_{\max}}{R_\ell} > (1+\delta)^2$, contradicting (22). □

Claim 6.5. *When the exploration phase is δ-approximate, we have $R(\widetilde{p}) \geq (1 - 7\delta)R(p^*)$.*

Proof. By Claim 6.4, we are done when $p^* = S^{-1}\left(\frac{k}{n}\right)$. Therefore, assume $p^* = r$, the myerson reserve price. It is easy to see that $R(p_{\ell+1}) \geq \frac{1}{1+\delta}\ R(p_\ell)$ for each ℓ. Let t be the first integer such that $p_t \leq p^* = r$. Note that $(1+\delta)^{-1}p^* \leq p_t \leq p^*$. Claim 6.4 says that $\widetilde{p} \leq p^* = r$, therefore $\widetilde{\ell} \geq t$ and by Claim D.2 $S(p_t) \geq S(r) \geq 1/e \geq \gamma$. It suffices to show that a stopping condition must be triggered before $R(p_\ell)$ gets too small.

Assume for a contradiction that the stopping condition is not triggered by phase $\ell \geq t$, for some ℓ such that $R(p_{\ell+1}) < (1 - 7\delta)R(p^*)$. Since R decreases slowly as described above, it follows that $t < \ell$. Moreover, since we assumed the exploration phase is δ-approximate, $S_t \geq \frac{1}{1+\delta}\ S(p_t) \geq \frac{1}{1+\delta}\ \gamma$. Therefore, during phase ℓ we have $R_{max} \geq R_t = S_t p_t \geq (\frac{1}{1+\delta})^2\ R(p^*)$. Since no stopping condition is triggered for phase ℓ, it must be that $R_\ell \geq (\frac{1}{1+\delta})^2\ R_{max} \geq (\frac{1}{1+\delta})^4\ R(p^*)$. Moreover $R(p_{\ell+1}) \geq \frac{1}{1+\delta}\ R(p_\ell) \geq (\frac{1}{1+\delta})^2\ R_\ell \geq (\frac{1}{1+\delta})^6 R(p^*)$, a contradiction. □

We can now complete the proof of Lemma 6.2.

We condition on the exploration phase being δ-approximate. Let n' and k' be the number of players and items left after the exploration phase, respectively. In the exploitation phase, we attain expected revenue $\mathtt{Rev}(A_{k'}^{n'}(\widetilde{p}))$. Moreover, in the exploration phase we attained revenue at least $(k - k')\widetilde{p}$, since we only used prices greater than or equal to $\widetilde{p}$. Therefore, the total revenue of our mechanism is at least $\mathtt{Rev}(\mathcal{A}_{k'}^{n'}(\widetilde{p})) + (k' - k)\widetilde{p}$. It is easy to see that this is at least $\mathtt{Rev}(\mathcal{A}_k^{n'}(\widetilde{p}))$.

It remains to bound the revenue of $\mathcal{A}_k^{n'}(\widetilde{p})$. Observe that $\frac{n'}{n} \geq 1-\delta$. For brevity, denote $\beta \triangleq (1 - \frac{1}{\sqrt{2\pi k}})$.

There are two cases. In the first case, $p^* = S^{-1}(k/n)$. Corollary C.3 implies that

$$\mathtt{Rev}(\mathcal{A}_k^{n'}(\widetilde{p})) \geq \beta\, \frac{n'}{n}\, \frac{\widetilde{p}}{p^*}\ \mathtt{Rev}(\mathcal{A}_n^n(p^*)).$$

Using Claim 6.4, this gives

$$\mathtt{Rev}(\mathcal{A}_k^{n'}(\widetilde{p})) \geq \beta\ (1 - 8\delta)\ \mathtt{Rev}(\mathcal{A}_n^n(p^*)).$$

The second case is $p^* = r$. By Claim 6.5 and unimodality of R, we have that

$$\mathtt{Rev}(\mathcal{A}_n^n(\max(S^{-1}(k/n), \widetilde{p}))) \geq \mathtt{Rev}(\mathcal{A}_n^n(\widetilde{p})) \geq (1 - 7\delta)\ \mathtt{Rev}(\mathcal{A}_n^n(p^*)).$$

Moreover, Corollary C.3 and Claim 6.4 show that

$$\mathtt{Rev}(\mathcal{A}_k^{n'}(\widetilde{p})) \geq \beta\ (1 - 8\delta)\ \mathtt{Rev}(\mathcal{A}_n^n(\max(S^{-1}(k/n), \widetilde{p}))).$$

Therefore we combine the above two equations to get

$$\mathtt{Rev}(\mathcal{A}_k^{n'}(\widetilde{p})) \geq \beta\ (1 - 15\delta)\ \mathtt{Rev}(\mathcal{A}_n^n(p^*).$$

Using Lemma B.1 shows that Mechanism 1 achieves, in expectation, at least the following fraction of the revenue of the offline optimal mechanism:

$$\beta\ (1 - O(\delta))\ \left(1 - 2\ \log_{1+\delta}(\tfrac{1}{\epsilon}) \exp(-\tfrac{1}{4}\, \delta^2 \gamma m)\right).$$

Now, plug in δ from (19), and m as defined in the mechanism. Note that $m = \Theta(\frac{\delta^2 n}{\log 1/\epsilon})$. We obtain the final bound replacing γ by the lesser quantity $\frac{k}{n}$, and using the fact that $\log_{1+\delta}(x) = \Theta(\frac{1}{\delta}\ \log x)$.

Acknowledgements. The authors are grateful to Robert Kleinberg and Assaf Zeevi for comments and suggestions.

References

[1] Rajeev Agrawal. The continuum-armed bandit problem. *SIAM J. Control and Optimization*, 33(6):1926–1951, 1995.

[2] Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2-3):235–256, 2002. Preliminary version in *15th ICML*, 1998.

[3] Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. The nonstochastic multiarmed bandit problem. *SIAM J. Comput.*, 32(1):48–77, 2002. Preliminary version in *36th IEEE FOCS*, 1995.

[4] Peter Auer, Ronald Ortner, and Csaba Szepesvári. Improved Rates for the Stochastic Continuum-Armed Bandit Problem. In *20th Conf. on Learning Theory (COLT)*, pages 454–468, 2007.

[5] Moshe Babaioff, Liad Blumrosen, Shaddin Dughmi, and Yaron Singer. Posting prices with unknown distributions. In *Symp. on Innovations in CS*, 2011.

[6] Moshe Babaioff, Nicole Immorlica, and Robert Kleinberg. Matroids, secretary problems, and online mechanisms. In *18th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 434–443, 2007.

[7] Moshe Babaioff, Robert Kleinberg, and Aleksandrs Slivkins. Truthful mechanisms with implicit payment computation. In *11th ACM Conf. on Electronic Commerce (EC)*, pages 43–52, 2010.

[8] Moshe Babaioff, Yogeshwer Sharma, and Aleksandrs Slivkins. Characterizing truthful multi-armed bandit mechanisms. In *10th ACM Conf. on Electronic Commerce (EC)*, pages 79–88, 2009.

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

[10] Dirk Bergemann and Juuso Välimäki. Bandit Problems. In Steven Durlauf and Larry Blume, editors, *The New Palgrave Dictionary of Economics, 2nd ed.* Macmillan Press, 2006.

[11] Omar Besbes and Assaf Zeevi. Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Oper. Res.*, 57:1407–1420, November 2009.

[12] Avrim Blum and Jason Hartline. Near-optimal online auctions. In *16th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, 2005.

[13] Avrim Blum, Vijay Kumar, Atri Rudra, and Felix Wu. Online learning in online auctions. In *14th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 202–204, 2003.

[14] Sébastien Bubeck, Rémi Munos, and Gilles Stoltz. Pure Exploration in Multi-Armed Bandit Problems. In *20th Intl. Conf. on Algorithmic Learning Theory (ALT)*, 2009.

[15] Sébastien Bubeck, Rémi Munos, Gilles Stoltz, and Csaba Szepesvari. Online Optimization in X-Armed Bandits. In *Advances in Neural Information Processing Systems (NIPS)*, pages 201–208, 2008.

[16] Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006.

[17] Tanmoy Chakraborty, Eyal Even-Dar, Sudipto Guha, Yishay Mansour, and S. Muthukrishnan. Approximation schemes for sequential posted pricing in multi-unit auctions. In *Workshop on Internet & Network Economics (WINE)*, pages 158–169, 2010.

[18] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *ACM Symp. on Theory of Computing (STOC)*, pages 311–320, 2010.

[19] Nikhil Devanur and Jason Hartline. Limited and online supply and the bayesian foundations of prior-free mechanism design. In *ACM Conf. on Electronic Commerce (EC)*, 2009.

[20] Nikhil Devanur and Sham M. Kakade. The price of truthfulness for pay-per-click auctions. In *10th ACM Conf. on Electronic Commerce (EC)*, pages 99–106, 2009.

[21] Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. In *ACM Conf. on Electronic Commerce (EC)*, pages 129–138, 2010.

[22] Ashish Goel, Sanjeev Khanna, and Brad Null. The Ratio Index for Budgeted Learning, with Applications. In *20th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 18–27, 2009.

[23] Mohammad T. Hajiaghayi, Robert Kleinberg, Mohammad Mahdian, and David C. Parkes. Adaptive limited-supply online auctions. In *Proc. ACM Conf. on Electronic Commerce*, pages 165–174, 2005.

[24] J.D. Hartline and T. Roughgarden. Optimal mechanism design and money burning. In *ACM Symp. on Theory of Computing (STOC)*, 2008.

[25] R. Kleinberg. A multiple-choice secretary problem with applications to online auctions. In *16th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 630–631, 2005.

[26] Robert Kleinberg. Nearly tight bounds for the continuum-armed bandit problem. In *18th Advances in Neural Information Processing Systems (NIPS)*, 2004.

[27] Robert Kleinberg, Aleksandrs Slivkins, and Eli Upfal. Multi-Armed Bandits in Metric Spaces. In *40th ACM Symp. on Theory of Computing (STOC)*, pages 681–690, 2008.

[28] Robert D. Kleinberg and Frank T. Leighton. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *IEEE Symp. on Foundations of Computer Science (FOCS)*, 2003.

[29] T.L. Lai and Herbert Robbins. Asymptotically efficient Adaptive Allocation Rules. *Advances in Applied Mathematics*, 6:4–22, 1985.

[30] Ron Lavi and Noam Nisan. Competitive analysis of incentive compatible on-line auctions. In *ACM Conference on Electronic Commerce*, pages 233–241, 2000.

[31] Colin McDiarmid. Concentration. In M. Habib. C. McDiarmid. J. Ramirez and B. Reed, editors, *Probabilistic Methods for Discrete Mathematics*, pages 195–248. Springer-Verlag, Berlin, 1998.

[32] R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

[33] Hamid Nazerzadeh, Amin Saberi, and Rakesh Vohra. Dynamic cost-per-action mechanisms and applications to online advertising. In *17th WWW*, 2008.

[34] Robert B. Wilson. Efficient and competitive rationing. *Econometrica*, 57:1–40, 1989.

[35] Qiqi Yan. Mechanism design via correlation gap. In *22nd ACM-SIAM Symp. on Discrete Algorithms (SODA)*, 2011.

A Tail bounds

We use the well-known Chernoff Bounds, in a formulation from (e.g.) Theorem 2.3 in [31].

Theorem A.1 (Chernoff Bounds)**.** *Consider n i.i.d. random variables $X_1 \dots X_n$ on $[0,1]$. Let X be their average, and let $\mu = \mathbb{E}[X]$. Then:*

(a) $\Pr[|X-\mu| > \delta\mu] < 2\, e^{-\mu n \delta^2/3}$ *for any $\delta \in (0,1)$.*

(b) $\Pr[X > a] < 2^{-an}$ *for any $a > 6\mu$.*

Further, we use a somewhat non-standard corollary which provides a sharper (i.e., smaller) confidence radius (as defined in (3)) when μ is small.

Lemma A.2. *Consider n i.i.d. random variables $X_1 \dots X_n$ on $[0,1]$. Let X be their average, and let $\mu = \mathbb{E}[X]$. Then for any $\alpha > 0$ the following holds:*

$$\Pr\left[\,|X-\mu| < r(\alpha, X) < 3\, r(\alpha,\mu)\,\right] > 1 - e^{-\Omega(\alpha)},$$

where $r(\alpha, x) = \frac{\alpha}{n} + \sqrt{\frac{\alpha x}{n}}$.

A proof of Lemma A.2 can be found in the full version of [27]; we provide it here for completeness.

Proof of Lemma A.2. First, suppose $\mu \geq \frac{\alpha}{6n}$. Apply Theorem A.1(a) with $\delta = \frac{1}{2}\sqrt{\frac{\alpha}{6\mu n}}$. Thus with probability at least $1 - e^{-\Omega(\alpha)}$ we have $|X - \mu| < \delta\mu \leq \mu/2$. Moreover, plugging in the value for δ,

$$|X-\mu| < \tfrac{1}{2}\sqrt{\tfrac{\alpha\mu}{n}} \leq \sqrt{\tfrac{\alpha X}{n}} \leq r(\alpha, X) < 1.5\, r(\alpha,\mu).$$

Now suppose $\mu < \frac{\alpha}{6n}$. Then using Theorem A.1(b) with $a = \frac{\alpha}{n}$, we obtain that with probability at least $1 - 2^{-\Omega(\alpha)}$ we have $X < \frac{\alpha}{n}$, and therefore

$$|X-\mu| < \tfrac{\alpha}{n} < r(\alpha, X) < (1+\sqrt{2})\,\tfrac{\alpha}{n} < 3\, r(\alpha,\mu). \quad \square$$

A.1 (Sharper) Azuma-Hoeffding inequality

We use a tail bound on the sum of n random variables $X_t \in \{0,1\}$ such that each variable X_t is a random coin toss with probability M_t that depends on the previous variables $X_1, \dots, X_{t-1}$. We are interested in bounding the deviation $|X - M|$, where $X = \sum_t X_t$ and $M = \sum_t M_t$. The well-known Azuma-Hoeffding inequality gives $|X - M| \leq O(\sqrt{n \log n})$ with high probability. However, we need a sharper bound: $|X - M| \leq O(\sqrt{M \log n})$. Moreover, we need an extension of such bound which considers deviation $|\sum_{t=1}^n \alpha_t (X_t - M_t)|$, where each multiplier $\alpha_t \in [0,1]$ is determined by $X_1, \dots, X_{t-1}$. The tail bound that we need is stated as follows:

Lemma A.3. *Let $X_1, \dots, X_n$ be 0-1 random variables. Let $M = \sum_{t=1}^n\ \mathbb{E}[X_t | X_1,\ \dots, X_{t-1}]$. For each t, let $\alpha_t \in [0,1]$ be the multiplier determined by $X_1, \dots, X_{t-1}$. Then for any $b \geq 1$ the event*

$$|\textstyle\sum_{t=1}^n\ \alpha_t (X_t - M_t)| \leq b(\sqrt{M\ \log n} + \log n).$$

holds with probability at least $1 - n^{-\Omega(b)}$.

We have not been able to find this exact formulation in the literature. Instead, we derive it as an easy corollary of a more general bound that can be found in [31].

Theorem A.4 (Theorem 3.15 in [31])**.** *Let $Z_1, \dots, Z_n$ be random variables which take values in $[-1,1]$. Let $Z = \sum_{t=1}^n\ Z_t$, $\mu = \mathbb{E}[Z]$. Let $V = \sum_{t=1}^n$ Var$(Z_t | Z_1,\ \dots, Z_{t-1})$. Then for any $a > 0, v > 0$*

$$Pr\left[(|Z-\mu| \geq a) \wedge (V \leq v)\right] \leq e^{-\Omega(\frac{a^2}{v+a})}.$$

Corollary A.5. *In the setting of Lemma A.3, let $Z_t = X_t - y_t$, where $y_t \in [0,1]$ is a function of $X_1, \dots, X_{t-1}$. Let $Z = \sum_{t=1}^n\ Z_t$. Let $M = \sum_{t=1}^n\ \mathbb{E}[X_t | X_1,\ \dots, X_{t-1}]$. Then for any $b \geq 1$ the event that*

$$|\textstyle\sum_{t=1}^n\ \alpha_t (Z_t - \mathbb{E}[Z_t])| \leq b(\sqrt{M\ \log n} + \log n)$$

holds with probability at least $1 - n^{-\Omega(b)}$.

Proof. Let $\mathcal{F}_t = \sigma(X_1, \ldots, X_t)$ be the σ-algebra generated by $X_1, \ldots, X_t$, and let $M_t = \mathbb{E}[X_t | X_1, \ldots, X_{t-1}]$. Then conditional on $\mathcal{F}_{t-1}$, Z_t is a random variable with expectation $M_t - y_t$ and two possible values, $-\alpha_t\, y_t$ and $\alpha_t\,(1 - y_t)$, where α_t and y_t are constants. It follows that $\mathrm{Var}(Z_t|\mathcal{F}_{t-1}) = \alpha_t^2(M_t - M_t^2) \leq M_t$, and therefore $V \triangleq \sum_{t=1}^n \mathrm{Var}(Z_t|\mathcal{F}_{t-1}) \leq M$.

Taking Theorem A.4 with $a = b(\sqrt{v\,\log n} + \log n)$ it follows that for any$b \geq 1$ the event

$$(|Z - E[Z]| \geq b(\sqrt{v\,\log n} + \log n)) \wedge (V \leq v).$$

holds with probability at most $n^{-\Omega(b)}$. Finally, we take the Union Bound over (say) all integer v between $\log n$ and n, noting that $V \leq M$. □

Proof of Lemma A.3. In Corollary A.5, set $y_t = M_t$ and note that $Z_t = X_t - M_t$ and so $\mathbb{E}[Z_t] = 0$. □

We note in passing that Corollary A.5 implies other useful tail bounds. For instance, by setting $\alpha_t \equiv 1$ and $y_t \equiv 0$ we obtain a tail bound for $|X - E[X]|$, where $X = \sum_{t=1}^n X_t$.

B A self-contained proof of Lemma 3.1

For the sake of convenience, let us restate Lemma 3.1:

Lemma (Yan [35])**.** *Assume that the demand distribution is regular. Then there exists a fixed-price mechanism whose expected revenue is at least the optimal offline revenue minus $O(\sqrt{k \log k})$.*

Recall that $A_k^n(p)$ denotes the fixed price mechanism with k items, n agents, and fixed price p. Let M_k^n denote the optimal offline auction with n-players and k-items. Use p^* and r as in Lemma 3.3, namely $p^* = \max(r,\, S^{-1}(\frac{k}{n}))$, where $r = \mathrm{argmax}_p\, p\, S(p)$ is the Myerson reserve price.

Claim B.1. *If the demand distribution is regular then*

$$\mathtt{Rev}(A_n^n(p^*)) \geq \mathtt{Rev}(M_k^n).$$

Proof. Let q_i be the probability that M_k^n sells to agent i. By symmetry, $q_i = q_j$ for all players i and j, so we simply denote this probability by q. Let $p = S^{-1}(q)$ be the single price we would need to offer a agent in order to sell to him with probability q. Since R is a concave function of the selling probability, Jensen's inequality implies that $R(p)$ is an upper bound on the revenue collected by the Myerson auction from a single agent. Equivalently: $nR(p) \geq \mathtt{Rev}(M_k^n)$.

Now, observe that the expected number of items sold by M_k^n is nq. Since M_k^n never sells more than k items, it must be that $q \leq \frac{k}{n}$. Therefore, $p \geq S^{-1}(\frac{k}{n})$. By definition of p^*, we deduce that there are two cases: (1) $p^* = r$, or (2) $r \leq p^* = S^{-1}(\frac{k}{n}) \leq p$. In case (1) it is clear that $R(p^*) \geq R(p)$. In case (2) we get that $R(p^*) \geq R(p)$ since $R(x)$ is decreasing for $x \geq r$. Then

$$\mathtt{Rev}(A_n^n(p^*)) = nR(p^*) \geq nR(p) \geq \mathtt{Rev}(M_k^n). \;\square$$

Lemma 3.1 follows from Claim B.1 and Claim 3.2 because for $p = p^*$ we have $S(p) \leq \frac{k}{n}$, and so in the right-hand side in (1) we have

$$p\,\min(k, n\,S(p))) = np^*\,S(p^*) = \mathtt{Rev}(A_n^n(p^*)) \geq \mathtt{Rev}(M_k^n).$$

C More on the offline benchmark

In this section, we strenghten the following result which is immediate from Yan [35].

Lemma C.1 ([35])**.** *Assume that the demand distribution is regular. Then there exists a fixed-price mechanism whose expected revenue approximates the optimal offline revenue up to a factor* $1 - \frac{1}{\sqrt{2\pi k}}$.

In particular, we prove Claim C.2, which — combined with Claim B.1 — implies Lemma C.1. We use the strengthened result in Section 6 (via Corollary C.3).

Claim C.2. *For any regular demand distribution and any* $p \geq S^{-1}(k/n)$ *it holds that*

$$\texttt{Rev}(\mathcal{A}_k^n(p)) \geq \left(1 - \tfrac{1}{\sqrt{2\pi k}}\right) \texttt{Rev}(\mathcal{A}_n^n(p))$$

Proof. As a thought experiment, consider an environment where agent valuations are *correlated* as follows: The joint distribution of agent valuations can be sampled by choosing a set S' of k players uniformly at random, then for each agent in S' sampling from the conditional distribution $F(x)|_{x \geq S^{-1}(k/n)}$, and for each agent not in S' sampling from the conditional distribution $F(x)|_{x < S^{-1}(k/n)}$. Observe that each agent's valuation is distributed according to F, yet at any point exactly k players have value exceeding $S^{-1}(k/n)$.

Let T' be the set of players in this correlated environment whose valuation exceeds p. The probability of a particular agent being included in T' is $S(p)$, and $\mathbb{E}[|T'|] = nS(p)$. Since $p \geq S^{-1}(k/n)$, it is clear that $T' \subseteq S'$ and therefore $0 \leq |T'| \leq k$.

Now consider our original environment where each agent's valuation is drawn i.i.d from F. Let T be the set of players in this environment whose valuations exceed p. The probability of a agent being included in T is $S(p)$ – the same as the probability of being included in T'. However, each agent is included in T *independently* with probability $S(p)$. As a result, some of the players in T do not win an item – this happens when $|T| > k$. We can write the revenue of $\mathcal{A}_k^n(p)$ in this i.i.d environment as follows.

$$\texttt{Rev}(\mathcal{A}_k^n(p))] = p\mathbb{E}[\min(|T|, k) \tag{27}$$

Now, observe that $r(Y) = \min(|Y|, k)$ is the rank function of the k-uniform matroid. Moreover, it was shown in [35] that the correlation gap of this function is $\beta \triangleq \left(1 - \frac{1}{\sqrt{2\pi k}}\right)$. Therefore, since each agent is included in T independently, we know by the definition of the correlation gap and the fact that T and T' have the same marginals that

$$\mathbb{E}[r(T)] \geq \beta\ \mathbb{E}[r(T')]. \tag{28}$$

Recall that T' is always bounded between 0 and k, therefore $r(T') = |T'|$. Combining (27) and (28), we get

$$\begin{aligned}
\texttt{Rev}(A_k^n(p)) &= p\ \mathbb{E}[\min(|T|, k)] \\
&\geq \beta\ p\mathbb{E}[|T'|] \\
&= \beta\ pnS(p)) \\
&= \beta\ \texttt{Rev}(\mathcal{A}_n^n(p)))
\end{aligned}$$

□

We use Claim C.2 in Section 6 via the following generalization:

Corollary C.3. *Assume the demand distribution is regular. Let* p, p' *be two prices with* $p' \leq p$ *and* $p \geq S^{-1}(k/n)$*. Let* $n' \leq n$*. Then*

$$\texttt{Rev}(\mathcal{A}_k^{n'}(p')) \geq \frac{n'}{n}\frac{p'}{p}\left(1 - \tfrac{1}{\sqrt{2\pi k}}\right) \texttt{Rev}(\mathcal{A}_n^n(p)).$$

Proof. Observe that $\mathcal{A}_k^n(p')$ sells at least as many items as $\mathcal{A}_k^n(p)$ for every realization of the bids, but at price p' instead of p. Therefore $\texttt{Rev}(A_k^n(p')) \geq \frac{p'}{p}\texttt{Rev}(A_k^n(p))$. Combining with Claim C.2 we get that

$$\texttt{Rev}(A_k^n(p')) \geq \frac{p'}{p}\left(1 - \tfrac{1}{\sqrt{2\pi k}}\right)\texttt{Rev}(\mathcal{A}_n^n(p)).$$

Next, a simple (omitted) argument shows that the revenue $\texttt{Rev}(A_k^n(p))$ of a fixed price auction exhibits diminishing marginal returns in the number n of players. Therefore, $\texttt{Rev}(A_k^{n'}(p)) \geq \frac{n'}{n}\texttt{Rev}(A_k^n(p))$, completing the proof. $\square$

D Properties of MHR distributions

Let us state and prove several properties of Monotone Hazard Rate (MHR) distributions. We use these properties in Section 6.

Throughout, for a distribution F we use $F(x)$ to denote the c.d.f, $S(x) = 1-F(x)$ to denote the survival rate, and $f(x)$ to denote the p.d.f. We begin with a simple known characterization of MHR distributions.

Fact D.1. *A distribution is MHR if and only if $S(\cdot)$ is log-concave (i.e. $\log S(x)$ is a concave function of x).*

Next, we bound the survival probability at the Myerson reserve price.

Claim D.2. *Let F be an MHR distribution with support on $[0,\infty]$, and let $S(x) = 1 - F(x)$. Let $r \in \operatorname{argmax} R(\cdot)$ where $R(x) = x\,S(x)$. Then $S(r) \geq 1/e$.*

Proof. We have $R'(r) = S(r) + rS'(r) = 0$. Moreover, by Fact D.1 we deduce that

$$\frac{\log S(r)}{r} \geq \frac{d}{dx}\log(S(x))|_r = \frac{S'(r)}{S(r)}$$

Combining with the previous equality, we have $\frac{-1}{r} \leq \frac{\log(S(r))}{r}$ which is equivalent to $S(r) \geq \frac{1}{e}$. $\square$

We now use log-concavity to bound the sensitivity of the inverse of the survival function.

Claim D.3. *Let F be a monotone hazard rate distribution with support on $[0,\infty]$, and let $\alpha, \beta \in [0,1]$ with $\beta \geq \alpha$. We have*

$$S^{-1}(\beta) \geq \frac{\log(\beta)}{\log(\alpha)} S^{-1}(\alpha)$$

Proof. By Fact D.1, $\log(S(x))$ is a concave, decreasing function of x with $\log(S(0)) = 0$ and $\lim_{x\to\infty}\log(S(x)) = -\infty$. Let $a = S^{-1}(\alpha)$ and $b = S^{-1}(\beta)$. By Jensen's inequality, this gives for every $a, b \in [0,\infty]$ with $b \leq a$

$$\frac{\log(S(b))}{\log(S(a))} \leq \frac{b}{a}$$

Let $a = S^{-1}(\alpha)$ and $b = S^{-1}(\beta)$. Plugging into the above inequality completes the proof. $\square$