Posting Prices with Unknown Distributions

Moshe Babaioff[1] Liad Blumrosen[2] Shaddin Dughmi[3] Yaron Singer[4]
[1]Microsoft Research, Mountain View, CA 94043
[2]Department of Economics, The Hebrew University, Jerusalem, Israel 91905
[3]Department of Computer Science, Stanford University, Stanford, CA 94305
[4]Computer Science Division, UC Berkeley, Berkeley, CA
moshe@microsoft.com blumrosen@huji.ac.il shaddin@cs.stanford.edu yaron@cs.berkeley.edu

Abstract: We consider a dynamic auction model, where bidders sequentially arrive to the market. The values of the bidders for the item for sale are independently drawn from a distribution, but this distribution is *unknown* to the seller. The seller offers a take-it-or-leave-it price for each arriving bidder (possibly different for different bidders), and aims to maximize revenue. We study how well can such sequential posted-price mechanisms approximate the optimal revenue that is achieved when the distribution is known to the seller. On the negative side, we show that sequential posted-price mechanisms cannot guarantee a constant fraction of this revenue when the class of candidate distributions is unrestricted. We show that this impossibility holds even if the set of possible distributions is very small, or when the seller has a prior distribution over the candidate distributions. On the positive side, we devise a posted-price mechanism that guarantees a constant fraction of the known-distribution revenue when all candidate distributions exhibit the monotone hazard rate property.

Keywords: mechanism design, revenue maximization, posted price.

1 Introduction

Two main points of view are prevalent in the literature regarding the distributional knowledge of sellers in markets. Economists have traditionally considered *Bayesian* models, where the players have accurate distributional beliefs about the uncertain information. This assumption is problematic in practice, as collecting the distributional data may be constrained by technical or operational reasons, and also by the fact that this data is elicited from market participants that may act strategically also during this preliminary phase of the mechanism. Computer Scientists, on the other hand, have adapted the worst-case approach, traditionally employed in analysis of algorithms, to designing mechanisms that are *prior-free* (see e.g. [10]). Here, the preferences of the players are assumed to be arbitrary, and the analysis compares the performance of the mechanism on a worst case instance to a carefully crafted benchmark. In reality, however, worst case instances are rarely representative of the real-world performance of a mechanism. Moreover, with no clear notion of *optimal auction* in a prior-free setting, benchmarks are often controversial, and yield worst-case competitive ratios that are often disappointing even for the best auctions.

In this paper, we consider a framework, proposed in [11] and further developed in [8], that bridges the worst case and Bayesian models, and enables competitive analysis of auctions in the sense traditionally employed in computer science. In this framework, we consider environments where the preferences of the customers are drawn from a distribution, but this distribution is unknown to the seller, and learning any information about this distribution is an integral part of the mechanism. In other words, the mechanism is *detail-free*, in the sense first proposed by Wilson [19]. The goal is to design such a detail-free mechanism that competes with the best Bayesian mechanism that knows the distribution. Mechanisms in this model preserve detail-freeness, yet compete against mechanisms with access to distributional knowledge. In this paper, we consider such mechanisms in an *online* setting with the following additional restriction: the mechanisms must be *posted price*. In contrast to traditional *direct-revelation* mechanisms, posted price mechanisms interact with a player by offering him a single take-it-or-leave it offer, and never learn a player's value directly.

We consider a dynamic single-item auction model. A seller is trying to sell a single good to a set of n bidders. The bidders arrive sequentially to the market in an order they cannot influence, and the seller interacts with each bidder before observing future bidders. The auction terminates once the item is sold to one of the bidders, but in case the bidder does not buy the item, she leaves the market and never returns. Each bidder i ($1 \le i \le n$) has a private value v_i for the item. All the values are independently drawn from the same distribution F. The distribution F is unknown to the seller, but the seller knows that F belongs to a known family of distributions $\mathcal{F}$, each with support $[1, h]$ for some $h \ge 1$ known to the seller. The seller aims to maximize revenue.

We are interested in designing truthful auctions where revealing their true value is a dominant strategy for the bidders. The truthfulness constraint implies that each bidder must face a price that is independent of his bid, and where a bidder wins if and only if his value is at least the offered price. Instead of revealing his exact value, the bidder can also reply by just accepting or rejecting the offered price. For the seller, however, it does matter whether the bidder reports his exact value or just an "accept"/"reject" message, as the information about the true value can reveal information on the unknown underlying distribution and thus be helpful for the seller when deciding on future prices. Yet, asking agents to report their values might be unrealistic in some settings. While an agent never gains anything from revealing his valuation (compared to just reporting if it above the offered price), he has no real incentive to reveal it as well. This especially holds for mechanisms with a sampling phase where bidders are asked to reveal their exact information although they have no chance of winning. Moreover, bidders may prefer revealing minimal information on their values if they plan to participate in similar markets in the future; In addition, figuring out the exact value of a bidder may require some efforts on her behalf, while answering a take-it-or-leave it offer is usually much easier. Consider, for instance, an online travel agency (e.g. *Expedia.com*) trying to sell an airline ticket to a sequence of bidders; asking the bidders to report their willingness to pay is unnatural in such environments, but it is customary to offer an arriving customer a price and only observe if this price is accepted or not. We therefore aim to design mechanisms that elicit information from the bidders while giving them real incentives to participate, and where bidders will not voluntarily disclose more than the necessary information about their private types.

This paper therefore considers a popular and natural family of mechanisms, which is the family of *posted-price* mechanisms. In such mechanisms, the seller offers a take-it-or-leave-it price to each bidder in his turn, and the bidder either accepts or rejects the offer. If he accepts the offer then he wins the item and the auction ends; otherwise, he leaves the market for good and the seller waits for the next bidder to come. In posted-price mechanisms, bidders have a dominant strategy to accept any offer which is below their values, and reject it otherwise. Note that the bidders are not expected to reveal their exact valuation, but only to send a "reject"/"accept" message. The seller has an opportunity to learn information on the underlying distribution; for example, if the first 10 bidders rejected a set of high prices, then some distributions in $\mathcal{F}$ may be more likely to be the actual distribution than others. However, since the auction is terminated with the first "accept", this learning ability is clearly limited. Moreover, the algorithm gains nothing from being adaptive as it can take into account that if it is posting a price to a bidder it means the item was not sold yet. Another reason that using posted prices is appealing in this model is that when the true distribution F is *known* to the bidder, the optimal dynamic auction is a posted-price mechanism (see, e.g., [5]) and can be computed by a dynamic programming algorithm.

In this paper we would like to measure how much revenue can be obtained using sequential posted-price mechanisms when bidders' valuations are identically distributed yet the prior distribution on the preferences of the bidders is unknown to the mechanism (but is known to exist).[1]

We would like to compare this revenue to the optimal expected revenue achievable in a dynamic mechanism when the distribution is known to the seller, and we denote this revenue by $R^{on}(F)$. One may also wish to compare this revenue to the optimal "offline" revenue, the revenue that is obtained when the distribution is known to the seller and all the bidders are simultaneously present in the market. For standard (Myerson-regular) distributions, this revenue is achieved by the Myerson auction [16] that is essen-

[1]While the assumption of identical distributions is strong, we note that if we assumed arbitrary non-identical distributions for the bidders that would yield very negative results. In this case the assumption that the distributions are not known is at least as strong as assuming adversarial input (each agent sampled from its own point distribution). Clearly, deterministic mechanisms cannot achieve any reasonable approximation (better than h) for such inputs. Moreover, we show via a simple proof in Claim 14 that randomized mechanisms cannot achieve a factor better than $\Omega(\log h / \log\log h)$. As both lower bounds are based on point distributions, which trivially have monotone hazard rate (MHR), we observe that adding the MHR assumption with unrestricted *non-identical* distributions does not make a reasonable upper bound possible. Given these negative results, in this paper we add the natural assumption that all agents distributions are *identical*.

tially a second-price (Vickrey) auction with an optimally chosen reserve price. When the distribution is known to the seller, it is known that sequential posted price mechanisms achieve in large markets at least 78% of the Myerson revenue ([4]) and at least half of the Myerson revenue when multiple items are for sale ([7]). When considering the above positive results regarding the power of posted-price mechanisms, one might hope that posted prices can work reasonably well even with unknown distributions. Nonetheless, the first main result in this paper is negative and shows that posted-price mechanisms can only obtain a diminishing fraction of the revenue that could have been achieved had the distribution been known to the seller. We first prove a hardness result for deterministic mechanisms when the class $\mathcal{F}$ of candidate distributions is unrestricted:

Theorem: *When $\mathcal{F}$ contains all possible distributions on the support $[1, h]$, no deterministic sequential posted-price mechanism obtains better than a $\Omega(\frac{\log h}{\log \log h})$-approximation to the revenue obtained with a known distribution ($R^{on}(F)$).*

We note that this impossibility result is nearly tight, as there is a simple, deterministic posted-price mechanism that achieves an $O(\log h)$ approximation to this revenue benchmark. We also note that our results hold for any number of agents n, and are only asymptotic in h.

The above theorem is proved by constructing a hard instance of $\frac{\log h}{\log \log h}$ distributions, where every sequence of prices will achieve poor revenue for at least one of the distributions. We can therefore strengthen the above theorem and claim that better than an $\Omega(\frac{\log h}{\log \log h})$-approximation is impossible using sequential posted prices, even when the set of candidate distributions is very small, that is, $|\mathcal{F}| \cong \frac{\log h}{\log \log h}$.

The above negative result has a worst-case nature, in the sense that we require that the mechanisms will perform well for any distribution in $\mathcal{F}$. A Bayesian approach to this problem would consider a case where the seller is still unaware to the actual distribution from $\mathcal{F}$ but has some prior distribution g over the set $\mathcal{F}$.[2] One would hope that this assumption would allow us to achieve better revenue guarantees. Unfortunately, we show that this case remains hard.

Theorem: *Consider a class of distributions $\mathcal{F}$, where the actual distribution is drawn from this class according to a known distribution g. Then, there exist $\mathcal{F}$ and g such that all deterministic mechanisms achieve at most an $\Omega(\frac{\log h}{\log \log h})$-approximation to the revenue obtained with a known distribution ($R^{on}(F)$).*

Using Yao's min-max principle, we conclude from the above theorem that our first theorem actually holds also for randomized mechanisms. We conclude that even randomized mechanisms cannot guarantee a better approximation than $\Omega(\frac{\log h}{\log \log h})$.

As the above hardness results show, one must restrict the set of possible distributions for obtaining positive results in our model. In our final main result, we construct a mechanism that guarantees a constant fraction of the known-distribution revenue when all candidate distributions have monotone hazard rate (that is, $\frac{f(x)}{1-F(x)}$ is non-decreasing).[3] For this approximation result, we require that n will large enough with respect to $\log h$, and their ratio will affect the approximation we obtain. For instance, the theorem show that when $\sqrt{n} > \log h$, our mechanism achieves at least $\frac{1}{4e}$ of the revenue achieved when the distribution is known. In general,

Theorem: *Assume that $n^{\epsilon} > \log h$ for some constant $1 > \epsilon > 0$. Then, there exists a deterministic mechanism that achieves a $\frac{1-\epsilon}{2e}$-fraction of the revenue obtained with a known distribution, $R^{on}(F)$, when $\mathcal{F}$ contains all distributions with monotone hazard rate.*

Our proposed mechanism is simple, while its analysis is quite evolved. We define $\log h$ price levels, $h/2, h/4, h/8, ..., h/2^i, ..., 2, 1$ and offer each one of them to $\frac{n}{\log h}$ bidders (from highest price level to lowest). The requirement that n is large enough with respect to $\log h$ is necessary, as we show that no deterministic mechanism can achieve an approximation ratio better than $h^{1/n}$ even for point distributions (which are trivially monotone hazard rate), thus constant approximation is impossible with n small relative to $\log h$.

[2] Note that a distribution over a class of distributions of types is *not* equivalent to just having a distribution over types, due to the repeated sampling. To see this consider a distribution over the family of 2 point distributions, one with point mass at 1 and the other with a point at 2. After setting the price to 2 for the first agent, the exact value of the second agent is complectly known.

[3] The non-decreasing hazard rate condition is standard in mechanism design (see, for example, [14] and in recent computer-science work of [6, 11]). It is satisfied by many natural distributions, including the exponential, uniform, and binomial distributions.

Related Work: We now survey some related papers on detail-free mechanism design, posted prices, dynamic mechanisms and secretary problems. First, we note that the compromise between worst-case and Bayesian competitive analysis was first proposed by Hartline and Roughgarden [11]. Dhangwatnotai et al [8] then applied this framework to design constant-approximate direct-revelation mechanisms in fairly general environments.

One closely related paper is by Gershkov and Moldovanu [9], who studied dynamic auctions settings where the distribution of the buyers' preferences is unknown to the seller; They characterized necessary and sufficient conditions for information-theoretic optimum to be implementable in equilibrium, and basically they showed that the first-best allocation should consist of threshold values (that correspond to posted prices in our model) to be implementable. [9] did not study the magnitude of inefficiency in this setting, and in this sense our work complements their work. Segal [18] and a sequence of papers in the CS literature (see survey in [12] studied prior-free environments where empirical distributions where used to obtain revenue guarantees. Blumrosen and Holenstein [4] studied posted-price mechanisms, both in static and dynamic environments, with commonly known distributions, computed their exact revenue and compared it to the optimal (Myerson) revenue. A recent paper by Chawla et al. [7] studies sequential posted pricing in more general models, of matroid-based allocation rules and in some multi-dimensional settings, and presented several constant approximations to the Myerson revenue. Kleinberg and Leighton [13] presented upper and lower bounds on the additive regret of posted-price auctions for unlimited supply of goods.

Several recent papers ([1, 2]) studied versions of the secretary problem, where an adversary fixes a set of n values, these values arrive in a random order, and stopping rules should be designed to approximate the full-information solution. We note that with i.i.d. samples any order of values is equally likely, thus the secretary model is weaker than our model[4] in the sense that any positive result for the secretary problem can be applied to our model, and any hardness result to the unknown i.i.d. distribution model holds for the secretary model. In the context of secretary problems, our paper studies stopping rules that are based on a threshold-based decisions at each stage, without observing the exact value of each secretary. Finally, online auctions were first studied by [15], and are surveyed in [17].

We proceed as follows. Section 2 briefly presents some definitions and notations. Section 3 describes our main impossibility result for deterministic mechanisms, and in Section 4 we extend this result to randomized mechanisms. We conclude in Section 5 by presenting a positive result when players have monotone hazard rate distributions.

2 Preliminaries

We consider a model where a seller has one item for sale. A set of n bidders arrive sequentially, and we index them by the order of their arrival (Bidder 1 arrives first, then Bidder 2, etc.). Each agent i has a private value v_i for the item. There is a publicly known $h \geq 1$ such that for every $1 \leq i \leq n$ it holds that $v_i \in [1, h]$. The n values $(v_1, ..., v_n)$ are sampled i.i.d. from a distribution $F \in \mathcal{F}$. Given a distribution F we establish the following notations.

- $W(F)$ denotes the expected social welfare, that is $E_{\mathbf{v}\sim F}[max_{i=1}^{n}\{v_i\}]$.
- Given a list of posted prices $\mathbf{p} = (p_1, ..., p_n)$, let $Rev(\mathbf{p}, F)$ be the expected revenue obtained in a posted-price mechanism that offers a price p_i to the i'th arriving bidder with value v_i, sampled from F.
- For a mechanism M that offers the prices $\mathbf{p}$ we denote its expected revenue by $R^M(F) = Rev(\mathbf{p}, F)$.
- Let $R^{on}(F)$ be the *optimal expected revenue* in a dynamic auction when the distribution F is *known* to the seller. Since it is known (e.g., [5]) that the optimal dynamic auction can be implemented by a posted-price mechanism, we have

$$R^{on}(F) = max_{\mathbf{p}\in\mathbb{R}^n} Rev(\mathbf{p}, F)$$

We observe that since the values are drawn from i.i.d. distributions, the revenue in any posted-price mechanism is dominated by another mechanism for which $p_1 \geq p_2 \geq ... \geq p_n$, that is, with decreasing posted prices.

For a benchmark B (e.g., welfare or revenue) we say that a posted price mechanism M is β-approximation if for every input it gets at least $1/\beta$ fraction of the benchmark B.

[4] Formally, this holds for continuous distributions where ties occur with probability zero.

3 A lower bound for deterministic mechanisms

In this section, we show that sequential posted-price mechanisms cannot obtain a good revenue approximation when the distribution on the bidders' preferences is unknown and unrestricted. We show that every posted price mechanisms can guarantee at most a fraction proportional to $\log\log h/\log h$ of the optimal revenue that is obtained by dynamic mechanisms with a known distribution ($R^{on}(F)$). In the next section, we extend this result to randomized mechanisms. We now present our lower bound.

Theorem 1. *When $\mathcal{F}$ contains all the distributions over $[1, h]$, every deterministic posted-price mechanism obtains a revenue approximation of no better than $\Omega(\frac{\log h}{\log\log h})$ for some $F \in \mathcal{F}$.*

In the rest of this section we first prove some simple bounds and then prove Theorem 1.

3.1 Few simple bounds

Before proving the above lower bound, we present some simple bounds on the revenue in our model (some bounds are given on the social welfare and are hence stronger). We first observe (in Proposition 2) that the impossibility result in Theorem 1 is almost tight as an upper bound of $2\log h$ can be trivially achieved with a randomized mechanism, and an upper bound of $4\log h$ can be achieved with a deterministic mechanism (as long as $n \geq \log h$). The condition that $n \geq \log h$ is essential in order to achieve such a deterministic $O(\log h)$ upper bound, as we show in Proposition 3 that no deterministic mechanism has revenue approximation better than $h^{1/n}$.

For a vector of realized values $\mathbf{v} = (v_1, v_2, ..., v_n)$ we define the *realized social welfare* to be $W(\mathbf{v}) = max_{i=1}^{n}\{v_i\}$ and the *realized revenue of mechanism M* by $R^M(\mathbf{v})$. With these notation we present the following proposition. Note that any revenue approximation to the social welfare obtains at least the same approximation factor to any revenue benchmark, since the revenue is bounded from above by the social welfare (in individually-rational mechanisms).

Proposition 2. *There exists a randomized posted-price mechanism M that achieves revenue that is a $2\log h$-approximation to the realized social welfare, i.e., $\frac{W(\mathbf{v})}{R^M(\mathbf{v})} \leq 2\log h$ for every vector $\mathbf{v} \in [1, h]^n$.*

In addition, there exists a deterministic mechanism M that achieves a $4\log h$-approximation to the expected social welfare when $n \geq \log h$, i.e., $\frac{W(F)}{R^M(F)} \leq 4\log h$ for every distribution F.

Proof. A randomized mechanism:

Pick a random price $p = 2^i$ for $i \in \{0, 1, ..., \log h - 1\}$ and set $p_i = p$ for all i. If v_{max} is the maximal value then with probability $1/\log h$ the price $p \in [v_{max}/2, v_{max}]$ and the approximation follows.

A deterministic mechanism:

Set the price $p = h/2^i$ for $i \in \{0, 1, ..., \log h - 1\}$ (in decreasing order) for $\lfloor n/\log h \rfloor$ times each price. Note that $\lfloor n/\log h \rfloor \geq max\{1, n/(2\log h)\}$ as $n \geq \log h$. Thus, with probability at least $1/2\log h$ the maximal value sampled from the distribution faces a price that is at least half the value, and the approximation follows.[5] □

We now show an easy lower bound on the approximation achieved by sequentially posting n prices. The main idea is that when the number of bidders n is small, the offered prices are sparsely scattered on the support and therefore a bad approximation is unavoidable for some singleton distributions. This proposition will thus allow us not to handle cases where n is small when proving our main results later in this section.

Proposition 3. *When $\mathcal{F}$ contains all possible point distributions over $[1, h]$, no deterministic posted-price mechanism obtains better than a $h^{1/n}$-approximation to the optimal revenue achievable with a known distribution; that is, for any $\epsilon > 0$ there exists a distribution $F \in \mathcal{F}$ such that*

$$\frac{R^{on}(F)}{R^M(F)} > h^{1/n} - \epsilon \qquad (1)$$

Proof. Let $p_1 \geq p_2 \geq ... \geq p_n$ be the posted prices published by the mechanism. We first observe that we must have that $p_n = 1$, otherwise if the whole mass of the distribution is on $1 + \epsilon < p_n$ the approximation ratio will be unbounded. A second observation is that the ratio between some pair of consecutive prices must be at least $h^{\frac{1}{n}}$; Otherwise,

$$h = \frac{h}{p_1} \cdot \frac{p_1}{p_2} \cdot \frac{p_2}{p_3} ... \cdot \frac{p_{n-1}}{p_n} < (h^{\frac{1}{n}})^n = h \qquad (2)$$

[5] We note that the above randomized mechanism has an advantage from a strategic point of view, as bidders have no reason to act strategically with respect to their arrival time as the price never changes. The deterministic mechanism, on the other hand, does not admit this property as it offers a decreasing sequence of prices. Moreover, the randomized mechanism is independent of the number of players n.

Let p_{i-1}, p_i be prices such that $\frac{p_{i-1}}{p_i} \geq h^{\frac{1}{n}}$. If the whole mass of the distribution lies on $p_{i-1}-\epsilon$, then our posted-price mechanism obtains revenue of p_i where a seller who is knowledgable about the true distribution can gain $p_{i-1}-\epsilon$. Overall, the approximation obtained is at least $\frac{p_{i-1}-\epsilon}{p_i} \geq h^{\frac{1}{n}} - \epsilon$. □

3.2 Proof of Theorem 1

To prove the theorem we define the following "hard" family of distributions $\mathcal{F} = F_1, F_2, ...$ Let $\alpha = \frac{\log h}{\log \log h}$. The $j-th$ distribution $F_j \in \mathcal{F}$ satisfies: $Pr[x = 1] = 1 - j/n$ and $Pr[x = \alpha^i] = 1/n$ for $i \leq j$.

It is easy to see that since $\alpha = \frac{\log h}{\log \log h}$, the size of $\mathcal{F}$ must be around α as the following observation shows (proof appears in Appendix A.2).

Observation 4. *For the above set of distributions $\mathcal{F}$, and for large enough h, we have that $\alpha - 1 \leq |\mathcal{F}| \leq 2\alpha$.*

Before presenting the formal proof we sketch the outline of the proof. We first show that for each distribution $F_j \in \mathcal{F}$ it holds that with constant probability one of the n sampled values is α^j and thus the optimal online mechanism, that knows the true distribution, obtains an expected revenue of at least a constant times α^j (Lemma 5). On the other hand, we show that the revenue of any deterministic posted-price mechanism is much smaller (about a $1/\alpha$ fraction of this revenue). This is done in few steps.

First, in Lemma 6 we show that if a price in $[\alpha^{j-1}, \alpha^j]$ is used r_j times out of n then the revenue on F_j is bounded by some function of α and r_j. As there are about α distributions in $\mathcal{F}$, for one of these distributions r_j must be small, at most $n/(\alpha - 1)$. For this distribution F_j, the revenue is small, in the order of α^{j-1}. Together with the fact that the optimum online mechanism obtains revenue of about α^j, we conclude that for F_j we only get about $1/\alpha$ fraction of the optimum.

We next move to the formal proof of the theorem. We assume that $n > 4\alpha = 4\frac{\log h}{\log \log h}$; Otherwise, we can invoke Proposition (3) that claims that no deterministic posted-price mechanism can achieve a better approximation than $h^{\frac{1}{n}} = \Omega(\log h)$ (when $n > \frac{\log h}{\log \log h}$).

For a sequence $(v_1, v_2, ..., v_n)$ of n independent samples from F_j define Y to be the random variable of the number of values α^j in the sequence. We first show that if F_j is known, then high expected revenue can be achieved by online mechanisms, that is, we show that $R^{on}(F_j)$ is proportional to α^j.

Lemma 5. *For any distribution $F_j \in \mathcal{F}$ it holds that*

$$\alpha^j \geq R^{on}(F_j) \geq (1 - e^{-1})\alpha^j \qquad (3)$$

Proof. The maximal value that can be sampled from F_j is α^j, thus $\alpha^j \geq R^{on}(F_j)$. $R^{on}(F_j)$ is the optimal online mechanism when it is known that the distribution is F_j. This mechanism has revenue at least as high as the mechanism that fixes a constant price of α^j for all agents. Such a mechanism will get a revenue of α^j whenever at least one value of α^j was sampled by one of the n agents. This happens with probability $1 - Pr[Y = 0]$. Since $Pr[Y = 0] = (1 - 1/n)^n \leq e^{-1}$, this event happens with probability at least $1 - e^{-1}$. □

We want to show that no posted-price mechanism can approximate this revenue, thus we bound the revenue of any mechanism from below. We first bound the revenue obtained on F_j as a function of r_j, the number of times the mechanism posts a price in $[\alpha^{j-1}, \alpha^j]$. In the rest of the proof we will use the notation $k_0 = \lceil \alpha \rceil$. Proof can be found in the appendix.

Lemma 6. *Assume that $n > 4\alpha$, and consider a deterministic posted-price mechanism that posts a price in $[\alpha^{j-1}, \alpha^j]$ for r_j times. Assume that $r_j \leq \frac{n-k_0}{2}$. For distribution F_j it holds that*

$$R^M(F_j) \leq \alpha^j \cdot \left(\frac{2}{\alpha} + \frac{4e \cdot r_j}{n}\right)$$

Using the above machinery, we can now complete the proof of Theorem 1.

Proof. (Of Theorem 1)

By Observation 4 there are at least $\alpha - 1 > \alpha/2$ distributions in $\mathcal{F}$. This implies that for at least one j it holds that $r_j < \frac{2n}{\alpha}$.

For h large enough we have that $4/\alpha < 1/2$. As $r_j \leq \frac{2}{\alpha} \cdot n$, $n > 4\alpha$ and $k_0 \leq 2\alpha$ it holds that

$$2r_j + k_0 \leq \frac{4}{\alpha} \cdot n + 2\alpha \leq \frac{1}{2} \cdot n + \frac{n}{2} = n$$

This implies that $r_j \leq \frac{n-k_0}{2}$. We can thus use Lemma 6 for distribution F_j to show that as $\frac{r_j}{n} \leq \frac{2}{\alpha}$ we have

$$R^M(F_j) \leq \alpha^j \cdot \left(\frac{2}{\alpha} + \frac{4e \cdot r_j}{n}\right)$$
$$\leq \alpha^j \cdot \left(\frac{2}{\alpha} + \frac{4e \cdot 2}{\alpha}\right) \quad (4)$$
$$= \alpha^j \cdot \frac{8e+2}{\alpha} < 26\alpha^{j-1}$$

By Lemma 5, $R^{on}(F_j) \geq (1-e^{-1})\alpha^j$. Taken together with Eq. (4) we have:

$$R^M(F_j) \cdot \alpha \cdot \frac{1-e^{-1}}{26} < R^{on}(F_j)$$

which concludes the proof of the theorem. □

4 Priors over distributions and randomized mechanisms

In this section, we extend the impossibility result presented in Theorem 1 to randomized mechanisms. For that, we first prove the limitations of deterministic mechanisms in the case where there is a prior distribution over the candidate distributions. We then use Yao's min-max principle to conclude our result for randomized mechanisms.

Let $\mathcal{F}$ be a family of distributions, and let g be a prior over $\mathcal{F}$. Define $R^{on}(g)$ to be the expected revenue (over g) of the optimal online mechanism that knows which distribution $F \in \mathcal{F}$ was realized. Define $R^M(g)$ to be the expected revenue (over g) of the mechanism M that knows $\mathcal{F}$ but does not know which distribution $F \in \mathcal{F}$ was realized. We show that the best online posted price mechanism that does not know which distribution was realized has much smaller expected revenue.

Theorem 7. *Every deterministic posted-price mechanism obtains expected revenue approximation (over g) of no better than $\Omega(\log h/\log\log h)$ when $n > 4\alpha$, i.e., there exists a constant $c > 0$, $\mathcal{F}$, and a prior g over $\mathcal{F}$, such that for every mechanism M it holds that*

$$\frac{R^{on}(g)}{R^M(g)} > c \cdot \frac{\log h}{\log\log h}$$

We define g to be a distribution over the "hard family of distributions" (presented in the beginning of Section 3.2) that picks F_j with probability proportional to $1/\alpha^j$. Formally, let $w_j = 1/\alpha^j$ and let $\sigma = \sum_{j:F_j\in\mathcal{F}} w_j$. By g the distribution F_j is picked with probability $Pr[F = F_j] = w_j/\sigma$.

The theorem directly follows from the two lemmas below. The first lemma shows that if the distributions was known to the seller, then an expected revenue of roughly $\frac{\log h}{\log\log h}$ could be achieved. The second lemma shows that with an unknown distribution, no posted-price mechanism can gain more than a constant expected revenue.

Lemma 8. *Let g be the prior over $\mathcal{F}$ defined above. It holds that $R^{on}(g) = \frac{1}{\sigma} \cdot \Omega(\frac{\log h}{\log\log h})$*

Proof. By Lemma 5, for $F_j \in \mathcal{F}$ it holds that $R^{on}(F_j) \geq (1-e^{-1})\alpha^j$. Thus, each F_j contributed at least $(1-e^{-1})\alpha^j \cdot w_j/\sigma = (1-e^{-1})\sigma$ to the expectation, and as by Observation 4 the family $\mathcal{F}$ is of size $\Omega(\frac{\log h}{\log\log h})$, we conclude that $R^{on}(g) = \frac{1}{\sigma} \cdot \Omega\left(\frac{\log h}{\log\log h}\right)$. □

The proof of the following lemma can be found at Appendix A.3

Lemma 9. *Assume that $n > 4\alpha$. Let g be the prior over $\mathcal{F}$ defined above. For any deterministic posted-price mechanism it holds that $R^M(g) = \frac{1}{\sigma} \cdot O(1)$.*

Using Yao's min-max lemma we conclude that randomized mechanisms cannot achieve good approximation on an adversarially chosen distribution. We note that this lower bound is almost tight, as we showed (Proposition 2) a simple mechanism that obtains an $O(\log h)$-approximation. Therefore, the following corollary strengthen Theorem 1 for randomized mechanisms.

Corollary 10. *When $\mathcal{F}$ contains all the distributions over $[1, h]$, every randomized mechanism has revenue approximation of no better than $\Omega(\log h/\log\log h)$ when $n > 4\alpha$. I.e., there exists a constant $c > 0$ such that for any randomized mechanism M there exists $F \in \mathcal{F}$ such that $c \cdot \frac{\log h}{\log\log h} \cdot R^M(F) < R^{on}(F)$.*

5 An upper bound for monotone hazard rate distributions

5.1 The environment

We consider player valuations drawn from a distribution F with support $[1, h]$. We use $F(x)$ to denote the c.d.f, $f(x) = dF(x)/dx$ to denote the p.d.f, μ to denote the expectation, $S(x) = 1 - F(x)$ to denote the survival probability, and $H(x) = f(x)/S(x)$ to denote the hazard rate of F. In this section, we will show an mechanism that attains a constant approximation ratio for distributions F with $H(x)$ monotone non-decreasing. This Monotone Hazard Rate (M.H.R) assumption is common in auction theory, and M.H.R distributions include most natural distributions in this setting. We emphasize that the mechanism has no knowledge of distribution F, yet achieves the claimed approximation ratio uniformly over all F with a monotone nondecreasing hazard rate.

5.2 The mechanism

We define the *Equal-Sample-of-Every-Scale* mechanism: The mechanism offers the price $h/2^i$ to $\lfloor n/(\log h)\rfloor$ agents, for every $i \in \{1, ..., \log h\}$ in that order. Next, we will show that this mechanism attains a constant factor approximation to the optimum possible revenue under mild assumptions.

5.3 Bounding the performance of the mechanism

Our main positive result is that the *Equal-Sample-of-Every-Scale* Mechanism achieves a constant approximation for every monotone hazard rate distribution.

Theorem 11. *Let $\log h \le n^\epsilon$ for $\epsilon \in (0,1)$, and consider player valuations drawn i.i.d from a monotone hazard rate distribution F. Let X_n denote the first order statistic of n samples from F. The expected revenue of the* Equal-Sample-of-Every-Scale *Mechanism is at least*

$$\frac{1-\epsilon}{2e} E[X_n]$$

In other words, the expected revenue of the mechanism is a constant factor of the expected social welfare.

We recall that the mechanism is deterministic. It is not surprising that we need n to be relatively large ($n^\epsilon > \log h$), as the lower bound of Proposition 3 shows that without n being at least $\log h$ a constant approximation is unachievable by a deterministic mechanism[6].

First, we show that expectation of the first order statistic X_n of an M.H.R distribution F, as a function of the number of samples n, exhibits diminishing marginal returns in a strong sense.

[6]We note that the lowerbound of Proposition 3 used point-distributions, which obey the monotone hazard rate assumption.

Lemma 12.

$$\frac{E[X_{n+1}] - E[X_n]}{E[X_n] - E[X_{n-1}]} \le \frac{n}{n+1}$$

Proof. First, we can write $E[X_n]$ as follows

$$\begin{aligned} E[X_n] &= \int_{x=0}^{\infty} 1 - F^n(x)dx \\ &= \int_{F(x)=0}^{1} \frac{1-F^n(x)}{f(x)} dF(x) \\ &= \int_{F(x)=0}^{1} \frac{1-F(x)}{f(x)} \left(\sum_{i=0}^{n-1} F^i(x)\right) dF(x) \\ &= \int_{F(x)=0}^{1} \frac{1}{H(x)} \left(\sum_{i=0}^{n-1} F^i(x)\right) dF(x) \end{aligned}$$

Let $\Delta_n = E[X_{n+1}] - E[X_n]$. By the above expression for $E[X_n]$, Δ_n can be written as follows.

$$\Delta_n = \int_{F(x)=0}^{1} \frac{1}{H(x)} F^n(x) dF(x)$$

F is an M.H.R distribution, therefore $1/H(x)$ is a non-increasing function of x, and therefore also of $F(x)$. Applying Lemma 21 in the appendix with $z = F(x)$ and $g(z) = \frac{1}{H(F^{-1}(z))}$ gives us that $\Delta_n/\Delta_{n-1} \le \frac{n}{n+1}$, as needed. ☐

This allows us to bound the growth of the first order statistic in terms of the number of samples. Here, $\mathcal{H}_n = \sum_{i=1}^n 1/i$ denotes the nth harmonic number

Lemma 13. *For $m \le n$ it holds that*

$$\frac{E[X_m]}{E[X_n]} \ge \frac{\mathcal{H}_m}{\mathcal{H}_n} \ge \frac{\log m}{\log n}$$

Proof. We show the second inequality in Lemma 22 in the appendix. To show the first inequality, by induction it suffices to show that $E[X_{n+1}]/E[X_n] \le \mathcal{H}_{n+1}/\mathcal{H}_n$ for all integers $n \ge 1$. Letting $\Delta_0 = E[X_1]$ and $\Delta_i = E[X_{i+1}] - E[X_i]$ for $i > 0$, this is equivalent to showing that

$$\frac{\sum_{i=0}^{n} \Delta_i}{\sum_{i=0}^{n-1} \Delta_i} \le \frac{\mathcal{H}_{n+1}}{\mathcal{H}_n}$$

This, in turn, is equivalent to showing

$$\frac{\Delta_n}{\sum_{i=0}^{n-1} \Delta_i} \le \frac{\mathcal{H}_{n+1}}{\mathcal{H}_n} - 1 = \frac{1}{(n+1)\mathcal{H}_n}$$

We rewrite the above condition as follows

$$\sum_{i=0}^{n-1} \frac{\Delta_i}{\Delta_n} \ge \sum_{i=1}^{n} (n+1)/i$$

Therefore, it suffices to show that $\Delta_i/\Delta_n \geq (n+1)/(i+1)$. This can be established by inductive application of Lemma 12, completing the proof. □

The above lemma implies that $\Pr[X_m \geq \frac{\log m}{\log n} E[X_n]] \geq \Pr[X_m \geq E[X_m]]$. Lemma 20 in the appendix implies that X_m is distributed as a monotone hazard rate distribution. Moreover, a result of Barlow and Marshall [3] implies–as a special case–that every monotone hazard rate distribution exceeds its expectation with probability at least $1/e$. This gives the following inequality.

$$\Pr\left[X_m \geq \frac{\log m}{\log n} E[X_n]\right] \geq 1/e \tag{5}$$

Proof of Theorem 11. The mechanism samples at least $m = n^{1-\epsilon}$ bidders for each price $2^i \in [1, h]$. Let $p = 2^i \in [E[X_n](1-\epsilon)/2, E[X_n](1-\epsilon)]$. The revenue of the algorithm is at least that attained had we simply tried to sell to m players using price p, which is at least:

$$\begin{aligned} p\Pr[X_m \geq p] &\geq p\Pr[X_m \geq E[X_n](1-\epsilon)] \\ &= p\Pr[X_m \geq \frac{\log m}{\log n} E[X_n]] \\ &\geq p/e \geq E[X_n](1-\epsilon)/2e \end{aligned}$$

□

References

[1] M. Babaioff, M. Dinitz, A. Gupta, N. Immorlica, and K. Talwar. Secretary problems: weights and discounts. In *SODA '09: Proceedings of the Nineteenth Annual ACM -SIAM Symposium on Discrete Algorithms*, pages 1245-1254, Philadelphia, PA, USA, 2009. Society for Industrial and Applied Mathematics.

[2] M. Babaioff, N. Immorlica, D. Kempe, and R. Kleinberg. A knapsack secretary problem with applications. In *APPROX-RANDOM*, pages 16-28, 2007.

[3] R. E. Barlow and A. W. Marshall. Bounds for distributions with monotone hazard rate. *Ann. Math. Statist.*, 35(3): 1258-1274, 1964.

[4] L. Blumrosen and T. Holenstein. Posted prices vs. negotiations: an asymptotic analysis. In *ACM Conference on Electronic Commerce*, 2008.

[5] L. Blumrosen, N. Nisan, and I. Segal. Auctions with severely bounded communications. *Journal of Artificial Intelligence Research*, 28:233-266, 2006.

[6] S. Chawla, J. Hartline, and R. Kleinberg. Algorithmic pricing via virtual valuations. In *ACM conference on Electronic Commerce (EC)*, pages 243-251, 2007.

[7] S. Chawla, J. Hartline, D. Malec, and B. Sivan. Sequential posted pricing and multi-parameter mechanism design. In *STOC'10: Proceedings of the 42nd ACM symposium on Theory of computing*, pages 311-320, 2010.

[8] P. Dhangwatnotai, T. Roughgarden, and Q. Yan. Revenue Maximization with Single Sample. In *ACM Conference on Electronic Commerce (EC)*, pages 129-138, 2010.

[9] A. Gershkov and B. Moldovanu. Learning about the future and dynamic efficiency. *American Economic Review*, 99(4):1576-1587, 2009.

[10] A. V. Goldberg, J. D. Hartline, A. R. Karlin, M. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242-269, 2006.

[11] J. Hartline and T. Roughgarden. Optimal mechanism design and money burning. In *ACM Symposium on Theory Of Computing (STOC)*, pages 75-84, 2008.

[12] J. D. Hartline and A. R. Karlin. *In N.Nisan, T.Roughgarden, E.Tardos, and V.Vazirani (eds.) Chapter 13, Profit Maximization in Mechanism Design.* Cambridge University Press., 2007.

[13] R. D. Kleinberg and F. T. Leighton. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *FOCS*, pages 594-605, 2003.

[14] V. Krishna. *Auction theory.* Academic press, 2002.

[15] R. Lavi and N. Nisan. Competitive analysis of incentive compatible on-line auctions. In *ACM Conference on Electronic Commerce*, pages 233-241, 2000.

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

[17] D. Parkes. *In N. Nisan, T. Roughgarden, E. Tardos, and V.Vazirani (eds.) Chapter 16, Online Mechanisms.* Cambridge University Press., 2007.

[18] I. Segal. Optimal pricing mechanisms with unknown demand. *American Economic Review*, 93(3):509-529, 2003.

[19] R. Wilson. Efficient and competitive rationing. *Econometrica*, 57:1-40, 1989.

A Missing claims and proofs

A.1 Section 1

Claim 14. *Let $\alpha = \log h / \log\log h$. Assume player values are drawn from unknown, non-identical point distributions with support in $[1, h]$. No randomized posted price mechanism achieves better than a $\frac{2}{\alpha}$ fraction of the optimal revenue.*

Proof. A randomized posted-price mechanism chooses a (possibly random) price p_i to offer to player i, who then arrives with value v_i. Observe that the distribution of p_i is independent of v_i, though may depend on $\{v_j\}_{j<i}$. We observe that this is an adversarial setting, where an adversary may set v_i depending on the *distribution* of p_i.

We consider an adversary who tries to minimize the mechanism's revenue, in the following manner. For each player i, choose an integer k_i such that $1 \le \alpha^{k_i-1} \le \alpha^{k_i} \le h$, and $\Pr[p_i \in [\alpha^{k_i-1}, \alpha^{k_i}]]$ is minimized. By observation 15, this probability is upper-bounded by $1/2\alpha$. Let $v_i = \alpha^{k_i}$. The revenue collected by the mechanism from player i is upper-bounded by

$$\frac{1}{2\alpha} v_i + \frac{v_i}{\alpha} < \frac{2}{\alpha} v_i$$

Where the first term of the sum upper-bounds the revenue attained when $p_i \in [\alpha^{k_i-1}, \alpha^{k_i}]$, and the second term upper-bounds the revenue otherwise. Summing over all players, the total revenue of the mechanism is at most $\frac{2}{\alpha}\sum_i v_i$. Since the player valuations are drawn from point-distributions, the optimal revenue is $\sum_i v_i$, completing the proof. □

A.2 Section 3

Observation 15. *Let $\alpha = \frac{\log h}{\log\log h}$. It holds that $\alpha^\alpha < h$. Additionally, if h is large enough then $h < \alpha^{2\alpha}$.*

Proof. We first show that $\alpha^\alpha < h$, that is, $(\frac{\log h}{\log\log h})^{\frac{\log h}{\log\log h}} < h$.

The claim is true if and only if

$$\frac{\log h}{\log\log h} \cdot \log\left(\frac{\log h}{\log\log h}\right) < \log h$$

which holds if and only if

$$\frac{\log h}{\log\log h} \cdot (\log\log h - \log\log\log h) < \log h$$

which clearly holds.

Next we show that if h is large enough then $h < \alpha^{2\alpha}$, or equivalently, $h < (\frac{\log h}{\log\log h})^{\frac{2\ \log h}{\log\log h}}$. This claim is true if and only if

$$\log h < \frac{2\ \log h}{\log\log h} \cdot log\left(\frac{\log h}{\log\log h}\right)$$

which holds if and only if

$$\log h < \frac{2\ \log h}{\log\log h} \cdot (\log\log h - \log\log\log h)$$

which holds if and only if

$$2\log\log\log h \le \log\log h$$

which clearly holds if h is large enough. □

Corollary 16. *If h is large enough then the family $\mathcal{F}$ has at most 2α and at least $\lfloor\alpha\rfloor$ distributions. Note that $\lfloor\alpha\rfloor \ge \alpha - 1$.*

A.2.1 Proof of Lemma 6

In this section we prove Lemma 6.

Lemma 17. *Assume that $n > 4\alpha$, and consider a deterministic posted-price mechanism that posts a price in $[\alpha^{j-1}, \alpha^j]$ for r_j times. Assume that $r_j \le \frac{n-k_0}{2}$. For distribution F_j it holds that*

$$R^M(F_j) \le \alpha^j \cdot \left(\frac{2}{\alpha} + \frac{4e \cdot r_j}{n}\right)$$

Proof. We need to bound $R^M(F_j)$. If the price is not in $[\alpha^{j-1}, \alpha^j]$ the revenue of the mechanism is smaller than α^{j-1}. Let $R(\mathbf{v})$ be the revenue of the mechanism that posts the price α^j for r_j times and always posts the price of 0 afterwards, when the vector of values is $\mathbf{v}$.

$R^M(F_j) \le \alpha^{j-1} + E[R(\mathbf{v})]$, where $E[R(\mathbf{v})]$ is the expectation of $R(\mathbf{v})$. Recall that Y counts the number of α^j in $\mathbf{v}$.

$$E[R(\mathbf{v})] = \sum_{k=1}^{n} E[R(\mathbf{v})|Y=k] \cdot Pr[Y=k]$$

We next split the sum into two terms.

$$\begin{aligned} E[R(\mathbf{v})] = &\sum_{k=1}^{k_0} E[R(\mathbf{v})|Y=k] \cdot Pr[Y=k] \\ &+ \sum_{k=k_0+1}^{n} E[R(\mathbf{v})|Y=k] \cdot Pr[Y=k] \end{aligned} \tag{6}$$

We observe the following easy bound on $Pr[Y=k]$.

$$Pr[Y=k] = \binom{n}{k} n^{-k}\left(1-\frac{1}{n}\right)^{n-k} \leq \frac{n^k}{k!}\cdot n^{-k}\cdot 1 \leq \frac{1}{k!} \tag{7}$$

We can now bound the latter term of Eq. (6). Clearly, $E[R(v)|Y=k] \leq \alpha^j$, thus

$$\sum_{k=k_0+1}^{n} E[R(v)|Y=k]\cdot Pr[Y=k] \leq \alpha^j \sum_{k=k_0+1}^{n}\frac{1}{k!} \leq \alpha^j \sum_{k=k_0+1}^{n}\frac{1}{2^k} \leq \alpha^j\cdot 2^{-k_0} \leq \frac{\alpha^j}{k_0} \tag{8}$$

We next move to bound the first term of Eq. (6). The following claim would be useful.

Claim 18. *For distribution F_j it holds that*

$$E[R(v)|Y=k] \leq \alpha^j \cdot \left(1-\left(1-\frac{r_j}{n-k}\right)^k\right) \tag{9}$$

Proof. Let Z be the the event that in none of the r_j times that the mechanism posts the price α^j, the realized value is α^j.

$$Pr[Z] = \frac{\binom{n-r_j}{k}}{\binom{n}{k}} = \prod_{i=0}^{k-1}\left(\frac{n-r_j-i}{n-i}\right) = \prod_{i=0}^{k-1}\left(1-\frac{r_j}{n-i}\right) \geq \left(1-\frac{r_j}{n-k}\right)^k \tag{10}$$

Therefore,

$$E[R(v)|Y=k] = \alpha^j\cdot(1-Pr[Z]) \leq \alpha^j\cdot\left(1-\left(1-\frac{r_j}{n-k}\right)^k\right) \tag{11}$$

□

Recall that $r_j \leq \frac{n-k_0}{2}$. For $k \leq k_0$ this implies that $\frac{1}{2} \geq \frac{r_j}{n-k_0} \geq \frac{r_j}{n-k}$. We use the fact that for $x \in [0,1/2]$ it holds that $e^{-2x} \leq 1-x \leq e^{-x}$ to conclude that

$$-\left(1-\frac{r_j}{n-k}\right)^k \leq 1-e^{-\frac{2\cdot r_j\cdot k}{n-k}} \leq \frac{2\cdot r_j\cdot k}{n-k} \tag{12}$$

As we assume that $n > 4\alpha$ and as $k \leq k_0 \leq 2\alpha$ it holds that $n/2 > 2\alpha \geq k$, thus $n-k \geq n/2$. As $n-k > n/2$ it holds that $\frac{2\cdot r_j\cdot k}{n-k} \leq \frac{4\cdot r_j\cdot k}{n}$. Combining this with Claim 18 and Eq. (12) we derive that for $k \leq k_0$ it holds that

$$E[R(v)|Y=k] \leq \frac{4\cdot r_j\cdot k}{n}\alpha^j$$

We use this and Eq. (7) to bound the first term of Eq. (6).

$$\sum_{k=1}^{k_0} E[R(v)|Y=k]\cdot Pr[Y=k] \leq \sum_{k=1}^{k_0}\alpha^j\cdot\frac{4\cdot r_j\cdot k}{n}\cdot\frac{1}{k!} \leq \alpha^j\cdot\frac{4\cdot r_j}{n}\cdot\sum_{k=1}^{k_0}\frac{1}{(k-1)!} \leq \alpha^j\cdot\frac{4e\cdot r_j}{n} \tag{13}$$

Combining Equations (6), (8) and (13) we conclude that

$$E[R(v)] \leq \alpha^j\cdot\left(\frac{1}{k_0}+\frac{4e\cdot r_j}{n}\right)$$

As $R^M(F_j) \leq \alpha^{j-1}+E[R(v)]$ and $k_0 \geq \alpha$, it follows that

$$R^M(F_j) \leq \alpha^{j-1}+E[R(v)] \leq \alpha^j\cdot\left(\frac{1}{\alpha}+\frac{1}{k_0}+\frac{4e\cdot r_j}{n}\right) \leq \alpha^j\cdot\left(\frac{2}{\alpha}+\frac{4e\cdot r_j}{n}\right) \tag{14}$$

□

A.3 Section 4

Lemma 19. *Assume that $n > 4\alpha$. Let g be the prior over $\mathcal{F}$ as defined in Section 4. For any deterministic posted-price mechanism it holds that $R^M(g) = \frac{1}{\sigma} \cdot O(1)$.*

Proof. Let J be the set of indices j such that $r_j > \frac{n-k_0}{2}$. Observe that as $n > 4\alpha$ and $k_0 \leq 2\alpha$ for large enough h, thus $k_0 < n/2$, therefore $r_j > n/4$ for every $j \in J$. Since $\sum_{j\in J} r_j \leq n$, it holds that $|J| \leq 4$. For every $j \in J$ we have $R^M(F_j) \leq \alpha^j \cdot w_j/\sigma = 1/\sigma$, and thus,

$$\sum_{j:F_j\in J} R^M(F_j) \cdot Pr[F = F_j] \leq \frac{4}{\sigma}$$

For $j \notin J$ we invoke Lemma 6. Recall that $\sum_j r_j = n$ and that by Observation 4 the family $\mathcal{F}$ is of size at most 2α.

$$\begin{aligned}
&\sum_{j:F_j\in \mathcal{F}\setminus J} R^M(F_j) \cdot Pr[F = F_j] \\
&\leq \sum_{j:F_j\in \mathcal{F}\setminus J} \alpha^j \cdot \left(\frac{2}{\alpha} + \frac{4e\cdot r_j}{n}\right) \cdot \frac{w_j}{\sigma} \\
&\leq \frac{1}{\sigma} \cdot \sum_{j:F_j\in \mathcal{F}\setminus J} \cdot \left(\frac{2}{\alpha} + \frac{4e\cdot r_j}{n}\right) \qquad (15)\\
&\leq \frac{1}{\sigma} \cdot \left(4 + \frac{4e}{n} \sum_{j:F_j\in \mathcal{F}\setminus J} r_j\right) \\
&= \frac{4(e+1)}{\sigma}
\end{aligned}$$

For the case that $n > 4\alpha$, by combining the bound for j such that $F_j \in J$ and for the complimentary set we complete the proof of this lemma.

$$\begin{aligned}
R^M(g) &= \sum_{j:F_j\in \mathcal{F}} R^M(F_j) \cdot Pr[F = F_j] \\
&= \sum_{j:F_j\in J} R^M(F_j) \cdot Pr[F = F_j] \\
&\quad + \sum_{j:F_j\in \mathcal{F}\setminus J} R^M(F_j) \cdot Pr[F = F_j] \qquad (16)\\
&\leq \frac{4(e+2)}{\sigma}
\end{aligned}$$

□

A.4 Section 5

First, we will show that the first order statistic of n i.i.d samples from F is also an M.H.R distribution.

Lemma 20. *Let F^n be the distribution of the first order statistic of n i.i.d samples from F. F^n has a monotone non-decreasing hazard rate.*

Proof. Our notation is no accident: it is easy to see that $F^n(x)$ is indeed the c.d.f of the first order statistic of n i.i.d samples from F. Let f_n denote the p.d.f, and H_n denote the hazard rate function of F^n. We can differentiate $F^n(x)$ to get

$$f_n(x) = nF^{n-1}(x)f(x)$$

We can now write and manipulate the hazard rate as follows.

$$\begin{aligned}
H_n(x) &= \frac{nF^{n-1}(x)f(x)}{1 - F^n(x)} \\
&= n\left(\frac{f(x)}{1 - F(x)}\right)\left(\frac{F^{n-1}(x)}{\sum_{i=0}^{n-1} F^i(x)}\right) \\
&= nH(x)\left(\frac{F^{n-1}(x)}{\sum_{i=0}^{n-1} F^i(x)}\right)
\end{aligned}$$

Note that $H(x)$ and $F(x)$ are nondecreasing. Therefore, by the above expression, in order to show that $H_n(x)$ is non-decreasing it suffices to show that $g(y) = y^{n-1}/\sum_{i=0}^{n-1} y^i$ is non-decreasing in y. To show this, we take $\alpha \geq 1$ and observe that $g(\alpha y) = \alpha^{n-1}y^{n-1}/\sum_{i=0}^{n-1} \alpha^i y^i \geq \alpha^{n-1}y^{n-1}/\sum_{i=0}^{n-1} \alpha^{n-1} y^i = g(y)$. □

Now, we show a bound on the integral of the product of a monomial and a non-increasing function that will prove useful.

Lemma 21. *Let $g : [0, 1] \to \mathbb{R}^+$ be a non-increasing function. For all integers $n \geq 1$ we have*

$$\frac{\int_{z=0}^{1} g(z)z^n dz}{\int_{z=0}^{1} g(z)z^{n-1} dz} \leq \frac{n}{n+1}$$

Proof. Let $\alpha_n = \int_{z=0}^{1} g(z)z^n dz$. We can integrate by parts using the rule $\int udv = uv - \int vdu$, and setting $dv = z^n dz$ and $u = g(z)$ to get

$$\begin{aligned}
\alpha_n &= \left[g(z)\frac{z^{n+1}}{n+1}\right]_{z=0}^{1} - \int_{z=0}^{1} \frac{z^{n+1}}{n+1} g'(z)dz \\
&= \frac{g(1)}{n+1} - \int_{z=0}^{1} \frac{z^{n+1}}{n+1} g'(z)dz
\end{aligned}$$

To complete the proof, it suffices to show that $(n+$

$1)\alpha_n \leq n\alpha_{n-1}$

$$\begin{aligned} & n\alpha_{n-1} - (n+1)\alpha_n \\ &= \int_{z=0}^{1} (z^{n+1} - z^n) g'(z) dz \geq 0 \end{aligned}$$

Where the inequality follows from the fact that $g'(z) \leq 0$ and $z^{n+1} - z^n \leq 0$. This completes the proof. □

Finally, we bound the ratio of two harmonic numbers in terms of the natural logarithm. Here, we use $\mathcal{H}_n$ to denote the nth harmonic number.

Lemma 22. *For $m \leq n$, $\frac{\mathcal{H}_m}{\mathcal{H}_n} \geq \frac{\log m}{\log n}$.*

Proof. Let $\delta_n = \mathcal{H}_n - \log n$. It is known that δ_n is nonnegative – for completeness we prove it here.

$$\begin{aligned} \delta_n &= \sum_{i=1}^{n} \frac{1}{i} - \int_{x=1}^{n} \frac{1}{x} dx \\ &\geq \sum_{i=1}^{n-1} \left(\frac{1}{i} - \int_{x=i}^{i+1} \frac{1}{x} dx \right) \\ &\geq \sum_{i=1}^{n-1} \left(\frac{1}{i} - \int_{x=i}^{i+1} \frac{1}{i} dx \right) = 0 \end{aligned}$$

Next, we prove that δ_n is a decreasing sequence.

$$\begin{aligned} \delta_n - \delta_{n+1} &= (\mathcal{H}_n - \log n) - (\mathcal{H}_{n+1} - \log(n+1)) \\ &= (\log(n+1) - \log n) - \frac{1}{n+1} \\ &= \int_{x=n}^{n+1} \frac{1}{x} dx - \frac{1}{n+1} \\ &> \int_{x=n}^{n+1} \frac{1}{n+1} dx - \frac{1}{n+1} = 0 \end{aligned}$$

Now we are ready to complete the proof.

$$\begin{aligned} \frac{\mathcal{H}_m}{\mathcal{H}_n} &= \frac{\log m + \delta_m}{\log n + \delta_n} \\ &= \frac{(1 + \frac{\delta_m}{\log m}) \log m}{(1 + \frac{\delta_n}{\log n}) \log n} \\ &\geq \frac{(1 + \frac{\delta_n}{\log n}) \log m}{(1 + \frac{\delta_n}{\log n}) \log n} = \frac{\log m}{\log n} \end{aligned}$$

The last inequality follows from $\delta_m \geq \delta_n \geq 0$ and $\log n \geq \log m \geq 0$. □