Posting Prices with Unknown Distributions

MOSHE BABAIOFF, Microsoft Research
LIAD BLUMROSEN, The Hebrew University
SHADDIN DUGHMI, University of Southern California
YARON SINGER, Harvard University

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 personalized take-it-or-leave-it price for each arriving bidder and aims to maximize revenue. We study how well can such sequential posted-price mechanisms approximate the optimal revenue that would be achieved if the distribution was 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 for randomized mechanisms and 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 simple posted-price mechanism that guarantees a constant fraction of the known-distribution revenue when all candidate distributions exhibit the monotone hazard rate property.

CCS Concepts: • **Theory of computation** → *Algorithmic mechanism design*; *Computational pricing and auctions;*

Additional Key Words and Phrases: Pricing, auctions, posted prices, mechanism design, hazard rate

ACM Reference Format:
Moshe Babaioff, Liad Blumrosen, Shaddin Dughmi, and Yaron Singer. 2017. Posting prices with unknown distributions. ACM Trans. Econ. Comput. 5, 2, Article 13 (March 2017), 20 pages.
DOI: http://dx.doi.org/10.1145/3037382

1. INTRODUCTION

Two extreme points of view are prevalent in the literature regarding the distributional knowledge of sellers in markets. In traditional *Bayesian* models, players have accurate distributional beliefs about the uncertain information. This assumption is problematic in many practical mechanism design settings, as collecting the distributional data may be constrained by technical or operational reasons and also by the fact that these data are elicited from market participants that may act strategically also during this preliminary phase of the mechanism. On the other hand, in the worst-case approach to mechanism design (also known as *detail-free* or *prior-free* mechanism design), the

Liad Blumrosen was supported by the Israeli Science Foundation grant number 230/10.
Authors' addresses: M. Babaioff, Microsoft Research, Herzliya Israel; email: moshe@microsoft.com; L. Blumrosen, School of Business administration, The Hebrew University, Mount Scopus, 91905, Jerusalem, Israel; email: blumrosen@gmail.com; S. Dughmi, Department of Computer Science, University of Southern California, Los Angeles, CA; email: shaddin@usc.edu; Y. Singer, School of Engineering and Applied Sciences, Harvard University, MA; email: yaron@seas.harvard.edu.

DOI: http://dx.doi.org/10.1145/3037382

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 approximation ratios that are often disappointing, even for the best auctions.

In this article, we consider a framework that bridges worst-case and Bayesian models. This is done in the spirit of work like that by Bulow and Klemperer [1996] and some recent work in computer science (see Hartline and Roughgarden [2008] and Dhangwatnotai et al. [2010]). In this framework, we study 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. Our goal is to design detail-free mechanisms, in the sense attributed to Wilson [1989], where, despite the lack of distributional knowledge, they perform well compared to the best Bayesian mechanism that knows the distribution. In this article, we consider such mechanisms in a *dynamic* ("online") setting with the following additional restriction: The mechanisms must use *posted prices*. In contrast to traditional *direct-revelation* mechanisms, posted-price mechanisms interact with bidders by offering each one of them a single take-it-or-leave-it offer so bidders do not need to reveal their exact private value.

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 has a private value v_i for the item, and all 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 in $[1, h]$ for some $h > 1$ known to the seller. The seller aims to maximize revenue.

When distributions are unknown, optimal mechanisms would typically run a sampling ("market research") phase and determine future prices according to the gleaned empirical distribution. However, asking the bidders to report their exact willingness to pay is unnatural in many practical cases. Consider, for instance, an online travel agency (e.g., *Expedia.com* or *Orbitz.com*) trying to sell an airline ticket to a sequence of bidders; the agency typically offers a price to each arriving customer and observes if customers accept or reject prices, but bidders are not expected to reveal the values they would have agreed to pay. Bidders may prefer revealing minimal information on their values if they plan to participate in similar markets in the future, and this holds especially for bidders who have no real chance of winning. This article tries to explore whether the seller can learn "effectively" even when bidders only accept or reject offers without revealing exact values.

This article therefore considers 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 without reporting the actual willingness to pay. 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 some circumstances, the seller might learn details about the underlying distribution; for example, after a series of high prices had been rejected, 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 very limited.

In this article, we would like to measure how much revenue can be obtained using sequential posted-price mechanisms with unknown distributions. We compare this revenue to the optimal expected revenue achievable in a dynamic mechanism when the distribution F is known to the seller, and we denote this revenue by $R^{on}(F)$.[1]

When considering existing positive results regarding the power of posted-price mechanisms (e.g., Blumrosen and Holenstein [2008], Chawla et al. [2010], and Adamczyk et al. [2015]), one might hope that posted prices can work reasonably well even with unknown distributions. Nonetheless, our first main result is negative and shows that posted-price mechanisms can only obtain a diminishing fraction of $R^{on}(F)$.

For a benchmark B, we say that a mechanism M achieves a β-approximation if for every possible input it obtains at least $1/\beta$ fraction of the benchmark B.

THEOREM. *When $\mathcal{F}$ contains all possible distributions with a support in $[1, h]$, no deterministic sequential posted-price mechanism guarantees better than a $\Omega(\frac{\log h}{\log\log h})$-approximation to $R^{on}(F)$.*

We show that this impossibility result is nearly tight, as there is a simple, deterministic posted-price mechanism that achieves a $O(\log h)$-approximation to this revenue benchmark.

Our results are given for every number of bidders n and are parametrized by h (the ratio between the highest possible value and the lowest one, where we normalize the lowest possible value to 1). Note that if h was a small constant, then a constant-factor approximation would be possible by simply posting the price of 1. An interesting aspect of our results is the tradeoff between n and h. It is quite intuitive that when n is very small with respect to h, then without any distributional knowledge only a small fraction of the optimal revenue can be achieved (a small number of offers are spread in a large interval, and thus it is inevitable that some bidders with high value will accept a much lower offer). Our impossibility results become more interesting and more difficult when n is large, where we still show that good approximation ratio is impossible.

The above theorem is proved by constructing a hard instance composed of a relatively small set of about $\frac{\log h}{\log\log h}$ different distributions. We show that every sequence of prices must achieve poor revenue for at least one of these distributions.

We next consider randomized mechanisms and show that they also cannot achieve good approximations.

THEOREM. *When $\mathcal{F}$ contains all possible distributions with a support in $[1, h]$, no randomized sequential posted-price mechanism guarantees better than a $\Omega(\log\log h)$-approximation to $R^{on}(F)$. Moreover, when $n > \frac{\log h}{\log\log h}$ the bound improves to $\Omega(\frac{\log h}{\log\log h})$.*

To prove the theorem, we first prove the same bounds for deterministic mechanisms that are given a Bayesian prior over the family of distributions and then use Yao's min-max principle. Specifically, we present a distribution g over a family of distributions $\mathcal{F}$ and show that any deterministic mechanism cannot achieve such approximations in expectation over the prior g on the family of distributions $\mathcal{F}$.

As the above hardness results show, one must restrict the set of possible distributions for obtaining positive results in our model. In our main positive result, we construct a mechanism that guarantees a constant fraction of the known-distribution

[1]One may also wish to compare this revenue to the optimal revenue in the setting where all bidders are simultaneously present in the market. However, it is known that this revenue and $R^{on}(F)$ differ by at most a small constant factor, and therefore our result applies to this benchmark as well: When the distribution is known, sequential posted-price mechanisms achieve for standard (Myerson-regular) distributions at least 78% of the optimal (Myerson) revenue in large markets ([Blumrosen and Holenstein 2008]) and at least half of the Myerson revenue when multiple items are for sale ([Chawla et al. 2010]).

Distribution Class	Impossibility Results	Approximation
All distributions with support in $[1, h]$	$\Omega\left(\frac{\log h}{\log\log h}\right)$ (deterministic, Sec 3)	$4\log h$ (deterministic,Sec 3.1)
	$\Omega(\log\log h)$ (randomized, Sec 4) $\Omega\left(\frac{\log h}{\log\log h}\right)$ (rand, large n, Sec 4)	$2\log h$ (randomized , [Balcan et al. 2008], Sec 3.1)
All monotone hazard rate distributions with support in $[1, h]$	$h^{\frac{1}{n}}$ (deterministic, Sec 3.2)	$\frac{2e}{1-\epsilon}$ (determ., for $\log h \le n^{\epsilon}$,Sec 5)

Fig. 1. The table summarizes the main results of this article. For each result, we mention if the result is given for deterministic or randomized mechanisms. We also point to the sections where the results are formally stated and proven.

revenue when all candidate distributions have monotone hazard rate[2] (that is, $\frac{f(x)}{1-F(x)}$ is non-decreasing, where f is the density function of F). For this approximation result, we require that n will be large enough with respect to $\log h$, and their ratio will affect the approximation we obtain. For instance, the theorem shows that when $\sqrt{n} > \log h$, our mechanism achieves at least $\frac{1}{4e} \cong 9\%$ of the revenue achieved when the distribution is known. In general,

THEOREM. *Assume that* $n^{\epsilon} > \log h$ *for some constant* $0 < \epsilon < 1$. *Then, there exists a deterministic mechanism that achieves a* $\frac{2e}{1-\epsilon}$*-approximation to* $R^{on}(F)$*, when all the distributions in* $\mathcal{F}$ *satisfy the monotone hazard-rate condition.*

Our proposed mechanism is simple, while its analysis is more involved. We define $\log h$ price levels, $h/2, h/4, h/8, \ldots, h/2^i, \ldots, 2, 1$, 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 trivially satisfy the monotone hazard rate condition), thus constant approximation is impossible with n small relative to $\log h$. Our results are summarized in Figure 1.

Related Work: The two main features of our model are the use of sequential posted prices and the lack of distributional knowledge. There is a recent line of research studying posted-price mechanisms with *known* priors. Blumrosen and Holenstein [2008] studied posted-price mechanisms with known distributions, both both static and dynamic environments, computed their exact revenue, and compared it to the optimal (Myerson) revenue. Chawla et al. [2010] studies sequential posted pricing with known distributions in more general models (matroid-based allocation rules and other multi-dimensional settings) and presented several constant approximations to the optimal revenue. A polynomial-time approximation mechanism for a seller that sells multiple copies with a known prior was given by Chakraborty et al. [2010]. It was shown recently (e.g., Feldman et al. [2015]) that posted prices can achieve good approximation for combinatorial auctions as well.

Several articles in the economic literature studied markets where an underlying distribution exists but is unknown. Rothschild [1974] integrated learning into the strategy of a buyer who is looking for the optimal price when no distribution is known. More recently, Gershkov and Moldovanu [2009a] studied dynamic auctions settings where the distribution of the buyers' preferences is unknown to the seller (but prior over possible distributions is known). They characterized necessary and sufficient conditions for information-theoretic optimum to be implementable in equilibrium.

[2]The non-decreasing hazard rate condition is standard in mechanism design (see, for example, Krishna [2002] and in recent computer-science work of Chawla et al. [2007] and Hartline and Roughgarden [2008]). It is satisfied by many natural distributions, including the exponential, uniform, and binomial distributions.

Gershkov and Moldovanu [2009a] did not study the magnitude of inefficiency in this setting, and in this sense our work complements their work. In subsequent work, Gershkov and Moldovanu [2009b] characterized the second-best solution (that maximizes efficiency under the incentive and information constraints) and showed that the optimal mechanism is deterministic. A similar model was studied by Wang and Chen [1999]; in their work, values are drawn from one of two possible distributions. They showed a sufficient condition (called hazard-rate dominance) for which the optimal posted prices are declining over time. A recent follow-up work to our article [Babaioff et al. 2012, 2015] extended our setting to multi-unit environments and showed mechanisms that achieve nearly tight approximate revenue as a function of the number of units and the number of buyers. Balcan et al. [2008] present revenue approximation results by online item pricing mechanisms for buyers with complex adversarial preferences. The lower bounds in Balcan et al. [2008] rely on the complex preference structure and thus do not apply in our simple model; their positive results are distribution free but degrade in n, while we achieve a constant approximation by assuming monotone hazard rate (MHR) prior distributions.[3]

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. Finally, in Section 5, we present a positive result for monotone hazard rate distributions.

2. PRELIMINARIES

We consider a model where a seller has one item for sale, and the seller's value (opportunity cost) for the item is 0. 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.). Bidder i arrives at period i and leaves the market for good before the next bidder arrives. 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 $(v_1, \ldots, v_n)$ are sampled i.i.d. from a distribution $F \in \mathcal{F}$.[4] Given a distribution F, we use the following notation.

—$W(F)$ denotes the maximum expected social welfare where bidders' values are drawn i.i.d. from F, that is, $E_{\mathbf{v}\sim F^n}[max_{i=1}^{n}\{v_i\}]$.

[3]Our article is also related to the work on *Prophet Inequalities*. Classic work (see Krengel and Sucheston. [1978] and Samuel-Cahn [1984]) showed stopping rules for values that arrive online and are drawn from independent (possibly non-identical) distributions. The connection to online auctions was made by Hajiaghayi et al. [Kleinberg and Sandholm 2007], and this connection was later generalized (see Kleinberg and Weinberg [2012] and Azar et al. [2014]). In particular, Azar et al. [2014] studied models where the seller does not know the underlying distribution of the buyers but it can use few samples from this distribution. Several recent articles (e.g., Babaioff et al. [2009] and Babaioff et al. [2007]) studied versions of the *secretary problem*, where an adversary fixes values that arrive in a random order, and stopping rules should be designed. We note that with i.i.d. samples (i.e., taken from identical and independent distributions), any order of values is equally likely, and thus the secretary model is weaker than our model (formally, this holds for continuous distributions where ties occur with probability zero) 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 article studies stopping rules that are based on a threshold-based decisions at each stage, without observing the exact value of each secretary.

[4]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 A.1 (in Appendix A.1) 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 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 article we add the natural assumption that all agents distributions are *identical*.

—Given a list of posted prices $\mathbf{p} = (p_1, \ldots, p_n)$, let $R^M(F)$ be the expected revenue obtained in a posted-price mechanism M that offers a price p_i to the ith arriving bidder with value v_i, sampled from F (if the item was not sold in previous periods). The price vector $\mathbf{p}$ is implicit in this notation but it will be clear from the context when used.

—Let $R^{on}(F)$ be the *optimal expected revenue* in a dynamic posted-price mechanism when the distribution F is *known* to the seller, that is,

$$R^{on}(F) = max_{M|M \text{ posts prices } \mathbf{p}\in\mathbb{R}^n} R^M(F).$$

Remark 2.1. Since the values are drawn i.i.d., the revenue in any posted-price mechanism is dominated by another mechanism for which $p_1 \geq p_2 \geq \cdots \geq p_n$, that is, with decreasing posted prices. We therefore implicitly assume in the proofs of some of our hardness results that prices are decreasing.

3. AN IMPOSSIBILITY RESULT 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 ($R^{on}(F)$). We now present our main impossibility result.

THEOREM 3.1. *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 other words, there is a constant c such that for every deterministic posted-price mechanism there exists a distribution F such that $\frac{R^M(F)}{R^{on}(F)} < c \cdot \frac{\log\log h}{\log h}$.

Let $\alpha = \frac{\log h}{\log\log h}$. To prove the theorem we define the following finite family $\mathcal{F}_\alpha$ of distributions with support in $[1, h]$. For every index j, we define the distribution F_j as follows: $Pr[x = 1] = 1 - j/n$ and $Pr[x = \alpha^i] = 1/n$ for $1 \leq i \leq j$. Let $\mathcal{F}_\alpha$ be the family of distributions that includes every such distribution F_j with support in $[1, h]$. Since $\alpha = \frac{\log h}{\log\log h}$, the size of $\mathcal{F}_\alpha$ must be approximately α as the following simple observation shows (proof appears in Appendix A.2).

OBSERVATION 3.2. *For the above set of distributions $\mathcal{F}_\alpha$, and for large-enough h, we have that $\alpha - 1 \leq |\mathcal{F}_\alpha| \leq 2\alpha$.*

Before presenting the formal proof we briefly sketch the outline of the proof. If a mechanism knew which one of the F_j is the true distribution, then it could easily obtain revenue of about α^j (this is shown in Lemma 3.3 below). Let r_j denote the maximal number of times that a price α^j is offered to bidders by some mechanism; We show that the overall revenue is at most $\alpha^j O(\frac{r_j}{n})$ if the true distribution is F_j (Lemma 3.4 below). However, as we only offer n prices and there are α relevant price levels, one of the prices must be offered at most n/α times so the mechanism achieves revenue of at most α^{j-1} for some distribution—and this revenue is a factor α away from α^j.

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 3.3. *For any distribution $F_j \in \mathcal{F}_\alpha$ it holds that*

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

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 of at least $1-(1-1/n)^n \leq 1-e^{-1}$. □

We want to show that no posted-price mechanism can approximate this revenue, and 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 $k_0 = \lceil \alpha \rceil$. A proof can be found in Appendix A.2.1.

LEMMA 3.4. *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 3.1.

PROOF (OF THEOREM 3.1). We assume that $n > 4\alpha = 4\frac{\log h}{\log\log h}$; otherwise, we can invoke Proposition 3.6 from Section 3.2 that claims that no deterministic posted-price mechanism can achieve a better approximation than $h^{\frac{1}{n}}$, which is $\Omega(\log h)$ when $n < 4\frac{\log h}{\log\log h}$.

By Observation 3.2, there are at least $\alpha - 1 > \alpha/2$ distributions in $\mathcal{F}_\alpha$. This implies that for at least one j it holds that $r_j < \frac{2n}{\alpha}$.

For h large enough,[5] 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 3.4 for distribution F_j to show that as $\frac{r_j}{n} \leq \frac{2}{\alpha}$ we have

$$\begin{aligned} R^M(F_j) &\leq \alpha^j \cdot \left(\frac{2}{\alpha} + \frac{4e \cdot r_j}{n}\right) \\ &\leq \alpha^j \cdot \left(\frac{2}{\alpha} + \frac{4e \cdot 2}{\alpha}\right) \\ &= \alpha^j \cdot \frac{8e+2}{\alpha} < 24\alpha^{j-1}. \end{aligned} \tag{2}$$

By Lemma 3.3, $R^{on}(F_j) \geq (1-e^{-1})\alpha^j$. Taken together with Equation (2), we have the following:

$$R^M(F_j) \cdot \alpha \cdot \frac{1-e^{-1}}{24} < R^{on}(F_j),$$

which concludes the proof of the theorem. □

[5]Note that the $\Omega(\cdot)$ notation in the theorem implies that the inapproximability result holds for large-enough h's.

3.1. A Simple Positive Result

We now show that Theorem 3.1 is nearly tight, as a logarithmic approximation can be achieved by a simple deterministic mechanism. We note that a logarithmic approximation can also be achieved by a single-price randomized mechanism as in Balcan et al. [2008] (the result in Balcan et al. [2008] holds for more general settings as well). For completeness, we present the randomized mechanism for our setting below.

For a vector of realized values $\mathbf{v} = (v_1, v_2, \ldots, v_n)$, 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 notations, we present the following proposition. Note that if a mechanism obtains expected revenue that approximates the optimal social welfare, then it obtains at least the same approximation factor to any revenue benchmark, as bidders never pay more than their value.

PROPOSITION 3.5. *There exists a deterministic mechanism M that achieves a* $4\log h$*-approximation to the optimal expected social welfare when* $n \geq \log h$*, that is,* $\frac{W(F)}{R^M(F)} \leq 4\log h$ *for every distribution F.*

In addition, from Balcan et al. [2008], there exists a randomized posted-price mechanism M that achieves revenue that is a $2\log h$*-approximation to the realized social welfare, that is,* $\frac{W(\mathbf{v})}{R^M(\mathbf{v})} \leq 2\log h$ *for every vector* $\mathbf{v} \in [1, h]^n$*.*

PROOF *The deterministic mechanism is as follows.*

(**Equal-Sample-of-Every-Scale Mechanism:**)
The mechanism offers the price $h/2^i$ *to* $\lfloor n/(\log h)\rfloor$ *agents, for every* $i \in \{1, \ldots, \log h\}$ *in that order.*

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.[6]

A randomized mechanism: Choose a price p a random from the set $\{2^j\}$, where $j \in \{1, \ldots, \log h\}$ 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. □

3.2. A Simple Impossibility Result

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. The following proposition allows us not to handle cases where n is small when proving our main results.

PROPOSITION 3.6. *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*

$$\frac{R^{on}(F)}{R^M(F)} > h^{1/n} - \epsilon. \quad (3)$$

PROOF. Let $p_1 \geq p_2 \geq \cdots \geq p_n$ 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 $1+\epsilon < p_n$, then the approximation ratio will be unbounded. A second

[6] 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.

observation is that the ratio between some pair of consecutive prices must be at least $h^{\frac{1}{n}}$; otherwise,

$$h = \frac{h}{p_1} \cdot \frac{p_1}{p_2} \cdot \frac{p_2}{p_3} \cdots \frac{p_{n-1}}{p_n} < (h^{\frac{1}{n}})^n = h. \tag{4}$$

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 $\frac{p_{i-1}-\epsilon}{p_i} \geq h^{\frac{1}{n}} - \epsilon$. □

4. PRIORS OVER DISTRIBUTIONS AND RANDOMIZED MECHANISMS

In this section, we extend the impossibility result presented in Theorem 3.1 to mechanisms in which bidders are offered random prices. 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 4.1. *No deterministic posted-price mechanism obtains a constant expected revenue approximation (over g). Specifically, no such mechanism achieves an approximation better than*

—$\Omega(\log h / \log\log h)$ *when* $n > 4\alpha$.
—$\Omega(\log\log h)$ *when* $n \leq 4\alpha$.

The hardness result can be read as follows: There exists a constant $c > 0$, $\mathcal{F}$, and a prior g over $\mathcal{F}$, such that for every deterministic posted-price mechanism M it holds that

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

when $n > 4\alpha$ (the right-hand side is $c \cdot \log\log h$ when $n \leq 4\alpha$).

The case $n \leq 4\alpha$ is proved in Proposition A.4 in the Appendix. We now present the proof for $n > 4\alpha$.

We define g to be a distribution over the "hard family of distributions" $\mathcal{F}_\alpha$ (presented in the beginning of Section 3) that picks F_j with probability proportional to $1/\alpha^j$. Formally, let $w_j = 1/\alpha^j$ and let $\sigma = \sum_{j:F_j \in \mathcal{F}_\alpha} w_j$. The distribution F_j is drawn 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 realization of the actual distribution 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 if the seller does not know the realization of the distribution but only knows the prior g from which it is drawn, then no posted-price mechanism can gain more than a constant expected revenue.

LEMMA 4.2. *Let g be the prior over* $\mathcal{F}_\alpha$ *defined above. It holds that* $R^{on}(g) = \frac{1}{\sigma} \cdot \Omega(\frac{\log h}{\log\log h})$.

Proof. By Lemma 3.3, for $F_j \in \mathcal{F}_\alpha$, 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 3.2 $\mathcal{F}_\alpha$ consists of $\Omega(\frac{\log h}{\log\log h})$ distributions, we conclude that $R^{on}(g) = \frac{1}{\sigma}\cdot\Omega(\frac{\log h}{\log\log h})$. □

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

Lemma 4.3. *Assume that $n > 4\alpha$. Let g be the prior over $\mathcal{F}_\alpha$ defined above. For any deterministic posted-price mechanism M, 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 bound is almost tight, as we showed (Proposition 3.5) a simple mechanism that obtains an $O(\log h)$-approximation. Therefore, the following corollary strengthens Theorem 3.1 for randomized mechanisms.

Corollary 4.4. *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 $n > 4\alpha$ (or $\Omega(\log\log h)$ when $n \leq 4\alpha$); that is, 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. MONOTONE HAZARD RATE DISTRIBUTIONS

In light of the impossibility results, one should restrict the class of possible distribution for having a constant approximation to the optimal revenue. In this section, we restrict attention to distributions that satisfy the *monotone hazard rate* assumption.

In this section, we assume that a density function $f(x) = dF(x)/dx$ exists for the underlying distribution and that this density function is always positive and differentiable. We let $S(x) = 1 - F(x)$ denote the survival probability and $H(x) = f(x)/S(x)$ denote the hazard rate of F. In this section, we will show a mechanism that attains a constant approximation ratio for distributions F with $H(x)$ monotone non-decreasing. This MHR assumption is common in auction theory, and MHR distributions include the most natural distributions in this setting. We emphasize that the mechanism has no knowledge of distribution F yet achieves the claimed approximation ratio for every F with non-decreasing hazard rate.

We consider the *Equal-Sample-of-Every-Scale* mechanism from Proposition 3.5. The mechanism offers the price $h/2^i$ to $\lfloor n/(\log h)\rfloor$ agents, for every $i \in \{1, \ldots, \log h\}$ in that order.

Despite its simplicity, our main positive result is that the *Equal-Sample-of-Every-Scale* mechanism achieves a constant approximation for every monotone hazard rate distribution (we already know from Proposition 3.5 that it achieves a $\log h$ approximation for general distributions).

Theorem 5.1. *Let $\log h \leq \lfloor n^\epsilon \rfloor$ 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 maximum social welfare. We note that if the mechanism *Equal-Sample-of-Every-Scale* used a different scale, for example, $1+\delta$ instead of 2, then the same analysis would prove a bound of $\frac{1-\epsilon}{(1+\delta)e}E[X_n]$ for n such that $\log_{1+\delta} h \leq n^\epsilon$.

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 3.6 (Section 3.2) shows that without n being at least $\log h$ a constant approximation is unachievable by deterministic mechanisms.[7]

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

LEMMA 5.2. *When F satisfies the monotone hazard rate condition, we have*

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

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

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

For the last inequality, note that we can write $H(x)$ as a function of $F(x)$, for example, $H(x) = H(F^{-1}(F(x)))$, which is well defined when the cdf f is always positive.

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:

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

where F is an MHR distribution, and therefore $1/H(x)$ is a non-increasing function of x and therefore also of $F(x)$. Applying Lemma A.8 in the appendix with $z = F(x)$ and $g(z) = \frac{1}{H(F^{-1}(z))}$ gives us that $\Delta_n/\Delta_{n-1} \leq \frac{n}{n+1}$, 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 5.3. *When F satisfies the monotone hazard rate condition and for $m \leq n$, it holds that*

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

PROOF. We show the second inequality in Lemma A.9 in the appendix. To show the first inequality, by induction it suffices to show that $E[X_{n+1}]/E[X_n] \leq \mathcal{H}_{n+1}/\mathcal{H}_n$ for all

[7]We note that the lower bound of Proposition 3.6 uses point distributions, which satisfies the monotone hazard rate assumption.

integers $n \geq 1$. 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

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

This, in turn, is equivalent to showing

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

We rewrite the above condition as follows:

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

Therefore, it suffices to show that $\Delta_i/\Delta_n \geq (n+1)/(i+1)$. This can be established by inductive application of Lemma 5.2, completing the proof. □

The above lemma implies that when $m \leq n$, we have that $E[X_m] \geq \frac{\log m}{\log n} E[X_n]$; thus, $\Pr[X_m \geq \frac{\log m}{\log n} E[X_n]] \geq \Pr[X_m \geq E[X_m]]$. Lemma A.7 in the appendix implies that X_m is distributed as a monotone hazard rate distribution. Moreover, a result of Barlow and Marshall [1964] implies—as a special case—that every monotone hazard rate distribution exceeds its expectation with a probability of 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 5.1. The mechanism samples at least $\lfloor n/(\log h) \rfloor$ bidders for each price $2^i \in [1, h]$. The theorem assumes that $\log h \leq \lfloor n^\epsilon \rfloor$, which implies that $\lfloor n/(\log h) \rfloor \geq n^{1-\epsilon}$. It follows that the mechanism samples at least $m = n^{1-\epsilon}$ bidders for each price.

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 the following:

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

where the equality holds as we have $m = n^{1-\epsilon}$, so $1 - \epsilon = \frac{\log m}{\log n}$. □

6. DISCUSSION

Our article considers sequential posted-price mechanisms in environments where the seller has uncertainty regarding the distribution of the bidders' preferences. We show that for general distributions no mechanism can achieve a constant approximation to the optimal revenue. On the other hand, we show that for the class of MHR distributions, this task is achievable, as long as the market is large enough compared to the largest possible value (otherwise, some impossibility results kick in).

We leave several interesting open questions. First, we do not prove any hardness results for MHR valuations. In particular, there may be a mechanism that, for such distributions, achieves an approximation ratio that approaches 1 in large markets.

Another interesting direction is to try to obtain positive results for wider classes of distributions, for example, regular distributions ([Myerson 1981]).

Our article focuses on the simple scenario where the seller has only one unit for sale. An immediate question is whether our approach can be extended to multi-unit settings. A follow-up work to our article [Babaioff et al. 2012, 2015] studied a similar setting but with multiple copies of the item. They designed mechanisms that approximate the optimal revenue, where the approximation factor depended on the market size n and the number of copies. Selling multiple item enables the sellers to use more complex learning techniques. Indeed, they showed how the seller's pricing problem resembles multi-arm bandit problems, and they used ideas from this literature in their analysis.

APPENDIX

A. MISSING CLAIMS AND PROOFS

A.1. Section 1

CLAIM A.1. *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, although it 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 $1 \leq \alpha^{k_i-1} \leq \alpha^{k_i} \leq h$ and $\Pr[p_i \in [\alpha^{k_i-1}, \alpha^{k_i}]]$ is minimized. By Observation A.2, 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 3.2 is a corollary of the following observation. Due to the definition of $\mathcal{F}_\alpha$, it follows from the observation below that if h is large enough, then the family $\mathcal{F}_\alpha$ consists of at most 2α and at least $\lfloor\alpha\rfloor$ distributions. Note that $\lfloor\alpha\rfloor \geq \alpha - 1$.

OBSERVATION A.2. *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 holds if $\frac{\log h}{\log\log h} \cdot (\log\log h - \log\log\log h) < \log h$, which holds, since $\frac{\log\log h - \log\log\log h}{\log\log h} < 1$.

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

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

which is analogous to $2\log\log\log h < \log\log h$, which clearly holds when h is large enough. □

A.2.1. Proof of Lemma 3.4. Following is the proof of Lemma 3.4.

Proof. We need to bound $R^M(F_j)$. If the price is not in $[\alpha^{j-1}, \alpha^j]$, then 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 $\mathbf{v}$.

$R^M(F_j) \leq \alpha^{j-1} + E[R(\mathbf{v})]$, where $E[R(\mathbf{v})]$ is the expectation of $R(\mathbf{v})$.

Let Y be 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,

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

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 Equation (6). Clearly, $E[R(v)|Y = k] \leq \alpha^j$, and 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 Equation (6). The following claim would be useful.

Claim A.3. *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,

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

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

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

As we assume that $n > 4\alpha$ and as $k \leq k_0 \leq 2\alpha$ it holds that $n/2 > 2\alpha \geq k$, and 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 A.3 and Equation (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 Equation (7) to bound the first term of Equation (6),

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

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

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

A.3. Section 4

Proposition A.4. *Let g be a prior distribution over $\mathcal{F}$ which is known to the seller. When $n \leq 4\alpha$, every deterministic posted-price mechanism obtains expected revenue approximation (over g) of no better than $\Omega(\log \log h)$.*

Proof. Let $\mathcal{F}$ be the set of singleton distributions on the values 2^j, for $j = \{0, \ldots, \log h - 1\}$. In other words, $\mathcal{F}$ contains $\log h$ possible distributions, where a distribution F_j in $\mathcal{F}$ selects the value 2^j with probability 1. Let $\sigma = \sum_{i=1}^{\log h - 1} \frac{1}{2^j}$ (a normalization factor). We define the prior g over $\mathcal{F}$ as follows: $g(F_j) = \frac{1}{\sigma \cdot 2^j}$.

A posted-price mechanism that always knows the realization F_j will post a price 2^j and will therefore gain an expected revenue of

$$\frac{1}{\sigma} \sum_{j=0}^{\log h - 1} 2^j \cdot \frac{1}{2^j} = \frac{\log h}{\sigma}. \tag{15}$$

The following claims characterize the optimal posted-price mechanism for this setting. A posted-price mechanism posts n prices, and the optimal offered prices will be of the form 2^j, where $j = \{0, \ldots, \log h - 1\}$ (if a price between such points was offered, then increasing it to the next exponent of 2 would strictly increase revenue). We will denote the prices posted by a mechanism by $2^{j_1}, 2^{j_2}, \ldots, 2^{j_n}$ $(j_1 \leq j_2 \cdots \leq j_n)$ (although the offers are given in decreasing order). We will first argue that the offers will be evenly spread across the possible singleton distributions.

CLAIM A.5. *Consider a posted-price mechanism M that achieves optimal expected revenue for the above setting with prices $2^{j_1}, 2^{j_2}, \ldots, 2^{j_n}$. For every $i, k \in \{0, \ldots, \log h - 2\}$, we have that $|(j_{i+1} - j_i) - (j_{k+1} - j_k)| \leq 1$.*

PROOF. We first observe that the expected payment is invariant of the actual price and depends only on the distance from the next highest price. That is, the expected revenue gained by the agent accepting the price 2^{j_k} depends only on $j_{k+1} - j_k$,

$$\frac{1}{\sigma} \sum_{i=j_k}^{j_{k+1}-1} 2^{j_k} \frac{1}{2^i} = \frac{1}{\sigma} \sum_{i=0}^{j_{k+1}-j_k-1} \frac{1}{2^i}. \tag{16}$$

Now, assume that the condition does not hold and that there exist j, k such that $(j_{i+1} - j_i) - (j_{k+1} - j_k) \geq 2$. We will assume that $k > i$ (the other case is treated analogously). We will show that, in this case, we can change the posted prices and strictly increase revenue. Consider the mechanism M' that posts the same prices but only shifts the prices $j_{i+1}, \ldots, j_k$ one notch to the left, that is, instead of posting the prices $2^{j_{i+1}}, \ldots, 2^{j_k}$, it posts the prices $2^{j_{i+1}-1}, \ldots, 2^{j_k-1}$ (all other prices are as in the M). As observed above, the expected revenue from each price only depends on the distance from the next price. We thus have that the only change in revenue is caused by increasing $j_{i+1} - j_i$ and decreasing $j_{k+1} - j_k$. Overall, the net change in revenue is $\frac{1}{\sigma}(\frac{1}{2^{j_{k+1}-j_k}} - \frac{1}{2^{j_{i+1}-j_i-1}})$, which is strictly positive when $(j_{i+1} - j_i) - (j_{k+1} - j_k) \geq 2$. □

CLAIM A.6. *All posted-price mechanisms for which the condition from Claim A.5 holds (that is, that for every $i, k \in \{0, \ldots, \log h - 2\}$, we have that $|(j_{i+1} - j_i) - (j_{k+1} - j_k)| \leq 1$) achieve the same expected revenue for the $\mathcal{F}$ and g defined above.*

PROOF. Due to the assumption, for some k and for every i, either $j_{i+1} - j_i = k$ or $j_{i+1} - j_i = k+1$. Let $K = |\{i | j_{i+1} - j_i = k+1\}|$, and, since there are n prices overall, K is the same for all mechanisms satisfying the given assumption. Since the expected revenue from each price in the mechanism depends only on the distance (in the exponent scale) from the next highest offer, we get that every two mechanisms satisfying the constraint obtain the same expected revenue. □

Due to the above two claims, the optimal posted-price mechanism achieves an expected revenue of at least (where $k = \lfloor \frac{\log h}{n} \rfloor$),

$$\frac{n}{\sigma} \sum_{i=0}^{k-1} 2^{-i}. \tag{17}$$

The ratio between the optimal revenue with known distributions (Equation (15)) and the optimal posted-price revenue is therefore

$$\frac{\log h}{n} \cdot \frac{1}{\sum_{i=0}^{k-1} 2^{-i}} \geq \frac{\log h}{2n} > 2 \log \log h. \tag{18}$$

The second inequality is due to the assumption that $n \leq 4\alpha$. □

Proof of Lemma 4.3:

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 $k_0 < n/2$, and 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 3.4. Recall that $\sum_j r_j = n$ and that by Observation 3.2 the family $\mathcal{F}$ is of size at most 2α.

$$\begin{aligned}
&\sum_{j:F_j\in\mathcal{F}\setminus J} R^M(F_j)\cdot Pr[F = F_j] \\
&\quad\leq \sum_{j:F_j\in\mathcal{F}\setminus J} \alpha^j \cdot \left(\frac{2}{\alpha} + \frac{4e\cdot r_j}{n}\right)\cdot\frac{w_j}{\sigma} \\
&\quad\leq \frac{1}{\sigma}\cdot\sum_{j:F_j\in\mathcal{F}\setminus J}\cdot\left(\frac{2}{\alpha} + \frac{4e\cdot r_j}{n}\right) \\
&\quad\leq \frac{1}{\sigma}\cdot\left(4 + \frac{4e}{n}\sum_{j:F_j\in\mathcal{F}\setminus J} r_j\right) \\
&\quad= \frac{4(e+1)}{\sigma}. \qquad (19)
\end{aligned}$$

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

$$\begin{aligned}
R^M(g) &= \sum_{j:F_j\in\mathcal{F}} R^M(F_j)\cdot Pr[F = F_j] \\
&= \sum_{j:F_j\in J} R^M(F_j)\cdot Pr[F = F_j] + \sum_{j:F_j\in\mathcal{F}\setminus J} R^M(F_j)\cdot Pr[F = F_j] \\
&\leq \frac{4(e+2)}{\sigma}. \quad \square \qquad (20)
\end{aligned}$$

A.4. Section 5

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

Lemma A.7. *Let F^n be the distribution of the first-order statistic of n i.i.d. samples from a distribution F. If F has a non-decreasing hazard rate, then F^n has a non-decreasing hazard rate.*

Proof. Our notation is no accident: It is easy to see that $F^n(x)$ is indeed the cumulative distribution of the first-order statistic of n i.i.d. samples from F. Let f_n denote the density function 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:

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

Note that $H(x)$ and $F(x)$ are non-decreasing. 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 A.8. *Let $g : [0, 1] \to \mathbb{R}_{>0}$ be a non-increasing, differentiable 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 $\int u dv = uv - \int v du$ 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 $(n+1)\alpha_n \leq n\alpha_{n-1}$

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

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 A.9. *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 non-negative; for completeness we prove it here,

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

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,

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

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

REFERENCES

Marek Adamczyk, Allan Borodin, Diodato Ferraioli, Bart de Keijzer, and Stefano Leonardi. 2015. Sequential posted price mechanisms with correlated valuations. In *Proceedings of the Web and Internet Economics: 11th International Conference (WINE'15)*. Evangelos Markakis and Guido Schäfer (Eds.). Springer, Berlin, 1–15.

Pablo Daniel Azar, Robert Kleinberg, and S. Matthew Weinberg. 2014. Prophet inequalities with limited information. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2014)*. 1358–1377.

Moshe Babaioff, Michael Dinitz, Anupam Gupta, Nicole Immorlica, and Kunal Talwar. 2009. Secretary problems: Weights and discounts. In *Proceedings of the 19th Annual ACM -SIAM Symposium on Discrete Algorithms (SODA'09)*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 1245–1254.

Moshe Babaioff, Shaddin Dughmi, Robert Kleinberg, and Aleksandrs Slivkins. 2012. Dynamic pricing with limited supply. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC'12)*. 74–91.

Moshe Babaioff, Shaddin Dughmi, Robert D. Kleinberg, and Aleksandrs Slivkins. 2015. Dynamic pricing with limited supply. *ACM Trans. Econ. Comput.* 3, 1 (2015), 4.

Moshe Babaioff, Nicole Immorlica, David Kempe, and Robert Kleinberg. 2007. A knapsack secretary problem with applications. In *APPROX-RANDOM*. 16–28.

Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. 2008. Item pricing for revenue maximization. In *Proceedings 9th ACM Conference on Electronic Commerce (EC-2008)*. 50–59.

Richard E. Barlow and Albert W. Marshall. 1964. Bounds for distributions with monotone hazard rate. *Ann. Math. Statist.* 35, 3 (1964), 1258–1274.

Liad Blumrosen and Thomas Holenstein. 2008. Posted prices vs. negotiations: An asymptotic analysis. In *Proceedings of the ACM Conference on Electronic Commerce*.

Jeremy Bulow and Paul Klemperer. 1996. Auctions versus negotiations. *Am. Econ. Rev.* 86, 1 (1996), 180–94.

Tanmoy Chakraborty, Eyal Even-Dar, Sudipto Guha, Yishay Mansour, and S. Muthukrishnan. 2010. Approximation schemes for sequential posted pricing in multi-unit auctions. In *Proceedings of the 6th International Conference on Internet and Network Economics (WINE'10)*. 158–169.

S. Chawla, J. D. Hartline, and R. Kleinberg. 2007. Algorithmic pricing via virtual valuations. In *ACM Conference on Electronic Commerce (EC)*. 243–251.

Shuchi Chawla, Jason Hartline, David Malec, and Balasubramanian Sivan. 2010. Sequential posted pricing and multi-parameter mechanism design. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC'10)*. 311–320.

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

Michal Feldman, Nick Gravin, and Brendan Lucier. 2015. Combinatorial auctions via posted prices. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'15)*. 123–135.

Alex Gershkov and Benny Moldovanu. 2009a. Learning about the future and dynamic efficiency. *Am. Econ. Rev.* 99, 4 (2009), 1576–1587.

Alex Gershkov and Benny Moldovanu. 2009b. Optimal Search, Learning and Implementation. *Journal of Economic Theory* 147, 3 (2012), 881–909.

J. D. Hartline and T. Roughgarden. 2008. Optimal mechanism design and money burning. In *ACM Symposium on Theory of Computing (STOC)*. 75–84.

M. Hajiaghayi R. Kleinberg and T. Sandholm. 2007. Automated online mechanism design and prophet inequalities. In *Proceedings of the 22nd Conference on Artificial Intelligence*.

Robert Kleinberg and S. Matthew Weinberg. 2012. Matroid prophet inequalities. In *Proceedings of the 44th Symposium on Theory of Computing Conference (STOC)*. 123–136.

U. Krengel and L. Sucheston. 1978. On semiamarts, amarts and processes with finite value. *Adv. Probab.* 4 (1978), 197–266.

V. Krishna. 2002. *Auction Theory*. Academic Press.

R. B. Myerson. 1981. Optimal auction design. *Math. Operat. Res.* 6, 1 (1981), 58–73.

Michael Rothschild. 1974. Searching for the lowest price when the distribution of prices is unknown. *J. Pol. Econ.* 82, 4 (1974), 689–711.

Ester Samuel-Cahn. 1984. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *Ann. Probab.* 12, 4 (11 1984), 1213–1216.

Ruqu Wang and Yongmin Chen. 1999. Learning buyers' valuation distribution in posted-price selling. *Econ. Theory* 14, 2 (1999), 417–428.

R. Wilson. 1989. Efficient and competitive rationing. *Econometrica* 57 (1989), 1–40.

Received May 2016; revised September 2016; accepted December 2016