Posted Price Mechanisms for a Random Stream of Customers

JOSÉ CORREA, Universidad de Chile
PATRICIO FONCEA, Universidad de Chile
RUBEN HOEKSMA, Universidad de Chile
TIM OOSTERWIJK, Maastricht University
TJARK VREDEVELD, Maastricht University

Posted price mechanisms constitute a widely used way of selling items to strategic consumers. Although suboptimal, the attractiveness of these mechanisms comes from their simplicity and easy implementation. In this paper, we investigate the performance of posted price mechanisms when customers arrive in an unknown random order. We compare the expected revenue of these mechanisms to the expected revenue of the optimal auction in two different settings. Namely, the nonadaptive setting in which all offers are sent to the customers beforehand, and the adaptive setting in which an offer is made when a consumer arrives. For the nonadaptive case, we obtain a strategy achieving an expected revenue within at least a $1 - 1/e$ fraction of that of the optimal auction. We also show that this bound is tight, even if the customers have i.i.d. valuations for the item. For the adaptive case, we exhibit a posted price mechanism that achieves a factor 0.745 of the optimal revenue, when the customers have i.i.d. valuations for the item. Furthermore, we prove that our results extend to the prophet inequality setting and in particular our result for i.i.d. random valuations resolves a problem posed by Hill and Kertz [13].

Additional Key Words and Phrases: Posted prices; Prophet inequality; Mechanism design

1 INTRODUCTION

Posted price mechanisms constitute an attractive and widely applicable way of selling items to strategic consumers. In this context, consumers are faced with take-it-or-leave-it offers, and therefore strategic behaviour simply vanishes. This type of mechanism has been vastly studied, particularly in the marketing community [5]. In recent years, there has been a significant effort to understand the expected revenue of the outcome generated by different posted price mechanisms when compared to that of the optimal auction [2, 4, 6, 23]. In addition, several companies have started to apply *personalized pricing* to sell their products. Under this policy, companies set different prices for different consumers based on purchase history or other factors that may affect their willingness to pay. For example, the online data provider Lexis-Nexis sells to virtually every user at a different price [22]. In 2012, Orbitz online travel agency found that people who use Mac computers spent as much as 30% more on hotels, so it started to show them different, and sometimes costlier, travel options than those shown to Windows visitors [18]. Similarly, retailers and supermarket chains such as Safeway are

Authors' email addresses: José Correa: correa@uchile.cl; Patricio Foncea: patricio.foncea@ug.uchile.cl; Ruben Hoeksma: rhoeksma@dim.uchile.cl; Tim Oosterwijk: t.oosterwijk@maastrichtuniversity.nl; Tjark Vredeveld: t.vredeveld@maastrichtuniversity.nl.

This work was supported by The Millennium Nucleus Information and Coordination in Networks ICM/FIC RC130003.

EC17, June 26–30, 2017, Cambridge, MA, USA. Copyright 2017 ACM ISBN 978-1-4503-4527-9/17/06...$15.00

 ACM ISBN 978-1-4503-4527-9/17/06
http://dx.doi.org/10.1145/3033274.3085137

using data culled from billions of purchases to offer deals tailored to specific shoppers [15]. Choudhary et al. [7] further investigated this issue, providing more examples and developing a theoretical framework to analyze equilibria between firms that apply personalized pricing and those who do not.

In its simplest form, the problem we consider is described as follows. A monopolist sells a single item to a set of known potential buyers. The seller places no value on the item, while the buyers have independent, not necessarily identical, random valuations for the item. The main question is to design a mechanism maximizing the revenue of the seller. This question was answered in a seminal paper by Myerson [19], and the solution is, in some situations, a remarkably simple mechanism. However, in many situations it is hard to implement, and the mechanism of choice turns out to be a simple posted price mechanism. A common example of this practice is that of *direct mail campaigns*, in which the seller contacts its potential buyers directly and offers each one a certain price for the item. The item is then sold to the first consumer who accepts the offer [5, 8].

In this paper, we investigate the performance of posted price mechanisms to sell a single item to a given set of customers who arrive in a random unknown order. We consider two different models which share the property that each customer is offered the item at most once. Upon receiving an offer, a customer immediately decides whether to buy the item at that price or to pass and simply not buy. The *nonadaptive* model considers the situation in which all offers have to be made simultaneously, and customers respond in random order, akin to direct mail campaigns. The *adaptive* model considers a situation in which the seller may adapt the offer. Here, customers again arrive in random order. Whenever a customer arrives, she is offered the item at a price, which the seller may base on the customer he is offering to, as well as the customers who already rejected earlier offers.

Problem description. A seller has a single item to sell to a given set of customers $\mathcal{I}$. We assume that the seller has no value for keeping the item. Customers have independent random valuations for the item with customer $i \in \mathcal{I}$ valuing the item at v_i, drawn from distribution $F_i(\cdot)$. The customers arrive in (uniform) random order, and the goal of the seller is to maximize his expected revenue. To this end, we consider a nonadaptive and an adaptive scenario.

Nonadaptive: The seller sets prices $p_i \geq 0$ for all $i \in \mathcal{I}$, with the goal of maximizing his expected revenue, defined as

$$\sum_{i \in \mathcal{I}} p_i \mathbb{P}_{\sigma, v} \left[i = \operatorname*{argmin}_{j \in \mathcal{I}} \{\sigma(j) \mid v_j \geq p_j\} \right] ,$$

where the probability is taken over the arrival permutation σ and the customers' valuations v.

Adaptive: The seller offers each customer a price as she arrives. So, the seller sets functions $p_i : 2^{\mathcal{I}} \to \mathbb{R}$ for each customer i, such that, if S is the set of customers who already arrived and declined the offer, $p_i(S)$ is the price offered to customer i if she is next to arrive. For an arrival permutation σ, we denote $p_i(\sigma) = p_i(\{\sigma^{-1}(1), \ldots, \sigma^{-1}(\sigma(i) - 1)\})$, and therefore we can write the seller's expected revenue as

$$\mathbb{E}_{\sigma} \left[\sum_{i \in \mathcal{I}} p_i(\sigma) \mathbb{P}_v \left[i = \operatorname*{argmin}_{j \in \mathcal{I}} \{\sigma(j) \mid v_j \geq p_j(\sigma)\} \right] \right] ,$$

where the expectation is taken over the arrival permutation σ, and the probability is taken over the customers' valuations v.

Our results. We present two posted price mechanisms: A nonadaptive posted price mechanism that guarantees an expected revenue within a factor $1 - 1/e$ of that of Myerson's optimal auction and an adaptive posted price mechanism for i.i.d. value distributions that has a guaranteed expected revenue of 0.745 of that of Myerson's optimal auction.[1]

[1] One may think here that the right benchmark should be the expectation of the maximum valuation. However, this cannot yield useful results. Consider a single customer whose valuation lies in $[1, +\infty)$ distributed according to $F(v) = 1 - 1/v$.

The factor achieved by our nonadaptive mechanism complements recent work of Alaei et al. [3] by showing that personalized pricing can increase the revenue from $1/e$ to $1 - 1/e$. On the other hand, the bound matches the well known result of Chawla et al. [6], who designed a *sequential* posted price mechanism with the same approximation guarantee. Although their mechanism is also nonadaptive in the sense that the selected prices are fixed a priori, it has the power to choose the arrival order of the customers. Thus, making it easier to extract revenue by offering to *good* customers first. Furthermore, this bound also matches the approximation guarantee obtained by Esfandiari et al. [10], who also consider the random arrival model, but in their mechanism the sequence of prices depends on the arrival order of customers, and it is therefore adaptive (according to the definition in this paper). Besides the natural application of our nonadaptive setting, it is interesting to note that one can achieve this approximation factor in the random arrival model without using adaptivity. Also, as opposed to previous results, we prove that the bound of $1 - 1/e$ is best possible for our setting.

Theorem 1.1. *For any given set of potential customers $\mathcal{I}$, there exists a nonadaptive posted price mechanism that achieves an expected revenue of at least a $1 - 1/e$ fraction of that of Myerson's optimal auction on $\mathcal{I}$.*

In the case of monotone virtual valuations[2], the algorithm that achieves this result becomes remarkably simple.

Input: Customers $i \in \mathcal{I}$ with valuation distributed according to F_i.
Algorithm 1:

(1) Compute q_i = probability that optimal auction assigns to i.
(2) Discard customer i with probability $1 - \frac{2}{2+(e-2)q_i}$.
(3) Offer non-discarded customers price $F_i^{-1}(1 - q_i)$.
(4) Item is allocated to a random customer accepting the offer.

The algorithm, while randomized, can be derandomized using standard techniques.

Algorithm 1 may seem counterintuitive since, in step (2), the higher the probability is that a customer wins the optimal auction, the higher the probability is that the algorithm discards her (though the probability of discarding any customer is at most $1 - \frac{2}{e}$). The following example gives some intuition on why agents need to be discarded. Consider just two customers: customer 1 who has deterministic value equal to 1, and customer 2 who values the item at 100 with probability 1/10 and at 0 with probability 9/10. In this situation, the optimal mechanism assigns the item to customer 2 with probability 9/10 and the total expected revenue is 10+9/10. Now, if a nonadaptive algorithm makes offers to both customers the expected revenue is (1/10)(50.5)+(9/10)(1)=5.95, which is not within the claimed ratio of the optimal mechanism. Another somewhat surprising element of Algorithm 1 is that the probability of not assigning the item can be computed as $\prod_{i \in I}(1 - 2q_i/(2 + (e - 2)q_i))) \geq 2/e$. Again, the previous example provides intuition to the fact that, if we shoot for an algorithm that assigns too frequently, we risk assigning the item for too low a price. This intuition does not hold in the adaptive case.

The cornerstone of our analysis is a basic result about Bernoulli random variables which may be of independent interest. The result states that if we are given a set of nonhomogeneous independent Bernoulli random variables with associated prizes, then there is a subset of variables so that the expected average prize of the successes is at least a factor $1 - 1/e$ of the expectation of the maximum prize over all random variables.

Clearly, if we charge price p the acceptance probability is $1/p$, for a total revenue of 1. On the other hand, the expectation of the valuation is actually $+\infty$. This example can easily be turned into one with finite expectation but arbitrarily large ratio between the optimal pricing and the expectation of the random variable.

[2]Recall that for a valuation distribution F, the virtual valuation is defined as $v - \frac{1-F(v)}{f(v)}$.

Lemma 1.2 (Bernoulli Selection Lemma). *Given a set $N = \{1, \ldots, n\}$ of independent Bernoulli random variables $X_1, \ldots, X_n$, where $X_i = 1$ with probability q_i and 0 otherwise, and associated prizes $b_1, \ldots, b_n$. The following inequalities hold:*

$$\frac{e}{e-1} \max_{S \subseteq N} \mathbb{E}\left[\frac{\sum_{i \in S} b_i X_i}{\sum_{i \in S} X_i}\right] \geq \max_{z_i \leq q_i} \left\{ \sum_{i \in N} b_i z_i \,\middle|\, \sum z_i \leq 1 \right\} \geq \mathbb{E}[\max_{i \in N} \{b_i X_i\}] \,.$$

Here, when evaluating the leftmost term, we define $0/0 = 0$.

To prove the lemma, we consider a continuous relaxation of the maximization problem, and then guess a solution in which each random variable is included in S with some instance-dependent probability. Then, we look for the worst possible instance by applying the first order optimality conditions of a nonlinear problem. These conditions reveal some structural insight on what a worst case instance looks like. Using this, we obtain the desired bound. Theorem 1.1 follows from Lemma 1.2 with fairly little extra work. The basic tool for this is a fundamental lemma by Chawla et al. [6, Lemma 4] that upper bounds the revenue of the optimal auction.

To complement our results, we provide instances that show that the bounds in Lemma 1.2 and Theorem 1.1 are tight. In particular, we show that even with *independent identically distributed* (*i.i.d.*) customer valuations the bound of Theorem 1.1 cannot be beaten. Therefore, adaptivity is necessary to go beyond $1 - 1/e$, even with i.i.d. distributions. For this setting we show the following.

Theorem 1.3. *For any given set of potential customers $\mathcal{I}$ whose values are independent and identically distributed, there exists an adaptive posted price mechanism that achieves an expected revenue of at least a $1/\beta > 0.745$ fraction of that of Myerson's optimal auction on $\mathcal{I}$, where β is the unique value such that*

$$\int_0^1 \frac{1}{y(1 - \ln(y)) + (\beta - 1)} dy = 1 \,.$$

To achieve this result we use a quite natural idea: as less customers are left, the price should decrease. Besides this the key ingredient of our algorithm is to use random prices drawn from a well chosen distribution that mimics an expression we obtain for the revenue of an optimal auction. Again in the case of monotone virtual valuations our algorithms is as follows:

Input: Customers $i \in \mathcal{I}$ with valuation i.i.d. according to F.
Algorithm 2:

(1) Partition the interval $[0, 1]$ into intervals $A_i = [a_{i-1}, a_i]$, s.t. $a_0 = 0$, $a_n = 1$.
(2) Sample q_i from A_i with an appropriately chosen distribution.
(3) When the i–th buyer comes, offer price $p_i = \max\{F^{-1}(1 - q_i), v^*\}$, where v^* is the reservation price of the optimal auction.

Like Algorithm 1, also Algorithm 2 can be derandomized using standard techniques.

We believe that the bound of Theorem 1.3 is tight. Although the best upper bound known for the i.i.d. case, due to Blumrosen and Holenstein [4], proves that no algorithm can achieve a fraction of at least 0.79, we believe that the family of instances provided by Hill and Kertz [13] in the context of prophet inequalities for i.i.d. random variables can be transformed into a tight family of instances for Theorem 1.3. We remark here that recent work of Dütting et al. [9] also studies the benefit of adaptivity in the i.i.d. case, but from a different perspective.

Due to space constraints many technical proofs have been omitted in this paper.

Prophet Inequalities. As it is common in the literature, we prove that our Theorems 1.1 and 1.3 also hold in the context of the vastly studied prophet inequalities [13, 14, 16, 17, 20, 21], whose study started in the sixties with the work of Gilbert and Mosteller [11].

In particular Theorem 1.1 becomes Corollary 2.2 and is related to recent work of Esfandiari et al. [10]. In this setting, known as *prophet secretary* or *full information secretary problem*, we are given n independent nonnegative random variables $X_1, \ldots, X_n$ that arrive in random order. Upon arrival of a random variable we see its value and may decide to keep it and finish or to drop it and continue. The goal is to design an algorithm that obtains a value within a large fraction of the expectation of the maximum of all X_i's. In Corollary 2.2 we use Lemma 1.2 to prove that, even nonadaptively, we may obtain a $1 - 1/e$ fraction, with an algorithm that sets a threshold for each random variable (independent of the arrival order) in advance and accepts the first random variable with value above the threshold. Moreover, this bound is also tight, which follows from the example in Section 2. In the case in which the probability of having two X_i's being the maximum is zero the algorithm is again quite simple:

(1) Compute q_i = probability that X_i is the maximum.
(2) Discard X_i with probability $1 - \frac{2}{2+(e-2)q_i}$.
(3) Set threshold $\tau_i = F_i^{-1}(1 - q_i)$.
(4) Keep first random variable whose realization is at least τ_i.

In the general situation we apply an arbitrary tie-breaking rule so that $\sum q_i = 1$.

Similarly in Corollary 4.7 (which follows from Theorem 1.3) we prove that a variant of Algorithm 2 gives a sequence of thresholds $\tau_1, \ldots, \tau_n$ such that, if we take the first of n i.i.d. random variables whose value is above the threshold, we obtain a value of at least a 0.745 fraction of the expectation of the maximum of the random variables. This result can be seen as a follow up on a result by Hill and Kertz [13] on the prophet inequality for i.i.d. random variables. They study the performance of the best stopping time when compared to a prophet that can extract the expectation of the maximum. The main result of Hill and Kertz is a recursive characterization of a_n, the best possible factor when faced with n random variables. More precisely, they prove that if $X_1, \ldots, X_n$ are i.i.d. nonnegative random variables and T_n denotes the set of stopping rules for $X_1, \ldots, X_n$ then

$$\mathbb{E}(\max\{X_1, \ldots, X_n\}) \leq a_n \sup\{\mathbb{E}(X_t) : t \in T_n\} .$$

Furthermore, Hill and Kertz find instances in which it is not possible to beat the factor a_n. They also prove that $a_n \leq e/(e-1)$, conjecture that the sequence is monotone, and leave open the existence and computation of its limit. The monotonicity together with the limit calculation would readily give a universal bound (valid for all n) on the performance of the best stopping rule. Shortly after, Samuel-Cahn [21] reports that Kertz proves existence of the limit a of the a_n sequence and conjectures that it equals 1.342 (obtained as the solution to $\int_0^1 (y - y\ln(y)) + a - 1)^{-1} dy = 1$). Finally, Kertz [14, Lemma 6.2] proves the latter conjecture (for which Saint-Mont [20] derives a simpler proof). However he is unable to prove that the sequence is monotone and therefore the best upper bound on the whole a_n sequence still stood at $e/(e-1) \approx 1.582$ [14, Lemma 3.4]. Very recently, and independently of our work, Abolhassani et al. [1] improved this upper bound to $1/0.738 \approx 1.355$. Our Corollary 4.7 closes this gap and implies that for all n, $a_n \leq a \approx 1.342$, and by the tight examples of Hill and Kertz [13] it turns out that this constant is best possible.

2 THE BERNOULLI SELECTION LEMMA

In this section we prove Lemma 1.2. Actually, we prove a slightly stronger version that will become clear at the end of the proof. We also provide a tight instance and discuss some generalizations.

The proof. The second inequality of Lemma 1.2 is trivial, as the expectation of the maximum is a sum over all values b_i weighed by the probability with which that value is the maximum. Since these probabilities sum to at most one, the inequality follows.

The proof for the first inequality has two main ingredients. First, we reformulate the left hand side in an appropriate way, and lower bound it by another function using KKT-conditions. Then, we show that this function is bounded from below by $1 - 1/e$.

As a warm up, we first show how to get a weaker result, that only gives us a factor of $\sqrt{e}$ instead of $\frac{e}{e-1}$, with more straightforward arguments.

Proof. We start the proof by rewriting the optimization problem:

$$\max_{S \subseteq N} \left\{ \mathbb{E}\left[\frac{\sum_{i \in S} b_i X_i}{\sum_{i \in S} X_i} \right] \right\} . \tag{P}$$

Instead of choosing a subset of N, we set for each $i \in N$ a value $\chi_i \in [0, 1]$, which represents the probability with which we actually choose i. Now, let $\pi_i = \chi_i q_i$ denote the probability of i being picked and having $X_i = 1$. So we can consider the following maximization problem, with decision variables π, as a relaxation of (P):

$$\max_{0 \le \pi_i \le q_i} \sum_{S \subseteq N} \left(\frac{b(S)}{|S|} \left(\prod_{i \in S} \pi_i \right) \left(\prod_{i \notin S} (1 - \pi_i) \right) \right) ,$$

where $b(S) = \sum_{i \in S} b_i$. Note that the previous objective is linear in each variable so that there is an extreme optimal solution [8]. Thus, the previous problem is in fact equivalent to (P). Now, by changing the order of the summations, we obtain

$$\max_{0 \le \pi_i \le q_i} \sum_{i \in N} b_i \pi_i \sum_{S \subseteq N \setminus \{i\}} \frac{1}{1 + |S|} \prod_{j \in S} \pi_j \prod_{j \in N \setminus (S \cup \{i\})} (1 - \pi_j) . \tag{1}$$

With (P) in this equivalent form, we now proceed to guess a feasible solution. To this end, consider an optimal solution z^* to

$$\max \left\{ \sum_{i \in N} b_i z_i \;\middle|\; \sum_{i \in N} z_i \le 1, z_i \le q_i \text{ for all } i \in N \right\} ,$$

and set $\pi_i = 2z_i^*/(2 + z_i^*)$. Note that $\pi_i \le q_i$, so that, substituting this in the objective of (1), we get

$$\sum_{i \in N} b_i z_i^* \prod_{j \in N} \frac{1}{1 + \frac{z_j^*}{2}} \sum_{S \subseteq N \setminus \{i\}} \frac{2^{|S|}}{1 + |S|} \prod_{j \in S} \frac{z_j^*}{2} \prod_{j \in N \setminus (S \cup \{i\})} \left(1 - \frac{z_j^*}{2} \right) . \tag{2}$$

It is easy to see that

$$\sum_{S \subseteq N \setminus \{i\}} \frac{2^{|S|}}{1 + |S|} \prod_{j \in S} \frac{z_j^*}{2} \prod_{j \in N \setminus (S \cup \{i\})} \left(1 - \frac{z_j^*}{2} \right) \ge 1 ,$$

since the left hand side corresponds to $\mathbb{E}[f(S)]$ over all $S \subseteq N \setminus \{i\}$ under probabilities $z_j^*/2$ for every element i and $f(S) = 2^{|S|}/(|S| + 1) \ge 1$. While for any values z_i such that $\sum_i z_i \le 1$, we have

$$\prod_{j=1}^{n} \frac{1}{1 + \frac{z_j}{2}} \ge e^{-\sum_{j=1}^{n} \frac{z_j}{2}} \ge \frac{1}{\sqrt{e}} ,$$

where the first inequality follows from $1 + x \le e^x$, concluding the proof. □

To obtain the improved factor $1 - 1/e$ we do the same as in the last proof, but make a subtle modification in the choice of π_i. We choose $\pi_i = \frac{2z_i^*}{2+(e-2)z_i^*}$,[3][4] such that $1 - \pi_i = \frac{2-(4-e)z_i^*}{2+(e-2)z_i^*}$. Note that this is a feasible choice of π_i for all $i \in N$, since for this choice $\pi_i \leq z_i^* \leq q_i$. We plug this back into (1), and obtain that (P) is lower bounded by

$$\sum_{i \in N} 2b_i z_i^* \left(\prod_{j \in N} \frac{1}{2 + (e-2)z_j^*} \right) \sum_{S \subseteq N \setminus \{i\}} \frac{2^{|S|}}{1 + |S|} \prod_{j \in S} z_j^* \prod_{j \in N \setminus (S \cup \{i\})} (2 - (4-e)z_j^*) \,. \tag{3}$$

We proceed to lower bound this quantity, where we use the following technical result.

Proposition 2.1. *Consider the problem* $\min_{x \in \mathbb{R}_+^M} \{ f_M(x) : \sum_{i \in M} x_i \leq a \}$, *where* $a \leq 1$ *and*

$$f_M(x) = \left(\prod_{j \in M} \frac{1}{2 + (e-2)x_j} \right) \sum_{S \subseteq M} \frac{2^{|S|}}{1 + |S|} \prod_{j \in S} x_j \prod_{j \in M \setminus S} (2 - (4-e)x_j) \,.$$

An optimal solution satisfies that all nonzero variables have to be equal and $\sum_{i \in M} x_i = a$.

Using Proposition 2.1, we lower bound (3) as follows. Consider the term

$$\left(\prod_{j \in N \setminus \{i\}} \frac{1}{2 + (e-2)z_j^*} \right) \sum_{S \subseteq N \setminus \{i\}} \frac{2^{|S|}}{1 + |S|} \prod_{j \in S} z_j^* \prod_{j \in N \setminus (S \cup \{i\})} (2 - (4-e)z_j^*) \,.$$

Note that this is equal to $f_{N \setminus \{i\}}(z_{-i}^*)$.[5] So, Proposition 2.1 can be applied with $a = 1 - z_i^*$. Thus,

$$f_{N \setminus \{i\}}(z_{-i}^*) \geq f_{N \setminus \{i\}}(x^*) \,,$$

with x^* the optimal solution to $\min_{x \in \mathbb{R}_+^{N \setminus \{i\}}} \{ f_{N \setminus \{i\}}(x) : \sum_{j \in N \setminus \{i\}} x_j \leq a \}$.

Proposition 2.1 states that $x_j^* = (1 - z_i^*)/k$, where $k \leq n - 1$ is the number of nonzero variables in x^*. Conditioning on the cardinality of the set S, and using the Binomial Theorem, a straightforward but tedious calculation shows that

$$f_{N \setminus \{i\}}(x^*) = \frac{2k + (e-2)(1 - z_i^*)}{2(k+1)(1 - z_i^*)} \left(1 - \left(1 - \frac{2(1 - z_i^*)}{2k + (e-2)(1 - z_i^*)} \right)^{k+1} \right) \,.$$

As this quantity only depends on k and z_i^*, we may define

$$\varphi_k(z_i^*) = \frac{2}{(2 + (e-2)z_i^*)} f_{N \setminus \{i\}}(x^*) \,,$$

to conclude that expression (3) (and in turn (P)) is lower bounded by

$$\sum_{i \in N} b_i z_i^* \varphi_{k(i)}(z_i^*) \,.$$

where the index $k(i)$, denoting the number of nonzero variables in x^*, is always at least 1, yet may vary, depending on i.

The remainder of the proof establishes that $\varphi_{k(i)}(z_i^*) \geq 1 - \frac{1}{e}$. Indeed, we show that, for all $y \in [0, 1]$ and for all $n \geq 1$, we have that $\varphi_n(y) \geq 1 - \frac{1}{e}$. □

[3] Because of the choice of π_i, we actually prove the slightly stronger bound where we maximize over $z_i \leq \frac{2q_i}{2-(e-2)q_i}$.

[4] The choice of π_i suggests that the random variables are not picked deterministically, but with probability less than 1, since $\pi_i < z_i^*$ if $z_i^* > 0$. However, as noted in the beginning of the proof, because of linearity of the objective in each variable, there is always an extreme optimal solution where the random variables are picked deterministically.

[5] x_{-i} denotes the vector x with coordinate i eliminated.

Tightness. We now provide a family of instances that show that the $1-1/e$ bound in the Bernoulli Selection Lemma is actually best possible. Consider n^2 independent identically distributed Bernoulli random variables with parameter $1/n$ and prizes $b_1 = n/(e-2)$ and $b_i = 1$ for $2 \le i \le n^2$. The expectation of the maximum prize is given by

$$\mathbb{E}\left[\max_{i\in N}\{b_i X_i\}\right] = \frac{1}{e-2} + \left(1-\frac{1}{n}\right)\left(1-\left(1-\frac{1}{n}\right)^{n^2-1}\right) \longrightarrow \frac{1}{e-2}+1\,.$$

In this particular setting, where the Bernoulli random variables are i.i.d., the best strategy is to sort by prize and take some subset with those of higher prize. This means to choose the first random variable and a subset of size $k-1$ of the rest for some $1 \le k \le n^2$. This yields an expected average value that can be upper bounded by

$$\left(1-\left(1-\frac{1}{n}\right)^k\right)\frac{\frac{n}{e-2}+k-1}{k} \le \left(1-\left(1-\frac{1}{n}\right)^k\right)\left(\frac{n}{k(e-2)}+1\right).$$

The above can be shown to converge, as $n \longrightarrow \infty$, to

$$\max_{0\le x\le n}(1-e^{-x})\left(\frac{1}{x(e-2)}+1\right),$$

where $x = \frac{k}{n}$. Moreover, this expression is maximized at $x = 1$. This yields the value $(1 - e^{-1})\left(\frac{1}{e-2}+1\right) = (1-1/e)\mathbb{E}[\max_{i\in N}\{b_i X_i\}]$.

Extension. The Bernoulli Selection Lemma can be used to prove a similar result for more general random variables. Suppose you can now choose one of n prizes whose values are random variables distributed according to n possibly different distributions. The prizes arrive in random order, and, upon arrival, we must decide whether we keep that prize, or we simply discard it and wait for the next. The goal is to maximize the expected value of the selected prize. Similar to the nonadaptive setting considered so far, we look for an acceptance criterion that is set beforehand, and only based on the distributions. This problem resembles the prophet inequality and also the prophet secretary problem [10].

Corollary 2.2. *Given n independent nonnegative random variables $X_1, \dots, X_n$ with $X_i \sim F_i$. There exist values $\tau_1, \dots, \tau_n$ such that*

$$\mathbb{E}\left[\frac{\sum_{i=1}^n X_i Y_i}{\sum_{i=1}^n Y_i}\right] \ge \left(1-\frac{1}{e}\right)\mathbb{E}\left[\max_{i=1,\dots,n} X_i\right],$$

where Y_i is a Bernoulli random variable that has value 1 if $X_i > \tau_i$.

In the context of the prophet inequality, note that the quantity on the left exactly corresponds to the expected value of the first X_i above τ_i, when the X_i's are ordered uniformly at random.

Proof. Assume first that the F_i are continuous for all i. Let $q_i = \mathbb{P}(X_i \ge X_j,\ \forall j = 1,\dots,n)$ be the probability that X_i is the largest and α_i be a value for which $1 - F_i(\alpha_i) = q_i$. Consider $b_i = \mathbb{E}[X_i \mid X_i > \alpha_i]$ and the Bernoulli random variables $Z_1, \dots, Z_n$ where Z_i has parameter q_i. We apply the Bernoulli Selection Lemma to this instance, and thus let $S \subseteq \{1,\dots,n\}$ be a set for which the lemma holds. Now define $\tau_i = \alpha_i$ for $i \in S$ and $\tau_i = \infty$ otherwise, and note that for $i \notin S$, we have $Y_i = 0$ almost surely, and for $i \in S$, we have $\mathbb{P}(X_i > \alpha_i) = \mathbb{P}(Y_i = 1) = q_i$. It follows that

$$\begin{aligned}
\mathbb{E}\left[\frac{\sum_{i=1}^n X_i Y_i}{\sum_{i=1}^n Y_i}\right] &= \sum_{i\in S}\mathbb{E}\left[\frac{X_i Y_i}{\sum_{j\in S} Y_j}\right] \\
&= \sum_{i\in S}\mathbb{E}[X_i \mid Y_i = 1]\mathbb{E}\left[(1+\sum_{j\in S\setminus\{i\}} Y_j)^{-1}\middle| Y_i = 1\right]\mathbb{P}(Y_i = 1) \\
&= \sum_{i\in S}\mathbb{E}[X_i \mid X_i > \alpha_i]\mathbb{E}\left[\frac{Y_i}{\sum_{j\in S} Y_j}\right] \\
&= \mathbb{E}\left[\frac{\sum_{i\in S}\mathbb{E}[X_i|X_i>\alpha_i]Z_i}{\sum_{i\in S} Z_i}\right] \\
&\geq \frac{e-1}{e}\max_{z_i\leq q_i}\left\{\sum_{i=1}^n \mathbb{E}[X_i \mid X_i > \alpha_i]z_i \,\middle|\, \sum_{i=1}^n z_i \leq 1\right\} \\
&\geq \frac{e-1}{e}\sum_{i=1}^n \mathbb{E}[X_i \mid X_i > \alpha_i]q_i \,,
\end{aligned}$$

where the second to last inequality follows from the Bernoulli Selection Lemma, while the last holds since $\sum_{i=1}^n q_i = 1$. Now note that $\mathbb{E}\left[\max_{i=1,\ldots,n} X_i\right] = \sum_{i=1}^n \mathbb{E}[X_i \mid X_i \geq X_j,\ \forall j = 1,\ldots,n]q_i$. To finish the proof, it suffices to show that

$$\mathbb{E}[X_i \mid X_i > \alpha_i] \geq \mathbb{E}[X_i \mid X_i \geq X_j,\ \forall j = 1,\ldots,n]\,.$$

Indeed, for $x > \alpha_i$, we have $\mathbb{P}(X_i > x \mid X_i > \alpha_i) = \int_x^\infty \frac{1}{q_i}dF_i(t)$, while, if $x \leq \alpha_i$, the previous probability equals 1. On the other hand,

$$\mathbb{P}(X_i > x \mid X_i \geq X_j\ \forall j = 1,\ldots,n) = \int_x^\infty \frac{\prod_{j\neq i} F_j(t)}{q_i}dF_i(t)\,.$$

From this, it follows that $\mathbb{P}(X_i > x \mid X_i > \alpha_i) \geq \mathbb{P}(X_i > x \mid X_i \geq X_j,\ \forall j = 1,\ldots,n)$ for all $x \geq 0$. Thus, $X_i \mid (X_i > \alpha_i)$ stochastically dominates $X_i \mid (X_i \geq X_j\ \forall j = 1,\ldots,n)$, and the conclusion follows.

When some F_i are not continuous, it could be the case that there is no α_i such that $1 - F_i(\alpha_i) = q_i$ or that $\sum q_i > 1$. If the former happens, the result still holds provided α_i is chosen randomly. The latter case is solved by slightly perturbing the support of the random variables in a way that the probability that two or more are the maximum simultaniously is negligible. □

3 NONADAPTIVE POSTED PRICE MECHANISMS

In this section we prove our main result, namely, Theorem 1.1. Recall that we have a single item on sale, a set of customers $\mathcal{I}$, and for customer $i \in \mathcal{I}$ her valuation for the item is $v_i \sim F_i$. As is standard in the literature, we say that a distribution F_i is *regular* if the *virtual value function* $c_i(v) = v - (1 - F_i(v))/f_i(v)$ is nondecreasing, where f_i is the density of F_i.

Besides the Bernoulli Selection Lemma, key to our analysis is the by now classic result of Chawla et al. [6].

Lemma 3.1 ([6, Lemma 4]). *If all value distributions are regular, then the expected value of Myerson's optimal auction is bounded from above by*

$$\sum_{i\in\mathcal{I}} F_i^{-1}(1 - q_i^M)q_i^M\,,$$

where q_i^M is the probability that the optimal auction assigns the item to i.

Furthermore, for every i (with regular or nonregular value distribution) there exist two prices $\underline{p_i}$ and $\overline{p_i}$, with corresponding probabilities $\underline{q_i}$ and $\overline{q_i}$, and a number $0 \leq x_i \leq 1$, such that $x_i\underline{q_i} + (1-x_i)\overline{q_i} = q_i^M$, and the expected revenue of Myerson's optimal auction is bounded from above by

$$\sum_{i\in\mathcal{I}} x_i\underline{p_i}\,\underline{q_i} + (1 - x_i)\overline{p_i}\,\overline{q_i}\,.$$

Theorem 1.1. *For any given set of potential customers $\mathcal{I}$, there exists a nonadaptive posted price mechanism that achieves an expected revenue of at least a $1 - 1/e$ fraction of that of Myerson's optimal auction on $\mathcal{I}$.*

Proof. We prove the regular case first. Let q_i^M denote the probability with which Myerson's optimal auction assigns the item to customer $i \in \mathcal{I}$, and set $b_i = F_i^{-1}(1 - q_i^M)$. The expected revenue of a nonadaptive posted price mechanism, that chooses to sell only to customers in $S \subseteq \mathcal{I}$ while offering prices b_i, is given by

$$\begin{aligned}\sum_{i \in S} b_i \mathbb{P}[i = \operatorname*{argmin}_{j \in S}\{\sigma(j) \mid v_j \geq b_j\}] &= \sum_{i \in S} b_i q_i^M \mathbb{P}[i = \operatorname*{argmin}_{j \in S}\{\sigma(j) \mid v_j \geq b_j\} \mid v_i \geq b_i] \\ &= \sum_{i \in S} b_i q_i^M \sum_{R \subseteq S \setminus \{i\}} \frac{1}{1 + |R|} \prod_{j \in R} q_j^M \prod_{j \in S \setminus (R \cup \{i\})} (1 - q_j^M) \\ &= \sum_{i \in S} b_i q_i^M \mathbb{E}\left[\tfrac{1}{1 + \sum_{j \in S \setminus \{i\}} X_j}\right] = \mathbb{E}\left[\tfrac{\sum_{i \in S} b_i X_i}{\sum_{i \in S} X_i}\right],\end{aligned}$$

where $\{X_i\}_{i \in \mathcal{I}}$ are Bernoulli random variables with $X_i = 1$ with probability q_i^M. By the Bernoulli Selection Lemma we can choose the set $S \subseteq \mathcal{I}$ to be such that the latter is lower bounded by

$$\left(1 - \frac{1}{e}\right) \max_{z_i \leq q_i^M} \left\{\sum_{i \in \mathcal{I}} b_i z_i \,\middle|\, \sum z_i \leq 1\right\} \geq \left(1 - \frac{1}{e}\right) \sum_{i \in \mathcal{I}} F_i^{-1}(1 - q_i^M) q_i^M .$$

Therefore, Lemma 3.1 leads to the desired conclusion.

In the nonregular case, the posted price mechanism runs a lottery between two prices to get the desired bound.[6] First, for every bidder with positive probability of winning the optimal auction, set

$$b_i' = \frac{x_i \underline{p_i}\, \underline{q_i} + (1 - x_i)\overline{p_i}\, \overline{q_i}}{q_i^M},$$

where the variables are defined as in the lemma. Also consider the same Bernoulli random variables presented in the first part of the proof. The nonadaptive posted price mechanism sells only to a set S' of customers (to be defined). For every $i \in S'$, it offers a random price p_i equal to $\underline{p_i}$ with probability x_i, and $\overline{p}_i$ otherwise. This way, the a priori probability that v_i is above the price offered is exactly $x_i \underline{q_i} + (1 - x_i)\overline{q_i} = q_i^M$, while the expected revenue of the mechanism can be evaluated as

$$\begin{aligned}&\sum_{i \in S'} x_i \underline{p_i}\, \underline{q_i} \mathbb{P}[i = \operatorname*{argmin}_{j \in S'}\{\sigma(j) \mid v_j \geq p_j\} \mid v_i \geq p_i,\ p_i = \underline{p_i}] \\ &\quad + (1 - x_i)\overline{p_i}\, \overline{q_i} \mathbb{P}[i = \operatorname*{argmin}_{j \in S'}\{\sigma(j) \mid v_j \geq p_j\} \mid v_i \geq p_i,\ p_i = \overline{p_i}] \\ &\quad = \sum_{i \in S'} (x_i \underline{p_i}\, \underline{q_i} + (1 - x_i)\overline{p_i}\, \overline{q_i}) \sum_{R \subseteq S' \setminus \{i\}} \frac{1}{1 + |R|} \prod_{j \in R} q_j^M \prod_{j \in S' \setminus (R \cup \{i\})} (1 - q_j^M) \\ &\quad = \sum_{i \in S'} b_i' q_i^M \mathbb{E}\left[\tfrac{1}{1 + \sum_{j \in S' \setminus \{i\}} X_j}\right] = \mathbb{E}\left[\tfrac{\sum_{i \in S'} b_i' X_i}{\sum_{i \in S'} X_i}\right].\end{aligned}$$

By the same argument as before, Lemma 1.2 implies that there exists $S' \subseteq \mathcal{I}$ such that the latter is lower bounded by $(1 - 1/e) \sum_{i \in \mathcal{I}} b_i' q_i^M$. Lemma 3.1 implies the bound over the optimal auction. □

[6] This lottery can be derandomized using standard techniques, since each combination of prices offered to the customers is a deterministic mechanism in itself and the random mechanism is simply a lottery over, and thus a convex combination of, those deterministic mechanisms.

Tight instance with i.i.d. valuations. We construct a family of instances for Problem 1 with i.i.d. customer valuations, such that, for all $\varepsilon > 0$, there is an instance from this family for which no nonadaptive strategy can achieve an expected revenue within a factor $(1+\varepsilon)(1-1/e)$ of the optimal expected revenue. The idea is to mimic the instance that makes the Bernoulli Selection Lemma tight, but here we achieve this with i.i.d. valuations. Consider n^2 customers whose values are independent identically distributed according to

$$V = \begin{cases} \frac{n}{e-2} & \text{w.p. } \frac{1}{n^3}\,, \\ 1 & \text{w.p. } \frac{1}{n}\,, \\ 0 & \text{w.p. } 1 - \frac{1}{n} - \frac{1}{n^3}\,. \end{cases}$$

Then, it is easy to design an auction that achieves a revenue approaching $(e-1)/(e-2)$ as $n \to \infty$. Indeed, consider the auction that offers the item for price $n/(e-2)-c$ (with c a small value, say $c = 2$) to any bid above that price (and assigns the item at random if more than one such offer is received), and if no such bid is received, then it runs a lottery at price 1 among all the bids above that price. As there are many buyers of value 1, a potential large value customer will prefer to make a revenue of c rather than risking to lose the item in the lottery. Therefore the revenue the auction will generate will approach $1/(e-2)+1$ as $n \to \infty$. Of course, the revenue of the optimal auction is then at least this quantity. On the other hand, the best posted price mechanism offers a price of 1 to, say, customers $1, \ldots, k$ and $n/(e-2)$ to the rest of the customers, for some well chosen value of k which turns out to be roughly n. One can show that in the limit the revenue approaches $(\frac{1}{e-2}+1)(1-e^{-1})$.

4 ADAPTIVE POSTED PRICE MECHANISM

In the previous section we considered the setting in which the posted price only depends on the customer, not on the order. In this section we consider the setting in which the posted price may depend both on the customer and on the customers that arrived before her. We design an adaptive posted price mechanism that achieves an expected revenue of at least a 0.745 fraction of that of the optimal mechanism. In particular we prove the following.

Theorem 1.3. *For any given set of potential customers $\mathcal{I}$ whose values are independent and identically distributed, there exists an adaptive posted price mechanism that achieves an expected revenue of at least a $1/\beta > 0.745$ fraction of that of Myerson's optimal auction on $\mathcal{I}$, where β is the unique value such that*

$$\int_0^1 \frac{1}{y(1-\ln(y)) + (\beta - 1)}\,dy = 1\,.$$

We assume that the valuations of the customers are i.i.d. with cumulative distribution function $F(\cdot)$ and probability density function $f(\cdot)$. For our analysis the bound provided by Lemma 3.1 is not enough, so we derive an exact expression for the expected revenue of the optimal auction.

Expected value of Myerson's optimal auction for i.i.d. customers. Following Myerson [19], we define the *virtual valuation* as $c(v) = v - \frac{1-F(v)}{f(v)}$ and the *ironed virtual valuation* as $\bar{c}(v) = G'(F(v))$, where $G = \mathrm{conv}(H)$[7] is the convexification of the negative revenue curve $H(q) = \int_0^q c(F^{-1}(\theta))d\theta$ as a function of the acceptance probability q. Also, let $\mathbb{E}(\mathrm{MY}(n,F))$ be the expected revenue of the optimal auction over n customers with values drawn from distribution F.

The expected profit of the optimal auction equals its expected virtual surplus (see, e.g., [12]), i.e., the sum over all customers of the expected value of the maximum of $\bar{c}$ above zero. Note that $\bar{c}$ is

[7] $G(q) = \min\{\gamma H(q_1) + (1-\gamma)H(q_2) : \gamma q_1 + (1-\gamma)q_2 = q,\ \gamma, q_1, q_2 \in [0,1]\}$. Note that if $c(v)$ is monotone (also known as *regular*), then $G(q) = H(q)$.

an increasing function, and let v^* be the value at which $\bar{c}(v^*) = 0$ or zero, if no such value exists. Then, the latter can be evaluated as:

$$\mathbb{E}(\mathrm{MY}(n, F)) = \int_{v^*}^{\infty} nF(v)^{n-1}\bar{c}(v)f(v)dv.$$

Performing the change of variables $q = 1 - F(v)$ and $\alpha^* = 1 - F(v^*)$, we obtain

$$\begin{aligned}
\mathbb{E}(\mathrm{MY}(n, F)) &= n\int_0^{\alpha^*} (1-q)^{n-1}\bar{c}(F^{-1}(1-q))dq \\
&= n\int_0^{\alpha^*} (1-q)^{n-1}G'(1-q)dq \\
&= -nG(1-q)(1-q)^{n-1}\Big|_0^{\alpha^*} - \int_0^{\alpha^*} n(n-1)(1-q)^{n-2}G(1-q)dq \\
&= nG(1) - nG(F(v^*))F(v^*)^{n-1} - n(n-1)\int_0^{\alpha^*} (1-q)^{n-2}G(1-q)dq\,.
\end{aligned}$$

Since $\bar{c}(v^*) = 0$, we know that G attains a minimum at $F(v^*)$ and, therefore, equals $H(F(v^*))$ at that point. Now, observe that

$$H(q) = \int_0^{q} F^{-1}(\theta) - \frac{1-\theta}{f(F^{-1}(\theta))}d\theta = -(1-q)F^{-1}(q)\,.$$

Therefore, we can conclude that

$$\begin{aligned}
\mathbb{E}(\mathrm{MY}(n, F)) &= -nH(F(v^*))F(v^*)^{n-1} - n(n-1)\int_0^{\alpha^*} (1-q)^{n-2}G(1-q)dq \\
&= nv^*(1-F(v^*))F(v^*)^{n-1} - n(n-1)\int_0^{\alpha^*} (1-q)^{n-2}G(1-q)dq\,.
\end{aligned}$$

Now, let

$$\bar{G}(1-q) = \begin{cases} -G(1-q) & \text{if } 1-q > F(v^*)\,, \\ v^*(1-F(v^*)) & \text{otherwise}\,. \end{cases}$$

Then, we can write the expected revenue of the optimal mechanism as

$$\mathbb{E}(\mathrm{MY}(n, F)) = n(n-1)\int_0^{1} (1-q)^{n-2}\bar{G}(1-q)dq\,. \tag{4}$$

We note that expression (4), although fairly natural to derive, appears to be new.

Adaptive posted price mechanism. In the adaptive setting, the price offered to a customer also depends on the set of customers that declined the offer. However, as the customers are i.i.d., an adaptive pricing mechanism only needs to know how many customers have received an offer and not exactly which customers. The intuition behind the mechanism is to start with a high price and decrease the price when the number of remaining customer decreases, such that the risk of not selling is mitigated. To describe the mechanism, we first restrict to monotone virtual valuations, and then describe how to deal with the general case.

We partition the interval $A = [0, 1]$ into n intervals $A_i = [\varepsilon_{i-1}, \varepsilon_i]$, with $0 = \varepsilon_0 < \varepsilon_1 < \ldots < \varepsilon_{n-1} < \varepsilon_n = 1$. The pricing mechanism will thus choose a price for customer i such that the probability that this customer accepts the offer lies in A_i, making sure that no customer will ever receive an offer lower than the reservation price v^*. To implement this idea we use Expression (4) as a guide and construct the offer for customer i by drawing a q_i from the interval A_i according to probability density function $f_i(q) = \frac{\psi(q)}{\alpha_i}$,[8] where $\psi(q) = (n-1)(1-q)^{n-2}$ and α_i is a normalization parameter equal to $\alpha_i = \int_{q \in A_i} \psi(q)dq$. This q_i is meant to be the acceptance probability of customer i, so the price offered to her is $\max\{F^{-1}(1-q_i), v^*\}$.

Note that due to the assumption that the virtual valuations are monotone, we know that when offering a price $F^{-1}(1-q_i)$, the expected profit of customer i is $q_i F^{-1}(1-q_i) = -H(1-q_i) = \bar{G}(1-q_i)$. For nonmonotone virtual valuations, it might be the case that $q_i F^{-1}(1-q_i) = -H(1-q_i) < \bar{G}(1-q_i)$, and by offering a price $F^{-1}(1-q_i)$ we might not get the best revenue. To circumvent this problem, we can randomize between two acceptance probabilities q_{i1} and q_{i2} such that $G(1-q_i) = \gamma H(1-q_{i1}) + (1-\gamma)H(1-q_{i2})$ and $q_i = \gamma q_{i1} + (1-\gamma)q_{i2}$.

With the mechanism in place we now establish the approximation guarantee. To this end we prove that the expected revenue of the adaptive posted price mechanism with n customers whose valuations are drawn from F, $\mathbb{E}(\text{ADAP}(n,F))$, satisfies

$$\mathbb{E}(\text{ADAP}(n,F)) \geq \sum_{i=1}^{n} \rho_i \int_{\varepsilon_{i-1}}^{\varepsilon_i} (n-1)(1-q)^{n-2}\bar{G}(1-q)dq\,,$$

where $\rho_1 = \frac{1}{\alpha_1}$ and $\rho_{i+1} = \frac{\rho_i}{\alpha_{i+1}} \int_{\varepsilon_{i-1}}^{\varepsilon_i} \psi(q)(1-q)dq$ for $i = 1, \ldots, n-1$.

By choosing $\varepsilon_1, \ldots, \varepsilon_{n-1}$ in such a way that $\rho_1 = \rho_2 = \ldots = \rho_n$, we have $\mathbb{E}(\text{ADAP}(n,F)) \geq \frac{1}{n\alpha_1}\mathbb{E}(\text{MY}(n,F))$, and thus we wrap-up by proving that the term $1/(n\alpha_1)$ is bounded by 0.745. For the latter we set up a recursion whose solution determines α_1 and then approximate the recursion with an ordinary differential equation.

Relating the optimal auction and the adaptive posted price mechanism.

Lemma 4.1. *Let $\rho_1 = \frac{1}{\alpha_1}$ and $\rho_{i+1} = \frac{\rho_i}{\alpha_{i+1}} \int_{\varepsilon_{i-1}}^{\varepsilon_i} \psi(q)(1-q)dq$ for $i = 1, \ldots, n-1$. Then the expected value of the adaptive posted price mechanism that faces n customers with value distribution F, $\mathbb{E}(ADAP(n,F))$, satisfies*

$$\mathbb{E}(ADAP(n,F)) \geq \sum_{i=1}^{n} \rho_i \int_{\varepsilon_{i-1}}^{\varepsilon_i} (n-1)(1-q)^{n-2}\bar{G}(1-q)dq\,.$$

Proof. First, assume that the item is offered to customer i. Let q_i denote the drawn acceptance probability for customer i. The expected revenue obtained from selling the item to customer i is $\bar{G}(1-q_i)$. To see this, suppose that $q_i < 1-F(v^*)$. Then, for monotone virtual valuations, the price offered to customer i is $F^{-1}(1-q_i)$, and thus the expected revenue is $q_i F^{-1}(1-q_i) = -G(1-q_i) = \bar{G}(1-q_i)$. On the other hand, if $q_i > 1-F(v^*)$, the price offered to customer i is v^* which is accepted with probability $1-F(v^*)$. Similar arguments hold when the virtual valuation is not monotone.

Now, for $j = 1, \ldots, n$, let q_j denote the drawn acceptance probability of customer j. Then, the probability that the item is offered to customer i is equal to $\Pi_{j=1}^{i-1} \max\{1-q_j, F(v^*)\}$. Hence, the

[8] Just as for the non-adaptive mechanism, this randomization can be derandomized using standard techniques, since each combination of set prices for the customers is in itself a deterministic mechanism and the random mechanism is simply a lottery over, and thus a convex combination of, those deterministic mechanisms.

expected value of our adaptive posted price mechanism is

$$\begin{aligned}
\mathbb{E}(\mathrm{ADAP}(n,F)) \quad &= \int_0^{\varepsilon_1} \frac{\psi(q_1)}{\alpha_1}\Bigg[\bar{G}(1-q_1) + \max\{1-q_1, F(v^*)\} \int_{\varepsilon_1}^{\varepsilon_2} \frac{\psi(q_2)}{\alpha_2}\Bigg[\bar{G}(1-q_2) \\
&+ \max\{1-q_2, F(v^*)\} \int_{\varepsilon_2}^{\varepsilon_3} \frac{\psi(q_3)}{\alpha_3}\Bigg[\bar{G}(1-q_3) + \dots \\
&+ \max\{1-q_{n-2}, F(v^*)\} \int_{\varepsilon_{n-2}}^{\varepsilon_{n-1}} \frac{\psi(q_{n-1})}{\alpha_{n-1}}\Bigg[\bar{G}(1-q_{n-1}) \\
&+ \max\{1-q_{n-1}, F(v^*)\} \int_{\varepsilon_{n-1}}^{1} \frac{\psi(q_n)}{\alpha_n}\bar{G}(1-q_n)\,dq_n\Bigg]\,dq_{n-1}\dots\Bigg]\,dq_3\Bigg]\,dq_2\Bigg]\,dq_1 \\
&\geq \int_0^{\varepsilon_1} \frac{\psi(q_1)}{\alpha_1}\Bigg[\bar{G}(1-q_1) + (1-q_1) \int_{\varepsilon_1}^{\varepsilon_2} \frac{\psi(q_2)}{\alpha_2}\Bigg[\bar{G}(1-q_2) + \\
&+ (1-q_2) \int_{\varepsilon_2}^{\varepsilon_3} \frac{\psi(q_3)}{\alpha_3}\Bigg[\bar{G}(1-q_3) + \dots \\
&+ (1-q_{n-2}) \int_{\varepsilon_{n-2}}^{\varepsilon_{n-1}} \frac{\psi(q_{n-1})}{\alpha_{n-1}}\Bigg[\bar{G}(1-q_{n-1}) \\
&+ (1-q_{n-1}) \int_{\varepsilon_{n-1}}^{1} \frac{\psi(q_3)}{\alpha_3}\bar{G}(1-q_n)\,dq_n\Bigg]\,dq_{n-1}\dots\Bigg]\,dq_3\Bigg]\,dq_2\Bigg]\,dq_1
\end{aligned}$$

and the latter equals

$$\begin{aligned}
&\frac{1}{\alpha_1}\int_0^{\varepsilon_1} \psi(q_1)\bar{G}(1-q_1)\,dq_1 + \frac{1}{\alpha_1}\frac{\int_0^{\varepsilon_1}\psi(q_1)(1-q_1)\,dq_1}{\alpha_2}\int_{\varepsilon_1}^{\varepsilon_2} \psi(q_2)\bar{G}(1-q_2)\,dq_2 + \dots \\
&\quad + \frac{1}{\alpha_1}\frac{\int_0^{\varepsilon_1}\psi(q_1)(1-q_1)\,dq_1}{\alpha_2}\cdot\ldots\cdot\frac{\int_{\varepsilon_{n-2}}^{\varepsilon_{n-1}}\psi(q_{n-1})(1-q_{n-1})\,dq_{n-1}}{\alpha_n}\int_{\varepsilon_{n-1}}^{1}\psi(q_n)\bar{G}(1-q_n)\,dq_n \\
&= \sum_{i=1}^{n}\rho_i\int_{\varepsilon_{i-1}}^{\varepsilon_i}(n-1)(1-q)^{n-2}\bar{G}(1-q)dq\,. \qquad \square
\end{aligned}$$

By choosing $\varepsilon_1, \dots, \varepsilon_{n-1}$ appropriately, we can lower bound the expected revenue of the adaptive posted price mechanism against that of Myerson.

LEMMA 4.2. *If we choose $\varepsilon_1, \dots, \varepsilon_{n-1}$ such that $\rho_1 = \rho_2 = \ldots = \rho_n$, then*

$$\mathbb{E}(ADAP(n,F)) \geq \frac{1}{n\alpha_1}\mathbb{E}(MY(n,F))\,.$$

PROOF. If we choose $\varepsilon_1, \dots, \varepsilon_{n-1}$ such that $\rho_i = \rho_1$ for all i, then by Lemma 4.1 we can bound the expected value of our mechanism by

$$\begin{aligned}
\mathbb{E}(\mathrm{ADAP}(n,F)) &\geq \sum_{i=1}^{n}\rho_i\int_{\varepsilon_{i-1}}^{\varepsilon_i}(n-1)(1-q)^{n-2}\bar{G}(1-q)\,dq \\
&= \frac{\rho_1}{n}\sum_{i=1}^{n}\int_{\varepsilon_{i-1}}^{\varepsilon_i}n(n-1)(1-q)^{n-2}\bar{G}(1-q)\,dq \\
&= \frac{1}{n\alpha_1}\int_0^1 n(n-1)(1-q)^{n-2}\bar{G}(1-q)\,dq = \frac{1}{n\alpha_1}\mathbb{E}(\mathrm{MY}(n,F))\,. \qquad \square
\end{aligned}$$

Bounding α_1 through a recursion. Since $\rho_{i+1} = \frac{\rho_i}{\alpha_{i+1}} \int_{\varepsilon_{i-1}}^{\varepsilon_i} \psi(q)(1-q)dq$ for all i, choosing $\varepsilon_1, \dots, \varepsilon_{n-1}$ such that all ρ_i are the same amounts to choosing them such that $\int_{\varepsilon_{i-1}}^{\varepsilon_i} \psi(q)(1-q)dq = \alpha_{i+1}$ for all i. By definition of α_{i+1} and $\psi(q)$, this is equivalent to choosing them such that for all i

$$\int_{\varepsilon_{i-1}}^{\varepsilon_i} \psi(q)(1-q)\,dq = \frac{n-1}{n}\left((1-\varepsilon_{i-1})^n - (1-\varepsilon_i)^n\right)$$

is equal to

$$\int_{\varepsilon_i}^{\varepsilon_{i+1}} \psi(q)\,dq = (1-\varepsilon_i)^{n-1} - (1-\varepsilon_{i+1})^{n-1}\,.$$

Now, substitute $x_i = 1 - \varepsilon_i$. Then, from Lemma 4.2, we obtain the following recursion on x_i:

$$\frac{{x_{i-1}}^n}{n} - \frac{{x_i}^n}{n} = \frac{{x_i}^{n-1}}{n-1} - \frac{{x_{i+1}}^{n-1}}{n-1}\,, \tag{5}$$

where $x_0 = 1$ and $x_n = 0$. Moreover,

$$\alpha_1 = \int_0^{\varepsilon_1} \psi(q)\,dq = 1 - {x_1}^{n-1}\,.$$

Combining this with Lemma 4.2, we see that if $\frac{1}{n(1-{x_1}^{n-1})} \geq \frac{1}{\beta}$, for some $\beta \geq 1$, then the adaptive posted price mechanism has an expected revenue that is at least a $\frac{1}{\beta}$ fraction of the expected revenue of Myerson's optimal auction.

Note that $\frac{1}{n(1-{x_1}^{n-1})} \geq \frac{1}{\beta}$ if and only if $x_1 \geq (1-\frac{\beta}{n})^{1/(n-1)}$. Thus, if we find the minimum value β for which $x_1 < (1-\frac{\beta}{n})^{1/(n-1)} \implies x_n < 0$, we know that $x_1 \geq (1-\frac{\beta}{n})^{1/(n-1)}$ for that value of β. Hence, we show an upper bound on the sequence x_i. For this, we use the following lemma.

Lemma 4.3. *For values $x_0, x_1, \dots, x_n$ satisfying (5) and $x_0 = 1$ and $x_n = 0$, we have for $i = 1, \dots, n-1$,*

$$x_{i+1} = \left(\frac{n-1}{n}{x_i}^n + {x_1}^{n-1} - \frac{n-1}{n}\right)^{1/(n-1)}.$$

Proof. We prove this lemma by induction. For $i = 1$ equation (5) gives

$$x_2 = \left({x_1}^{n-1} + \frac{n-1}{n}{x_1}^n - \frac{n-1}{n}\right)^{1/(n-1)}.$$

Now, suppose the claim is true for $i = 1, \dots, j$. From (5), we know that

$$\begin{aligned}
{x_{j+1}}^{n-1} &= {x_j}^{n-1} + \frac{n-1}{n}{x_j}^n - \frac{n-1}{n}{x_{j-1}}^n \\
&= \frac{n-1}{n}{x_{j-1}}^n + {x_1}^{n-1} - \frac{n-1}{n} + \frac{n-1}{n}{x_j}^n - \frac{n-1}{n}{x_{j-1}}^n \\
&= \frac{n-1}{n}{x_j}^n + {x_1}^{n-1} - \frac{n-1}{n}\,,
\end{aligned}$$

where the second equality is due to the induction hypothesis. □

Bounding the recursion through a differential equation. In the following, we show that each of the terms x_i in the recursion can be bounded with a function $y(t) : [0,1] \to \mathbb{R}$, defined through the following differential equation. All derivatives of y are with respect to t.

$$\begin{aligned}
y' &= y(\ln(y) - 1) - (\beta - 1)\,, \\
y(0) &= 1\,.
\end{aligned} \tag{ODE}$$

Furthermore, we define $y(1) = \lim_{t\uparrow 1} y(t)$ as the continuous extension of $y(t)$. Later on, we will choose $\beta \approx 1.34$. For this β, we have $y \in [0,1]$, so we restrict our analysis of (ODE) to this interval. We assume $\beta > 1.25$ and $y \in [0,1]$. We validate these assumptions at the end of our analysis.

We now proceed to prove that the solution of (ODE) dominates the terms of the recurrence. In this proof, we make use of the following technical results.

Lemma 4.4. *Differential equation* (ODE) *has a unique solution* $y(t)$*, which is a decreasing and strictly convex function on the interval* $[0,1]$*. Furthermore,* $y'''(t) > 0$ *for* $y \in (0,1)$*.*

Proposition 4.5. *If* $x \in (0,1]$ *and* $n \geq 2$*, then*

$$x + \frac{x(\ln(x)-1)}{n} + \frac{x(\ln(x)-1)-(\beta-1)}{2n^2}\ln(x) \geq \frac{n-1}{n}x^{\frac{n}{n-1}}.$$

Using this inequality, we bound the recurrence by the function $y(t)$ in the following way.

Lemma 4.6. *If* $x_1 < (1-\frac{\beta}{n})^{\frac{1}{n-1}}$*, then* ${x_i}^{n-1} < y(\frac{i}{n})$ *for* $i = 1,\dots,n$*, where* $y(t)$ *is the unique solution of* (ODE)*.*

Proof. First note that $x_0 = y(0) = 1$, by definition. Moreover, a straightforward computation shows that $y'(0) = -\beta$. As $y(t)$ is strictly convex, we know that $y(1/n) > y(0) - \frac{1}{n}\beta > {x_1}^{n-1}$, where the last inequality follows by assumption. Now assume that ${x_i}^{n-1} < y(\frac{i}{n})$, then we prove ${x_{i+1}}^{n-1} < y(\frac{i+1}{n})$. First observe that the Taylor expansion of $y(\frac{i+1}{n})$ around $\frac{i}{n}$ is

$$y(\tfrac{i+1}{n}) = y(\tfrac{i}{n}) + \frac{1}{n}y'(\tfrac{i}{n}) + \frac{1}{2n^2}y''(\tfrac{i}{n}) + \frac{1}{6n^6}y'''(\zeta),$$

with $\zeta \in [\frac{i}{n}, \frac{i+1}{n}]$. As $y''' > 0$, it follows that

$$\begin{aligned} y(\tfrac{i+1}{n}) &> y(\tfrac{i}{n}) + \frac{1}{n}y'(\tfrac{i}{n}) + \frac{1}{2n^2}y''(\tfrac{i}{n}) \\ &= y(\tfrac{i}{n}) + \frac{y(\frac{i}{n})(\ln(y(\frac{i}{n}))-1)-(\beta-1)}{n} + \frac{y(\frac{i}{n})(\ln(y(\frac{i}{n})-1)-(\beta-1))}{2n^2}\ln(y(\tfrac{i}{n})) \\ &\geq \frac{n-1}{n}y(\tfrac{i}{n})^{\frac{n}{n-1}} - \frac{\beta-1}{n} > \frac{n-1}{n}x_i^n - \frac{\beta-1}{n} > x_{i+1}^{n-1}, \end{aligned}$$

where the second inequality is due to Proposition 4.5 and the last inequality follows from Lemma 4.3 and the assumption that $x_1^{n-1} < 1 - \frac{\beta}{n}$. □

We now finish the proof of Theorem 1.3.

Proof of Theorem 1.3. Choosing $0 = \varepsilon_0 < \varepsilon_1 < \dots < \varepsilon_{n-1} < \varepsilon_n = 1$ such that for all i $\int_{\varepsilon_{i-1}}^{\varepsilon_i} \psi(q)(1-q)\,dq = \alpha_{i+1}$, we know by Lemma 4.2 that

$$\mathbb{E}(\mathrm{ADAP}(n,F)) \geq \frac{1}{n\alpha_1}\mathbb{E}(\mathrm{MY}(n,F)),$$

where $\alpha_1 = 1-(1-\varepsilon_1)^{n-1} = 1-x_1^{n-1}$. Hence, we want to show $\frac{1}{n(1-x_1^{n-1})} \geq \frac{1}{\beta}$ for $\beta \approx 1.3415 < \frac{1}{0.745}$.

We prove by contradiction and assume $x_1^{n-1} < 1 - \frac{\beta}{n}$. Then Lemma 4.6 yields $x_n < y(1)$, so we choose β such that $y(1) = 0$ to reach a contradiction with the fact that $x_n = 0$. Note that this indeed implies $y \in [0,1]$ as we assumed. Hereto, note that $y(t)$ is invertible by Lemma 4.4. Hence, we can consider t as a function of y, for which we know $t(1) = 0$, and we want to choose β such that $t(0) = 1$. In particular, we have that

$$t(1) = t(0) + \int_0^1 \frac{dt}{dy} dy = 1 + \int_0^1 \frac{1}{\frac{dy}{dt}} dy = 1 + \int_0^1 \frac{1}{y(\ln(y) - 1) - (\beta - 1)} dy .$$

So β is the value such that the last integral equals -1. This yields $\beta \approx 1.3415 < \frac{1}{0.745}$. □

Extension. As we did in Corollary 2.2 of Section 2, we can extend our result in Theorem 1.3 to an adaptive version of the prophet inequality.

Corollary 4.7. *Given nonnegative i.i.d. random variables, $X_1, \dots, X_n$, with $X_i \sim F$ for all i, there exist thresholds, $\tau_1, \dots, \tau_n$, such that for a sequence σ, drawn uniformly at random, the expected value of the first variable that exceeds its threshold according to that sequence, $X_{\sigma(i)} \geq \tau_i$, is at least a $1/\beta$ fraction of the expected value of* $\max\{X_1, \dots, X_n\}$.

Proof. When F is continuous and strictly increasing, we express the expected value of the maximum in a convenient way:

$$\begin{aligned}\mathbb{E}\left(\max_{i=1,\dots,n} X_i\right) &= n \int_0^1 tF^{n-1} dF(t) = n \int_0^1 (1-q)^{n-1} F^{-1}(1-q) dq \\ &= n \int_0^1 (n-1)(1-q)^{n-2} \left(\int_0^q F^{-1}(1-\theta) d\theta \right) dq = n \int_0^1 (n-1)(1-q)^{n-2} R(q) q dq ,\end{aligned}$$

where $R(q) = \mathbb{E}(X \mid X > F^{-1}(1-q))$ is the expected value of a random variable given that the probability that that variable attains a larger value is at most q (here, $X \sim F$). As we did in the previous mechanism, we partition $[0, 1]$ into n intervals $A_i = [\varepsilon_{i-1}, \varepsilon_i]$, $0 = \varepsilon_0 < \varepsilon_1 < \dots < \varepsilon_{n-1} < \varepsilon_n = 1$, and we set a threshold $\tau_i = F^{-1}(1 - q_i)$ for the random variable arriving at step i, where q_i is drawn from the interval A_i according to the probability density $f_i(q) = \frac{(n-1)(1-q_i)^{n-2}}{\alpha_i}$ with $\alpha_i = \int_{\varepsilon_{i-1}}^{\varepsilon_i} \psi(q) dq$. Since $R(q)$ corresponds to the expected value we get when setting threshold $F^{-1}(1 - q)$, one can prove – following the same reasoning as the analogous part in the proof of Theorem 1.3 – that the expected revenue of this strategy is equal to

$$\sum_{i=1}^{n} \rho_i \int_{\varepsilon_{i-1}}^{\varepsilon_i} (n-1)(1-q)^{n-2} R(q) q \, dq , \tag{6}$$

where $\rho_1 = \frac{1}{\alpha_1}$ and $\rho_{i+1} = \frac{\rho_i}{\alpha_{i+1}} \int_{\varepsilon_{i-1}}^{\varepsilon_i} (n-1)(1-q)^{n-1} dq$ for $i = 1, \dots, n-1$. Again, if we choose $\varepsilon_1, \dots, \varepsilon_{n-1}$ such that $\rho_1 = \rho_2 = \dots = \rho_n$, and solve the recurrence $\{\varepsilon_i\}_{i=1}^{n-1}$ satisfies, then expression (6) equals $\frac{1}{n\alpha_1} \mathbb{E}\left(\max_{i=1,\dots,n} X_i\right) \geq \frac{1}{\beta} \mathbb{E}\left(\max_{i=1,\dots,n} X_i\right)$.

If F is not strictly increasing, replacing $F^{-1}(1-q)$ by $\arg\max F^{-1}(1-q)$ maintains the correctness of the proof. In the case that F is not continuous, one must be more careful when expressing the expected value of the maximum. The result still holds provided the thresholds are chosen randomly, after sampling the probability of acceptance q. □

Remark. A routine exercise shows that the sequence a_n defined by Hill and Kertz [13] exactly equals our $n\alpha_1$. Note here that our α_1 does depend on n, though we have omitted this dependency for simplicity of notation. Thus our result implies that $a_n \leq \beta$ and by the work of Kertz [14] we know that this bound is tight.

Acknowledgments

We thank Bobby Kleinberg for pointing out the connection of our Theorem 1.3 and Corollary 4.7 to the work of Hill and Kertz [13].

REFERENCES

[1] M. Abolhassani, S. Ehsani, H. Esfandiari, M.T. Hajiaghayi, R. Kleinberg, and B. Lucier. Beating 1-1/e for Ordered Prophets. In *STOC 2017*.
[2] M. Adamczyk, A. Borodin, D. Ferraioli, B. de Keijzer, and S. Leonardi. Sequential Posted Price Mechanisms with Correlated Valuations. In *WINE 2015*.
[3] Saeed Alaei, Jason Hartline, Rad Niazadeh, Emmanouil Pountourakis, and Yang Yuan. Optimal Auctions vs. Anonymous Pricing. In *FOCS 2015*.
[4] L. Blumrosen and T. Holenstein. Posted Prices vs. Negotiations: An Asymptotic Analysis. In *EC 2008*.
[5] J.R. Bult and T. Wansbeek. Optimal selection for direct mail. *Marketing Science* 14(4):378–394, 1995.
[6] S. Chawla, J.D. Hartline, D.L. Malec, and B. Sivan. Multi-parameter Mechanism Design and Sequential Posted Pricing. In *STOC 2010*.
[7] V. Choudhary, A. Ghose, T. Mukhopadhyay, and U. Rajan. Personalized Pricing and Quality Differentiation. *Management Science* 51(7):1120–1130, 2005.
[8] R. Cominetti, J. Correa, T. Rothvoß, and J. San Martín. Optimal Selection of Customers for a Last-Minute Offer. *Operations Research* 58(4):878–888, 2010.
[9] P. Dütting, F. Fischer, and M. Klimm. Revenue Gaps for Discriminatory and Anonymous Sequential Posted Pricing. arXiv:cs.GT/1607.07105, 2016.
[10] H. Esfandiari, M. Hajiaghayi, V. Liaghat, and M. Monemizadeh. Prophet Secretary. In *ESA 2015*.
[11] J.P. Gilbert and F. Mosteller. Recognizing the Maximum of a Sequence. *J. Amer. Statist. Assoc.* 61 (1966), 35–76.
[12] J. Hartline. Mechanism Design and Approximation, 2016. http://jasonhartline.com/MDnA/
[13] T.P. Hill and R.P. Kertz. 1982. Comparisons of stop rule and supremum expectations of i.i.d. random variables. *The Annals of Probability*10(2):336–345, 1982.
[14] R.P. Kertz. Stop rule and supremum expectations of i.i.d. random variables: A complete comparison by conjugate duality. *Journal of Multivariate Analysis* 19:88–112, 1986.
[15] O. Kharif. 2013. Supermarkets Offer Personalized Pricing. *Bloomberg* (15 Nov. 2013).
[16] R. Kleinberg and S.M. Weinberg. Matroid Prophet Inequalities. In *STOC 2012*.
[17] U. Krengel and L. Sucheston. Semiamarts and finite values. *Bull. Amer. Math. Soc.* 83:745–747, 1977.
[18] D. Mattioli. 2012. On Orbitz, Mac Users Steered to Pricier Hotels. *The Wall Street Journal* (23 Aug. 2012).
[19] R.B. Myerson. Optimal Auction Design. *Mathematics of Operations Research* 6(1):58–73, 1981.
[20] U. Saint-Mont. A simple derivation of a complicated prophet region. *Journal of Multivariate Analysis* 80:67–72, 2002.
[21] E. Samuel-Cahn. Comparisons of threshold stop rule and maximum for independent nonnegative random variables. *The Annals of Probability* 12(4):1213–1216, 1983.
[22] C. Shapiro and H.R. Varian. 1999. *Information Rules: A Strategic Guide to the Network Economy.* Harvard Business School Press, Cambridge, MA.
[23] Q. Yan. Mechanism design via correlation gap. In *SODA 2011*.