



Operations Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Static Pricing for Multi-unit Prophet Inequalities

Shuchi Chawla, Nikhil Devanur, Thodoris Lykouris

To cite this article:

Shuchi Chawla, Nikhil Devanur, Thodoris Lykouris (2023) Static Pricing for Multi-unit Prophet Inequalities. Operations Research
Published online in Articles in Advance 01 Nov 2023
. <https://doi.org/10.1287/opre.2023.0031>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2023, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Crosscutting Areas

Static Pricing for Multi-unit Prophet Inequalities

Shuchi Chawla,^a Nikhil Devanur,^b Thodoris Lykouris^{c,*}

^aDepartment of Computer Science, The University of Texas at Austin, Austin, Texas 78712; ^bAmazon, Seattle, Washington 98109; ^cSloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts 02142

*Corresponding author

Contact: shuchi@cs.texas.edu,  <https://orcid.org/0000-0001-5583-2320> (SC); iam@nikhildevanur.com,

 <https://orcid.org/0009-0005-4406-5935> (ND); lykouris@mit.edu,  <https://orcid.org/0000-0002-3375-5579> (TL)

Received: January 18, 2023

Revised: June 20, 2023

Accepted: August 28, 2023

Published Online in Articles in Advance:
November 1, 2023

Area of Review: Market Analytics and
Revenue Management

<https://doi.org/10.1287/opre.2023.0031>

Copyright: © 2023 INFORMS

Abstract. We study a pricing problem where a seller has k identical copies of a product, buyers arrive sequentially, and the seller prices the items aiming to maximize social welfare. When $k = 1$, this is the so-called *prophet inequality* problem for which there is a simple pricing scheme achieving a competitive ratio of $1/2$. On the other end of the spectrum, as k goes to infinity, the asymptotic performance of both static and adaptive pricing is well understood. We provide a static pricing scheme for the small-supply regime: where k is small but larger than one. Prior to our work, the best competitive ratio known for this setting was the $1/2$ that follows from the single-unit prophet inequality. Our pricing scheme is easy to describe as well as practical; it is anonymous, nonadaptive, and order oblivious. We pick a single price that equalizes the expected fraction of items sold and the probability that the supply does not sell out before all customers are served; this price is then offered to each customer while supply lasts. This extends an approach introduced by Samuel-Cahn for the case of $k = 1$. This pricing scheme achieves a competitive ratio that increases gradually with the supply. Subsequent work shows that our pricing scheme is the optimal static pricing for every value of k .

Funding: This work was supported by the National Science Foundation [Grants CCF-2008006 and SHF-1704117]. T. Lykouris would like to acknowledge funding from Google Research.

Keywords: static threshold policies • optimal pricing • social welfare maximization

1. Introduction

The *prophet inequality* problem of Krengel and Sucheston (1977) constitutes one of the cornerstones of online decision making. A designer knows a set of n distributions $\mathcal{F}_1, \dots, \mathcal{F}_n$ from which random variables $X_t \sim \mathcal{F}_t$ are sequentially realized in an arbitrary order. Once a random variable is realized, the designer decides whether to accept it or not; at most one realized random variable can be accepted. The objective is to maximize the value of the variable accepted, and the performance of the algorithm is evaluated against the ex post maximum realized. In a beautiful result, Samuel-Cahn (1984) showed that a simple static threshold policy achieves the optimal competitive ratio for this problem. The algorithm of Samuel-Cahn (1984) determines a threshold p such that the probability that there exists a realization exceeding the threshold is exactly $1/2$ and then, accepts the first random variable that exceeds the threshold. This algorithm achieves a competitive ratio of $1/2$ against the ex post optimum; no online algorithm, even one with adaptive thresholds, can obtain better performance.

Over the last few years, many extensions of the basic prophet inequality to more general feasibility constraints

have been studied, and tight bounds on the competitive ratio have been established. However, one simple natural extension has largely been overlooked: where the designer is allowed to accept $k > 1$ random variables for some small value of k . This is called the *multiunit prophet inequality*. When k is relatively large, then it is known that static threshold policies can achieve a competitive ratio of $1 - O(\sqrt{\log(k)/k})$ (Hajiaghayi et al. 2007), which goes to one as $k \rightarrow \infty$. However, (for example) for $k=2$ or 3 , prior to our work, the best known competitive ratio of static thresholds remained $1/2$. Our work addresses this gap by posing and answering the following questions.

Can a static threshold policy achieve a better competitive ratio than $1/2$ for small $k = 2, 3, \dots$?

How should it be computed as a function of k ? How does its performance scale with k ?


A primary motivation for our work is its connection to welfare maximization in mechanism design. In this application, a seller has one or more units of an item to sell. The distributions correspond to known priors on

the valuations of different customers (possibly heterogeneous), and the realizations correspond to the actual valuation of an incoming customer. The seller's goal is to maximize the social welfare or the sum of values of the customers that obtain the item. Any online strategy for the prophet inequality problem corresponds to selecting prices for customers; customers buy if any units of the item are still available and their valuation is higher than the price. Static threshold policies correspond to static pricings, where the seller simply places a fixed price on the item and customers can purchase the item at that price while supply lasts. Static pricings have many nice properties that make them practical and suitable for real-world contexts. They are *non-adaptive* (the price does not depend on which customers have already arrived) and *order oblivious* (the price does not depend on the order of customers). This makes their implementation simpler and removes the incentive on customers to strategize on the arrival order to obtain a better price, enhancing the customer experience. We note that, although the order of customers does not affect the price assuming that the supply is not depleted, it *does* affect the probability that the supply is depleted; this probability is zero for the first customer and increases as customers arrive. Finally, static pricing is *anonymous* (it does not discriminate based on which customer arrives), which is typically regarded as a more fair pricing scheme. We, therefore, focus on static pricings in this work.

1.1. Our Results


We answer the questions by developing an algorithm for finding a static threshold policy for the multiunit prophet inequality that is sensitive to the supply k . Our algorithm is very simple and practical. For any fixed price p , it estimates two statistics based on the given prior: (1) the fraction of items expected to be sold at that price, $\mu_k(p)$, illustrated in Figure 1, and (2)

Figure 1. (Color online) For a Price p^* , We Cannot Hope for More Revenue Than Selling All Units at p^* (i.e., the Red Dashed Area: $R_k(p^*) = kp^*$)



Note. However, a price p^* thins the demand and results in an expected number of sales that is equal to $k \cdot \mu_k(p^*)$ and thereby, an expected revenue equal to the area below the blue dotted region: $\mu_k(p^*) \cdot R_k(p^*)$.

Figure 2. (Color online) For a Price p^* , We Cannot Hope for More Consumer Surplus (Utility) Than the One of All Customers with Value Higher Than p^* : $U(p^*) := \sum_t \max(0, v_t - p^*)$



Notes. However, we do not receive utility from customer t if there is no unit when she arrives; a lower bound on the probability of having a unit available is the probability $\delta_k(p^*)$ that not all units are depleted anyway. This lower bounds the welfare we collect from consumer surplus by $\delta_k(p^*) \cdot U(p^*)$.

the probability that not all units will sell out before all the customers have been served, $\delta_k(p)$, illustrated in Figure 2. We then pick the static price p^* at which these two quantities are equal: $\mu_k(p^*) = \delta_k(p^*)$.

The competitive ratio of this static pricing increases gracefully as the supply increases and approaches one at the rate of $1 - O(\sqrt{\log(k)/k})$ as $k \rightarrow \infty$, as we will explain next. The precise competitive ratio at any particular value of k can be determined as the solution to a particular equation; let X be a Poisson random variable with a rate defined such that the following equation holds:

$$\frac{1}{k} \mathbb{E}[\min(X, k)] = \mathbb{P}[X \leq k - 1]. \quad (1)$$


The worst-case competitive ratio of our algorithm is then given by the value of either side of the equation, say $\mathbb{P}[X \leq k - 1]$. Note that this quantity is well defined because on the one hand, the truncated expectation $\frac{1}{k} \mathbb{E}[\min(X, k)]$ increases with the rate of the Poisson variable X ; it is zero for rate equal to zero and one for rate equal to infinity. On the other hand, the probability $\mathbb{P}[X \leq k - 1]$ decreases with the rate; it is one when the rate is zero and zero when the rate is infinity. In effect, our analysis shows that the worst case for our static pricing occurs when the number of customers with value exceeding the price is given precisely by the Poisson variable X . As $k \rightarrow \infty$, the competitive ratio for this instance tends to $1 - O(\sqrt{\log(k)/k})$.

To obtain a better sense of the exact quantities the equation leads to, Figure 3 depicts the ratio as a function of k , and Table 1 instantiates it for small values of k .

1.2. Our Techniques

1.2.1. The Approach of Samuel-Cahn (1984): Balancing Revenue and Utility Contribution to Welfare. Our methodology is inspired by the approach of Samuel-Cahn (1984) for the single-unit prophet inequality ($k = 1$). First,

Figure 3. (Color online) Competitive Ratio of Our Static Pricing as a Function of the Number of Supply Units k




the social welfare obtained by any static price p can be expressed in two parts: (1) the expected revenue the seller obtains from selling any units of the item and (2) the expected utility the buyers obtain from purchasing any units of the item. What is the most revenue and utility that we can expect at a particular price p ? The most revenue the seller can obtain is simply the price p , $R(p) := p$. On the other hand, the most total utility the buyers can obtain at a price of p is $U(p) := \sum_t \max(0, v_t - p)$ or the total excess value of the buyers above price p assuming that *everyone who wants the item gets it*. (In the pricing application, customers are assumed to have *quasilinear utilities* (i.e., they buy when their value is above the price and there is an available item and obtain payoff equal to their value minus the price).) It turns out that no matter what p is, $R(p) + U(p)$ is an upper bound on the optimal-in-hindsight social welfare. Samuel-Cahn (1984) observed that when $k=1$, with the right choice of p both the seller and the buyers can in expectation each obtain at least a half of these revenue and utility upper bounds, respectively. One way to choose such a price p is to ensure that the probability of selling the item is exactly $1/2$. At that price, on the one hand, the buyer sells $1/2$ units in expectation, and on the other hand, each buyer has a probability at least $1/2$ of being offered the item and contributing to the total utility, resulting in the competitive ratio of $1/2$. Figure 4 illustrates the idea behind the approach of Samuel-Cahn (1984).

Table 1. Competitive Ratio of Our Static Pricing for the Small Number of Supply Units k

Number of units	Competitive ratio
$k = 1$	0.5
$k = 2$	0.585
$k = 3$	0.630
$k = 4$	0.660
$k = 5$	0.682
$k = 6$	0.698

Figure 4. (Color online) This Figure Is Similar to Figure 1 as, for $k = 1$, $\mu_1(p) = \mathbb{E}[\min(\sum_{t=1}^n \mathbf{1}\{X_t \geq p\}, k)]$ and the Complement of Figure 2 as $\delta_1(p) = \mathbb{P}[\sum_{t=1}^n \mathbf{1}\{X_t \geq p\} \leq k - 1]$



Notes. For $k = 1$, $\mu_1(p) = 1 - \delta_1(p)$ for all prices and selecting p^* such that $\mu_1(p^*) = \delta_1(p^*)$ leads to competitive ratio of $1/2$. The root behind our algorithm is to extend this revenue-utility decomposition to $k > 1$ where it no longer holds that $\mu_k(p) = 1 - \delta_k(p)$ for all prices p .

1.2.2. Our Extension to Multiple Units. Extending this approach beyond a single unit, we similarly define $R_k(p) := pk$ to be the revenue obtained if all k units of the item get sold at price p and $U(p) := \sum_t \max(0, v_t - p)$ to be the total excess value of the buyers above price p assuming that everyone who wants the item at price p gets it. Then, $R_k(p) + U(p)$ is an upper bound on the optimal in hindsight social welfare. Letting $\mu_k(p)$ denote the expected fraction of the supply sold at the price p , the seller's expected revenue is $\mu_k(p)R_k(p)$. On the other hand, the probability that a buyer is offered an unsold unit is at least as large as the probability that not all units are sold out at the end of the process; we call this probability $\delta_k(p)$. Then, the total utility contributed by the buyers is at least $\delta_k(p)U(p)$. The static pricing p , therefore, obtains at least a $\phi_k := \min\{\mu_k(p), \delta_k(p)\}$ fraction of the upper bound $R_k(p) + U(p)$. Our pricing scheme selects the price that maximizes this quantity ϕ_k . Because $\mu_k(p)$ is a decreasing function of p and $\delta_k(p)$ is an increasing function, their minimum is maximized when the two are equal. Note that ϕ_k only depends on the buyers' value distributions and is independent of their order as both $\mu_k(p)$ and $\delta_k(p)$ are also order-oblivious quantities.

1.2.3. Crux of Our Analysis: Characterizing Worst-Case Performance of Our Scheme. The description quantifies the competitive ratio of our scheme for any known distribution. To characterize its worst-case performance, we need to also identify worst-case distributions (i.e., those resulting to the lowest ϕ_k). The crux of our analysis is a series of reductions eventually showing that Poisson distributions are these worst-case distributions. As a result, the competitive ratio of our scheme is ϕ_k for Poisson distributions, and this leads to the competitive ratio we illustrated in (1).

1.3. Related Work

As already discussed, prophet inequalities were introduced by Krenkel and Sucheston (1977); Samuel-Cahn

(1984) provided a very clean analysis that our work builds upon. In the last decade, there has been a tremendous amount of work on extending prophet inequalities to different feasibility constraints over buyers (Kleinberg and Weinberg 2012, Dütting and Kleinberg 2015, Rubinstein and Singla 2017) as well as to pricing with heterogeneous items where buyers have more complex valuations (Feldman et al. 2014; Chawla et al. 2017, 2019; Dütting et al. 2020). The reader is referred to Lucier (2017) for a general survey.

1.3.1. Balanced Prices. The dominant approach for establishing prophet inequalities in combinatorial settings is by constructing so-called *balanced prices*, a technique introduced by Kleinberg and Weinberg (2012) and further developed in Feldman et al. (2014) and Dütting et al. (2020). This approach also has its roots in the work of Samuel-Cahn (1984). Recall that the optimal social welfare is bounded by $R_k(p) + U(p)$ and that the pricing p obtains a $\mu_k(p)$ fraction of the first term and a $\delta_k(p)$ fraction of the second. Feldman et al. (2014) choose a price $p_{1/2}$ that balances the revenue and utility upper bounds: $R_k(p_{1/2}) = U(p_{1/2})$. They accordingly obtain a competitive ratio of $\frac{1}{2}(\mu_k(p_{1/2}) + \delta_k(p_{1/2}))$. Noting that $\mu_k(p) + \delta_k(p) \geq 1$ at any price p , this competitive ratio is always at least $1/2$, but in general, it is no better even when k is large. In contrast, our approach picks a price where μ and δ are simultaneously larger than $1/2$ (Figure 5).

1.3.2. Adaptive Pricing and the Magician's Problem. Multiunit prophet inequalities were also previously studied by Alaei (2014) in the context of revenue optimal mechanism design. Alaei (2014) provided a competitive ratio of $\alpha_k := 1 - \sqrt{1/(k+3)}$ for a more general problem called the magician's problem, which also applies to the multiunit prophet inequality. The pricing scheme of Alaei (2014) is not static in the sense of a price that you "set it and forget it." It uses a single price for all buyers but probabilistically skips some buyers in order to maintain a certain probability of not


running out of items. This is equivalent to offering a price of ∞ to some users. As such, it is more powerful than the class of static pricing policies that we consider, and it does not satisfy two of the three properties mentioned: nonadaptivity and order obliviousness. The static pricing we develop provides a strictly better competitive ratio than that of Alaei (2014) for $k \in [2, 20]$. Subsequent to our work, Jiang et al. (2022) improved upon the bound of Alaei (2014) to obtain tight competitive ratios for the magician's problem at all values of k . This bound is again achieved by a dynamic pricing and therefore, strictly exceeds the bound achieved by our static pricing.

1.3.3. Tightness of Our Pricing Scheme. Subsequently to our work, Jiang et al. (2023) established that our static pricing policy is worst-case optimal across all static threshold policies. In particular, Jiang et al. (2023) provide an instance where the gap between any static policy and the ex post optimal is exactly ϕ_k . We note that, for specific instances, one can obtain improved static prices. As a result, our policy is not instance optimal. For example, consider the policy that selects the price p that directly maximizes $\mu_k(p)R_k(p) + \delta_k(p)U(p)$; such a policy would also enjoy the competitive ratio guarantees that we provide while obtaining improved performance on some distributions. However, this policy requires the evaluation of $U(p)$ for all prices p , which makes it less simple (especially when one needs to estimate these quantities from data). In contrast, our policy only requires learning the price that equates $\mu_k(p)$ and $\delta_k(p)$, which can be done via binary search on the prices (as both can be evaluated with only samples from price p). Finally, although we present our results as comparing against the ex post optimum, the same guarantees also hold against the corresponding ex ante relaxation (see Remark 2).

1.3.4. Other Subsequent Work. In another subsequent work, Arnosti and Ma (2022) build on our techniques to study the performance of static threshold policies in the prophet-secretary setting with k units, where customers arrive in uniformly random order (rather than in a worst-case order as what we consider). For this special case, Arnosti and Ma (2022) provide tight competitive ratio guarantees; in doing so, they heavily rely on the structure of our analysis and use the random order property to establish a better competitive ratio guarantee for the final step of our analysis. The performance of the worst-case distribution is the Poisson distribution, which is similar to our result.

1.3.5. Pricing with Limited Supply Beyond Prophet Inequalities. Our work lies in the general theme of providing supply-dependent guarantees for pricing with

Figure 5. (Color online) The Price $p_{1/2}$ Is Chosen Such That the Shaded Red Area Is Half of the Area Under the Solid Blue Curve



Note. This price is in general smaller than the price p^* our approach chooses, and its worst-case competitive ratio is $1/2$ for any k .

known priors and limited supply. Beyond prophet inequalities, such guarantees have also been provided in ride-sharing settings (Balseiro et al. 2021, Banerjee et al. 2022). The latter works typically make a stronger assumption that the system is in steady state but has more complex state externalities. In multiunit prophet inequalities, the supply just decreases when items are sold; in ridesharing, it is reallocated across the network. To the best of our knowledge, these are the only two pricing settings where such supply-dependent guarantees with known priors and limited supply are provided; most prior work focuses on asymptotic optimality guarantees when the supply is large.

When the priors are not known in advance, a few other lines of work attempt to address these settings with the additional complication of learning information about the distributions. For example, dynamic pricing with limited supply has been studied in the context of prior-independent mechanisms (i.e., those that do not have distributional knowledge (Babaioff et al. 2015)); this work has been then extended in more general bandit settings under knapsack constraints (Badanidiyuru et al. 2018, Agrawal and Devanur 2019). On the positive side, these approaches do not assume knowledge of the distributions; on the negative side, the guarantees they provide become meaningful only when the supply is large (e.g., $\sqrt{n}$, where n is the number of buyers).

2. Model

An instance of the prophet inequality problem consists of a set of n distributions supported on nonnegative real numbers $\mathcal{F} = \{\mathcal{F}_t : t \in [n]\}$. To ease of presentation, we denote by $\mathcal{F}_t$ the t th arriving distribution; this order is not known to the seller. In multiunit prophet inequalities, there is also a supply k that determines the number of units available for purchase at the beginning of time.

A static price is defined by a single number $p \in \mathbb{R}$. The pricing works as follows. Buyers arrive one by one and are offered a copy of the item at price p as long as there is available supply. Buyer t has a value v_t drawn independently from the distribution $\mathcal{F}_t$. The buyer purchases a unit of the item if and only if her value is above the price and there is an available item. In this case, the available supply decreases by one; otherwise, the buyer leaves the system without an item, and the available supply remains unaltered.

A static pricing scheme π maps the supply k and the distributions $\mathcal{F}$ to a static price $\pi(k, \mathcal{F}) \in \mathbb{R}$.

The welfare of a static price $p \in \mathbb{R}$ on a particular realization of buyer values is the total value of the buyers who purchase a unit of the item. We denote its expected welfare by $\text{WELFARE}(p, k, \mathcal{F})$, where the expectation is over the randomness in buyer values drawn from $\mathcal{F}$.

The benchmark we compare with is the expected optimal welfare in hindsight and is denoted by $\text{OPT}(k, \mathcal{F})$ (i.e., $\text{OPT}(k, \mathcal{F})$ is the expected sum of the k -highest realized values drawn from the set of distributions $\mathcal{F}$). The *competitive ratio* for a static pricing scheme π on supply k is the worst-case welfare-to-optimum ratio across all the possible set of distributions $\mathcal{F}$: that is,

$$\text{COMPRATIO}(\pi, k) = \inf_{\mathcal{F}} \frac{\text{WELFARE}(\pi(k, \mathcal{F}), k, \mathcal{F})}{\text{OPT}(k, \mathcal{F})}.$$

Our goal is to identify a static pricing scheme π that maximizes this worst-case competitive ratio. In the remainder of the presentation, we omit the arguments of $\pi(k, \mathcal{F})$ when clear from the context.

Without loss of generality, we assume that each distribution has $\mathbb{P}[v_t > 0] > 0$ (otherwise, we can ignore it) and further assume that $n > k$ (otherwise, setting a price of zero is optimal).

To ease the presentation of our scheme, we assume that the distributions are atomless. Remark 1 shows how our results extend to general distributions.

3. Our Pricing Scheme and Its Performance Guarantee

3.1. Decomposing to Revenue and Utility Contributions

For any fixed price $p \in \mathbb{R}$ and distributions $\mathcal{F}$, let X_p denote the number of buyers who have value higher than the price. This is a random variable because the buyers' values are drawn from the distributions $\mathcal{F}$; in particular, it is equal to

$$X_p := |\{t : v_t \geq p\}|.$$

As in the approach of Samuel-Cahn (1984), we decompose the welfare into two components: the total utility obtained by the buyers and the total revenue obtained by the seller. We now define some quantities of interest that determine these components. The first quantity is the probability that the seller runs out of units to sell or in other words, that X_p is at least k . We use $\delta_k(X_p)$ to denote one minus this probability:

$$\delta_k(X) := \mathbb{P}[X \leq k - 1]. \quad (2)$$

The second quantity is the expected fraction of units sold and is directly related to the revenue obtained by the seller. We use $\mu_k(X_p)$ to denote this truncated expectation:

$$\mu_k(X) := \frac{1}{k} \mathbb{E}[\min\{X, k\}]. \quad (3)$$

The first important lemma that drives the design of our pricing scheme is that, for any distributions $\mathcal{F}$, the welfare-to-optimum ratio is at least the minimum of these two quantities.

Lemma 1. For any supply k , set of distributions $\mathcal{F}$, and any price $p \in \mathbb{R}$,

$$\frac{\text{WELFARE}(p, k, \mathcal{F})}{\text{OPT}(k, \mathcal{F})} \geq \min(\delta_k(X_p), \mu_k(X_p)).$$

This lemma is the main structural contribution of our work to the prophet inequality literature. Its proof is deferred to Section 4.

3.2. Our Pricing Scheme

For a given set of distributions $\mathcal{F}$ and supply k , our pricing scheme π outputs a static price $\pi(k, \mathcal{F}) = \pi$ that ensures that the two quantities in Lemma 1 are equalized:

$$\delta_k(X_\pi) = \mu_k(X_\pi). \quad (4)$$

The atomless assumption ensures that such a price always exists (see Remark 1 on how the results extend beyond the atomless assumption). Observe that $\delta_k(X_p)$ is monotone nondecreasing in p and $\mu_k(X_p)$ is monotone nonincreasing in p . Moreover, $\mu_k(X_p)$ goes from one to zero as the price goes from zero to ∞ (because $n > k$). The atomless assumption ensures that both $\delta_k(X_p)$ and $\mu_k(X_p)$ are continuous. The intermediate value theorem then guarantees the existence of π .

We now define the competitive ratio of our pricing scheme for any distributions $\mathcal{F}$ and supply k as

$$\phi_k(\pi, \mathcal{F}) := \delta_k(X_{\pi(k, \mathcal{F})}) = \mu_k(X_{\pi(k, \mathcal{F})})$$

and the worst-case competitive ratio of $\mathbf{p}$ as

$$\phi_k(\pi) := \inf_{\mathcal{F}} \phi_k(\pi, \mathcal{F}). \quad (5)$$

The second important lemma that enables our competitive ratio guarantee is that, for any fixed supply k , the minimum of $\delta_k(X_{\pi(k, \mathcal{F})})$ and $\mu_k(X_{\pi(k, \mathcal{F})})$ attains its lowest value when $\mathcal{F}$ consists of infinitely many Bernoulli random variables, all with equal bias; in this case, X_p is a Poisson distribution. This is formalized in the following lemma.

Lemma 2. For any supply k and any set of distributions $\mathcal{F}$, $\phi_k(\pi, \mathcal{F})$ attains its lowest value $\phi_k(\mathbf{p})$ when $\mathcal{F}$ is a collection of infinitely many Bernoulli distributions with equal bias (i.e., $X_{\pi(k, \mathcal{F})}$ is a Poisson distribution).

The proof of this lemma stems from a series of reductions and is the main technical contribution of our analysis. Its proof is deferred to Section 5.

3.3. Competitive Ratio of Our Pricing Scheme

The two lemmas seamlessly establish the competitive ratio of our pricing scheme as demonstrated in the following theorem, which is the main result of our work. The competitive ratio ϕ_k as a function of k is illustrated in Figure 3.

Theorem 1. Let X^λ be a Poisson random variable with rate λ and set λ_k such that $\delta_k(X^{\lambda_k}) = \mu_k(X^{\lambda_k})$. The competitive ratio of our pricing scheme π is at least $\phi_k := \delta_k(X^{\lambda_k}) = \mu_k(X^{\lambda_k})$.

Proof. The proof of the theorem comes directly by combining Lemmas 1 and 2. $\square$

Remark 1. If there are point masses in the distributions at price $\pi(k, \mathcal{F})$, we still obtain the same results provided we can break ties at random. A buyer with value $\pi(k, \mathcal{F})$ is allocated the item with a probability such that Equality (4) holds. The definition of $X_{\pi(k, \mathcal{F})}$ is adjusted accordingly; if $v_t = \pi(k, \mathcal{F})$, then t is counted only with some probability. The same effect can be achieved by randomly perturbing the price by an infinitesimal amount (although not static, this is still anonymous, nonadaptive, and order oblivious).

4. Welfare-to-Optimum Lower Bound for Any Price (Lemma 1)

This section proves Lemma 1. For any supply k , set of distributions $\mathcal{F}$, and any price $p \in \mathbb{R}$, we have

$$\frac{\text{WELFARE}(p, k, \mathcal{F})}{\text{OPT}(k, \mathcal{F})} \geq \min(\delta_k(X_p), \mu_k(X_p)),$$

where $\delta_k(X) := \mathbb{P}[X \leq k - 1]$ and $\mu_k(X) := \mathbb{E}[\min\{X, k\}/k]$ as introduced in Equations (2) and (3).

Proof of Lemma 1. The proof follows the approach of Samuel-Cahn (1984) for the single-unit prophet inequality. We first bound the hindsight optimal welfare from above in terms of the price p by bounding both the maximum possible revenue generated for the seller and the maximum possible utility generated for the buyers when posting price p . Let Z_p denote the (random) set of buyers whose value exceeds the price p . Then, we have

$$\begin{aligned} \text{OPT}(k, \mathcal{F}) &= \mathbb{E} \left[\max_{S \subseteq [n]; |S| \leq k} \sum_{t \in S} v_t \right] \\ &\leq \mathbb{E} \left[\max_{S \subseteq [n]; |S| \leq k} \sum_{t \in S} (p + \max(0, v_t - p)) \right] \\ &\leq kp + \mathbb{E} \left[\sum_{t \in Z_p} (v_t - p) \right] \\ &\leq kp + \sum_{t \in [n]} \mathbb{P}[t \in Z_p] \mathbb{E}[v_t - p | t \in Z_p]. \quad (6) \end{aligned}$$

We note that the first summand in the last term corresponds to what we referred to in Section 1 as $R_k(p) = kp$, whereas the second summand corresponds to the expected value of what we referred in Section 1 as $U(p) = \sum_t \mathbb{E}[\max(0, v_t - p)]$.

We now decompose the expected welfare generated by price p to a revenue and a utility component. The expected revenue of the seller upon setting price p is

$$\text{REVENUE}(p, k, \mathcal{F}) = \mathbb{E}[\min(|Z_p|, k)] p = \mu_k(X_p)kp. \quad (7)$$

On the other hand, a buyer t receives utility of $v_t - p$ if and only if (1) v_t is at least p (that is, $t \in Z_p$) and (2) the item is still available when the buyer t arrives. Regardless of the order in which buyers arrive, the latter event happens with probability at least as large as the probability that the item is not sold out at the end of the process. Recall that this latter probability is $\delta(X_p) = \mathbb{P}[X_p \leq k - 1]$. We, therefore, get the following lower bound on the utility generated by the pricing p :

$$\text{UTILITY}(p, k, \mathcal{F}) \geq \sum_{t \in [n]} \delta_k(X_p) \cdot \mathbb{P}[t \in Z_p] \mathbb{E}[v_t - p | t \in Z_p]. \quad (8)$$

The proof of the lemma is completed by putting Equations (6), (7), and (8) together, which yields

$$\begin{aligned} \text{WELFARE}(p, k, \mathcal{F}) &= \text{REVENUE}(p, k, \mathcal{F}) + \text{UTILITY}(p, k, \mathcal{F}) \\ &\geq \mu_k(X_p)kp + \delta_k(X_p) \sum_{t \in [n]} \mathbb{P}[t \in Z_p] \\ &\quad \mathbb{E}[v_t - p | t \in Z_p] \\ &\geq \min(\mu_k(X_p), \delta_k(X_p)) \text{OPT}(k, \mathcal{F}). \quad \square \end{aligned}$$

Remark 2. Although the result on Lemma 1 bounds the expected welfare of any static price p to the expected value of $\text{OPT}(k, \mathcal{F})$, we note that its proof also implies that the same guarantee holds against the ex ante relaxation.

5. Establishing Poisson as Worst-Case Distribution (Lemma 2)

This section proves Lemma 2. Recall that for any supply k and any set of distributions $\mathcal{F}$, $\pi(k, \mathcal{F})$ is the price that satisfies $\delta_k(X_{\pi(k, \mathcal{F})}) = \mu(X_{\pi(k, \mathcal{F})})$. We show that the corresponding competitive ratio $\phi_k(\pi)$ attains its lowest value when $\mathcal{F}$ is a collection of infinitely many Bernoulli distributions with equal bias (i.e., $X_{\pi(k, \mathcal{F})}$ is a Poisson distribution).

Proof Sketch. To prove the lemma, we progressively refine our understanding of the worst-case distributions, as outlined in the following three steps.

1. We reduce the problem of finding the worst distribution to a finite-dimensional problem searching only over Bernoulli distributions (Section 5.1). Intuitively, our analysis is only affected by the probability that $v_t \geq \pi(k, \mathcal{F})$ corresponding to the bias of a Bernoulli distribution.

2. We show that all the Bernoulli biases are equal unless they are either zero or one (Section 5.2).

3. We show that the Bernoullis, in fact, must all have the same bias (Section 5.3).

The lemma then follows by considering n Bernoullis with the same bias and letting n tend to infinity. The complete proof is provided at the end of the section. $\square$

5.1. Reducing to Bernoulli Distributions

5.1.1. Reducing Worst-Case Distributions to Bernoulli Distributions. A Bernoulli random variable with bias b takes on the value of one with probability b and zero otherwise. We reduce the problem of finding the worst-case distribution to the following finite-dimensional problem:

$$\begin{aligned} \phi_k^* &:= \min_{b_1, b_2, \dots, b_n, \phi} \phi \quad \text{s.t.} \\ &\quad (\text{GENBERN} : \min \phi \quad \text{s.t.} \quad \delta = \mu = \phi) \end{aligned}$$

X is the sum of n Bernoullis with bias

$$b_1, b_2, \dots, b_n$$

$$\delta_k(X) = \mu_k(X) = \phi.$$

Lemma 3. For any supply $k > 0$ and number $n > k$ of customers, the worst-case competitive ratio $\phi_k(\pi)$ of our pricing scheme is equal to the optimal value of $(\text{GENBERN} : \min \phi \quad \text{s.t.} \quad \delta = \mu = \phi)$, ϕ_k^* .

Proof. The worst-case competitive ratio $\phi_k(\pi)$ optimizes the objective of the problem $(\text{GENBERN} : \min \phi \quad \text{s.t.} \quad \delta = \mu = \phi)$ across any set of prior distributions, whereas the optimization problem $(\text{GENBERN} : \min \phi \quad \text{s.t.} \quad \delta = \mu = \phi)$ optimizes only over Bernoulli distributions. We show that for every set of prior distributions, there exists a corresponding set of Bernoulli distributions that are feasible for $(\text{GENBERN} : \min \phi \quad \text{s.t.} \quad \delta = \mu = \phi)$ and obtain the same objective function value.

The reduction is relatively simple. For any set of distributions $\mathcal{F}$ (not necessarily Bernoulli), we first compute the price $\pi(k, \mathcal{F})$ of our pricing scheme (i.e., the one that makes $\delta_k(X_{\pi(k, \mathcal{F})}) = \mu_k(X_{\pi(k, \mathcal{F})})$). Subsequently, for each distribution $\mathcal{F}_t$, we compute an equivalent Bernoulli bias $b_t = \mathbb{P}[v_t \geq \pi(k, \mathcal{F})]$. The probability that any Bernoulli random variable is one is, therefore, equal to the probability that its original counterpart is higher than the price $\pi(k, \mathcal{F})$. As a result, both $\delta_k(\cdot)$ and $\mu_k(\cdot)$ are the same for the resulting sum as in the original problem, which proves the lemma. $\square$

5.1.2. A Simpler Equivalent Way to Express the Resulting Optimization Problem. We now define a slightly different form of the objective function, which makes analyzing the optimal setting of the biases easier.

Here, ϕ_k^* is the optimal value of problem (GENBERN : $\min\phi$ s.t. $\delta = \mu = \phi$):

$$\begin{aligned} \min \mu_k(X) \text{ s.t. } & \text{(GENBERN : } \min\mu \text{ s.t. } \delta = \phi) \\ & \delta_k(X) = \phi_k^* \\ & X \text{ is the sum of Bernoulli r.v.s.} \end{aligned}$$

Lemma 4. *The optimal values of the optimization problems (GENBERN : $\min\phi$ s.t. $\delta = \mu = \phi$) and (GENBERN : $\min\mu$ s.t. $\delta = \phi$) are equal.*

Proof. The optimal solution ϕ_k^* for the optimization problem (GENBERN : $\min\phi$ s.t. $\delta = \mu = \phi$) is feasible for (GENBERN : $\min\mu$ s.t. $\delta = \phi$) as the latter program needs to satisfy a subset of the former program's constraints; thus, the optimum of (GENBERN : $\min\mu$ s.t. $\delta = \phi$) is no larger than the one of (GENBERN : $\min\phi$ s.t. $\delta = \mu = \phi$).

For the opposite direction, assume that the optimum of (GENBERN : $\min\mu$ s.t. $\delta = \phi$) is strictly smaller than the one of (GENBERN : $\min\phi$ s.t. $\delta = \mu = \phi$). Because $\delta_k(X)$ and $\mu_k(X)$ are both continuous and they are monotone decreasing and increasing, respectively, by increasing any of the biases, starting from the optimal solution of (GENBERN : $\min\mu$ s.t. $\delta = \phi$), we arrive to a new solution X' , with $\delta_k(X') = \mu_k(X') < \delta_k(X) = \phi_k^*$. This contradicts the fact that ϕ_k^* was the optimum for (GENBERN : $\min\phi$ s.t. $\delta = \mu = \phi$) and establishes that its optimal value is no larger than the one of (GENBERN : $\min\mu$ s.t. $\delta = \phi$). $\square$

5.2. Reducing to Bernoulli Distributions with Equal Bias Unless Degenerate

We now show that the optimum of the problem (GENBERN : $\min\mu$ s.t. $\delta = \phi$) is attained when all the Bernoulli distributions either have equal bias or are degenerate (with bias 0 or 1).

Lemma 5. *The optimization problem (GENBERN : $\min\mu$ s.t. $\delta = \phi$) is minimized when all nondegenerate Bernoulli distributions (that do not have bias 0 or 1) have equal bias.*

We note that Hoeffding (1956) provides a similar result but without the constraint on $\delta_k(X)$. A generalization of Lemma 5 was obtained via a similar case analysis by subsequent work of Arnosti and Ma (2022) for the case where the objective and the constraint involve the expectation of an arbitrary nonnegative integer-valued function of X ; in our setting, those functions are $\mu_k(X)$ and $\delta_k(X)$, respectively. Readers familiar with these results can safely skip this section and move to Section 5.3.

5.2.1. High-Level Structure of the Reduction. The key idea of the proof is to fix all but two of the biases and consider the problem of minimizing $\mu_k(X)$ subject to $\delta_k(X)$ being fixed as a function of these two biases. This

is a problem in two dimensions, and we can characterize the optimal solutions to this problem. We then set aside these two distributions and assume by the principle of deferred decisions that they are instantiated in the end. The eventual goal is to establish that $\mu_k(X)$ is minimized when these two biases are equal or degenerate (either zero or one). By working inductively on the number of biases that are not equal and are nondegenerate, we eventually establish that all biases should be equal or degenerate.

Formally, assume that we have n Bernoulli distributions and fix all but two biases; let b_1 and b_2 be these two biases, and refer to $r_1 = 1 - b_1$ and $r_2 = 1 - b_2$ as the rates of the respective random variables. Denote by $\bar{X}$ the sum of random variables drawn from the remaining $n - 2$ distributions. Let Q_1 be the probability that $\bar{X} = k - 1$ (equivalently, exactly one unit is left available for the last two distributions), Q_2 be the probability that $\bar{X} = k - 2$ (i.e., two units are left available), and $Q_{\geq 3}$ be the probability that $\bar{X} \leq k - 3$ (i.e., more than two units are left available). Finally, recall that X is the sum of random variables drawn from all distributions.

The following two claims enable the proof of Lemma 5.

Claim 1. *The problem of minimizing $\mu_k(X)$ as a function of r_1 and r_2 subject to $\delta_k(X) = \phi_k^*$ for a constant value ϕ_k^* is captured by the following program:*

$$\begin{aligned} & \text{Maximize } (Q_2 + Q_{\geq 3}) \cdot (r_1 + r_2) + Q_1 r_1 r_2 \quad (\text{Min-Revenue}) \\ & \text{subject to } r_1, r_2 \in [0, 1]^2 \text{ and } r_1 r_2 \cdot (Q_1 - Q_2) \\ & \quad + (r_1 + r_2) \cdot Q_2 = \phi_k^* - Q_{\geq 3}. \end{aligned} \quad (9)$$

Claim 2. *There always exists an optimal solution for the (Min-Revenue) program that satisfies $r_1 = r_2$, $r_1 \in \{0, 1\}$, or $r_2 \in \{0, 1\}$. Moreover, when $(Q_2)^2 - Q_{\geq 3}(Q_1 - Q_2) > 0$, the unique optimal solution satisfies $r_1 = r_2$.*

Using the two claims, we can directly provide the proof of the lemma.

Proof of Lemma 5. Among all the optimal solutions for problem (GENBERN : $\min\mu$ s.t. $\delta = \phi$), consider the one that satisfies the following conditions. First, it has the fewest Bernoulli variables that are nondegenerate. Second, among those, it has the smallest difference between the largest nondegenerate bias and the smallest nondegenerate bias; call these biases h and s accordingly. Third, among those, it has the smallest number of variables with bias that is equal to either h or s . We show, by contradiction, that $h = s$, establishing that all nondegenerate biases are equal.

Among all the optimal solutions, we select the one satisfying these criteria, and we select two distributions with bias $b_1 = h$ and $b_2 = s$, respectively. We apply Claim 1 with these two and express (GENBERN : $\min\mu$

s.t. $\delta = \phi$) as a function of $r_1 = 1 - b_1$, $r_2 = 1 - b_2$, and quantities Q_1 , Q_2 , $Q_{\geq 3}$ that are independent of r_1 and r_2 . Claim 2 establishes that there exists another optimal solution with all the other biases the same, and the biases b'_1 and b'_2 corresponding to the two distributions either satisfy $b'_1 = b'_2$ or have one them degenerate. The latter is a contradiction as we assumed that the solution has the smallest number of nondegenerate Bernoulli distributions. This means that $b'_1 = b'_2$. However, in order for the new solution to induce the same $\delta_k(\cdot)$ and $\mu_k(\cdot)$, it means that $b_2 \leq b'_1 = b'_2 \leq b_1$. Unless $b_1 = b_2$, we have, therefore, identified a new optimal solution with a smaller number of variables with bias equal to h or s , which again would induce contradiction. As a result, $b_1 = b_2$, and because $b_1 = h$ and $b_2 = s$, this means that $h = s$ and that all nondegenerate biases are equal. $\square$

5.2.2. Expressing (GENBERN : $\min \mu$ s.t. $\delta = \phi$) as a Function of Two Biases (Claim 1).

Proof of Claim 1. We first write the utility component $\delta_k(X)$ (i.e., the probability that not all k items are sold in terms of the probabilities Q_1 , Q_2 , and $Q_{\geq 3}$ and the biases r_1 and r_2 :

$$\delta_k(X) = Q_{\geq 3} + Q_2(r_1 + r_2 - r_1 r_2) + Q_1 r_1 r_2.$$

Applying this equation in the constraint of optimization problem (GENBERN : $\min \mu$ s.t. $\delta = \phi$) leads to Constraint (9) of the (Min-Revenue) optimization problem.

Our goal is now to minimize the objective of (GENBERN : $\min \mu$ s.t. $\delta = \phi$), which corresponds to the revenue component $\mu_k(X)$ (i.e., the expected fraction of items sold). Because all but biases b_1 and b_2 are kept constant, this is equivalent to minimizing the number of items sold to the remaining two buyers. Our objective is, therefore, to minimize the following expression:

$$(r_1(1 - r_2) + r_2(1 - r_1)) \cdot (Q_1 + Q_2 + Q_{\geq 3}) \\ + (1 - r_1) \cdot (1 - r_2) \cdot (Q_1 + 2Q_2 + 2Q_{\geq 3}).$$

The first term corresponds to the contribution of the two remaining buyers when only one of the two realized random variables is nonzero; then, $\mu_k(X)$ is increased if there exists at least one item that is left available from the other buyers. The second term corresponds to the event that both buyers have nonzero realized random variables; in this case, $\mu_k(X)$ is increased by one if exactly one available item is left from the other buyers and two if at least two available items are left. Simplifying the expression leads to the following minimization objective:

$$(Q_1 + 2Q_2 + 2Q_{\geq 3}) - (Q_2 + Q_{\geq 3}) \cdot (r_1 + r_2) - Q_1 r_1 r_2.$$

Eliminating constant terms and negating the objective offer the maximization objective of the claim. $\square$

5.2.3. Auxiliary Geometric Interpretation of the Constraint Intersection with $[0, 1]^2$. We now provide a geometric fact about hyperbolas of a particular form that is useful in characterizing the optimal solutions of the Min-Revenue program.

Claim 3. Let $a, b \in \mathbb{R}$, and consider the hyperbola in the x, y plane given by $xy + a(x + y) = b$. Expressing it as $y = g(x)$ and considering the segment(s) of g that intersects the region $[0, 1]^2$, the following holds.

- When $a, b > 0$, and $2a + 1 - b > 0$, the segment of g intersecting $[0, 1]^2$ is convex and decreasing.
- When $a, b < 0$, and $2a + 1 - b < 0$, the segment of g intersecting $[0, 1]^2$ is concave and decreasing.

Proof. In both cases considered, $a^2 + b > 0$. In the first case, this is straightforward as $a, b > 0$. In the second case, $2a + 1 < b$, and adding a^2 , we obtain $a^2 + 2a + 1 < a^2 + b$, implying $a^2 + b > 0$. If $x = -a$, then it means that $ax = b$ and hence, $-a^2 = b$; this contradicts the fact that $a^2 + b > 0$.

As a result, $x \neq -a$, and the hyperbola can be expressed as $y = \frac{b - ax}{x + a}$. The first derivative of this function is $\frac{dy}{dx} = \frac{-a(x+a) - (b-ax)}{(x+a)^2} = \frac{(-b-a^2)}{(x+a)^2}$. This is negative as $a^2 + b > 0$, which establishes that y is decreasing in x . The second derivative of this function is $\frac{d^2y}{dx^2} = -2 \frac{(-b-a^2)}{(x+a)^3} = 2 \frac{a^2+b}{(x+a)^3}$. The convexity or concavity is determined by the sign of this derivative, which is determined by the sign of $x + a$. Hence, the resulting hyperbola has two segments: one convex for $x > -a$ and one concave for $x < -a$.

- In the first case, because both $x \geq 0$ and $a > 0$, this sign is positive, and the segment is convex.
- In the second case, because of the symmetry, if the convex segment intersects $[0, 1]^2$, it should also intersect it at $x = y$. Hence, the point of this intersection is given by the equation $x^2 + ax = b$ whose roots are $x = -a \pm \sqrt{a^2 + b}$. The negative root is $-a - \sqrt{a^2 + b} < 0$, and the positive root is $-a + \sqrt{a^2 + b} > -a + \sqrt{a^2 + 2a + 1} > 1$ (because $2a + 1 < b$). As a result, the convex segment does not intersect with $[0, 1]^2$, and if the hyperbola does intersect, this happens with its concave segment. $\square$

5.2.4. Characterizing the Optimal Solutions of the Min-Revenue Program (Claim 2).

Proof of Claim 2. We first start from two corner cases.

Ia. We start from the simplest case where $Q_2 = 0$. This transforms the Min-Revenue program to

$$\begin{aligned} &\text{maximize} && Q_{\geq 3}(r_1 + r_2) + Q_1 r_1 r_2 \\ &\text{subject to} && Q_1 r_1 r_2 = \phi_k^* - Q_{\geq 3}. \end{aligned}$$

In this case, depending on the sign of $\phi_k^* - Q_{\geq 3}$, the optimal solution satisfies $r_1 = r_2$, $r_1 \in \{0, 1\}$, or $r_2 \in \{0, 1\}$ (if $Q_1 \neq 0$) or is independent of r_1, r_2 (otherwise).

Ib. We then consider the case where $Q_1 = Q_2 > 0$ (where it also holds that $(Q_2)^2 - Q_{\geq 3}(Q_1 - Q_2) > 0$); then, Constraint (9) is a linear constraint symmetric in r_1 and r_2 , whereas the objective becomes equivalent to maximizing $r_1 r_2$. The optimum is, therefore, uniquely achieved at $r_1 = r_2$.

For the remaining cases, we assume that $Q_2 > 0$ and $Q_1 \neq Q_2$. We can, therefore, rewrite the objective in terms of $r_1 + r_2$, obtaining the following equivalent formulation:

$$\begin{aligned} & \text{Minimize } (r_1 + r_2) \cdot \frac{(Q_2)^2 - Q_{\geq 3}(Q_1 - Q_2)}{Q_1 - Q_2} \\ & \text{subject to } r_1, r_2 \in [0, 1]^2 \text{ and} \\ & \quad r_1 r_2 + (r_1 + r_2) \cdot \frac{Q_2}{Q_1 - Q_2} = \frac{\phi_k^* - Q_{\geq 3}}{Q_1 - Q_2}. \end{aligned}$$

The constraint can be rewritten as $r_1 r_2 + a(r_1 + r_2) = b$ for $a = \frac{Q_2}{Q_1 - Q_2}$ and $b = \frac{\phi_k^* - Q_{\geq 3}}{Q_1 - Q_2}$.

Observe that unless $r_1 = r_2 = 1$ is the only feasible solution (in which case the claim is already proven), it holds that $Q_1 - Q_2 + 2Q_2 > r_1 r_2(Q_1 - Q_2) + (r_1 + r_2)Q_2 = \phi_k^* - Q_{\geq 3} > 0$. (Note that for any feasible solution, we have $\phi_k^* = \delta_k(X) \geq Q_2 + Q_{\geq 3} > Q_{\geq 3}$, so $\phi_k^* - Q_{\geq 3} > 0$.) Hence, it holds that a , b , and $1 + 2a - b$ all have the same sign, and this is the sign of $Q_1 - Q_2$. By Claim 3, the hyperbola $r_2 = g(r_1)$ corresponding to the constraint is always decreasing and is (a) convex when $Q_1 > Q_2$ ($a > 0$) and (b) concave when $Q_1 < Q_2$ ($a < 0$); it is also always symmetric in r_1, r_2 .

IIa. If $Q_1 < Q_2$, which implies $(Q_2)^2 - Q_{\geq 3}(Q_1 - Q_2) > 0$, the multiplier of the objective is negative in this case; therefore, our goal is to maximize $r_1 + r_2$ subject to a concave and decreasing constraint. As a result, the optimum is uniquely determined on the line $r_1 = r_2$.

IIb. If $Q_1 > Q_2$ and $(Q_2)^2 - Q_{\geq 3}(Q_1 - Q_2) > 0$, the multiplier of $(r_1 + r_2)$ in the objective is positive; therefore, our goal is to minimize $(r_1 + r_2)$ subject to a convex and decreasing constraint. As a result, the optimum is uniquely determined on the line $r_1 = r_2$.

IIc. If $Q_1 > Q_2$ and $(Q_2)^2 - Q_{\geq 3}(Q_1 - Q_2) < 0$, the multiplier of $(r_1 + r_2)$ in the objective is negative; therefore, our goal is to maximize $(r_1 + r_2)$ subject to a convex and decreasing constraint. As a result, the optimum lies on the boundary of $[0, 1]^2$.

The last case is that $(Q_2)^2 - Q_{\geq 3}(Q_1 - Q_2) = 0$. Then, if the program is feasible, there is a feasible solution with $r_1 = r_2$ or with one of r_1, r_2 in $\{0, 1\}$ via a similar reasoning.

Finally, $(Q_2)^2 - Q_{\geq 3}(Q_1 - Q_2) = 0 > 0$ corresponds to cases I(b), II(a), and II(b). In all three cases, $r_1 = r_2$ is the unique optimal solution. $\square$

5.3. Reducing to the Poisson Distribution

Lemma 5 states that the optimal solution to $(\text{GENBERN} : \min \mu \text{ s.t. } \delta = \phi)$ is such that some of the biases are zero or one, and the remainder has equal bias $b_i = 1 - r$. We now show that, in fact, there cannot be any zero or one biases.

Lemma 6. *The optimization problem $(\text{GENBERN} : \min \mu \text{ s.t. } \delta = \phi)$ is minimized when all of the Bernoulli distributions have equal bias (and there are no degenerate distributions).*

Proof. By Lemma 5, we know that there exists an optimal solution that consists of Bernoulli distributions with bias 0, 1, or $1 - r$ for a fixed r . Consider any such optimal solution. We first observe that this solution cannot have k or more Bernoullis with bias 1; otherwise, $\delta_k(X) = 0 < \phi_k^*$. (This holds because ϕ_k^* is at least $1/2$ as it is the intended competitive ratio.)

Now, we pick two specific Bernoulli variables with unequal bias and reoptimize the objective over these, keeping the rest fixed. If there is a Bernoulli with bias 1 and another with bias $1 - r \neq 1$, we pick two such variables. Otherwise, we pick one variable with bias $1 - r$ and another with bias 0. We define Q_1 , Q_2 , and $Q_{\geq 3}$ as in Section 5.2. In either case, we argue that $(Q_2)^2 - Q_{\geq 3}(Q_1 - Q_2) > 0$. Then, applying Claim 2, we arrive at a contradiction to the claim that our initial solution was optimal.

Among the biases left fixed, let n_1 denote the number of Bernoullis with bias 1 and n_r denote the number with bias $1 - r$. As discussed previously, we note $n_1 \leq k - 2$. Let $k' = k - n_1 \geq 2$ denote the number of item units left available once the n_1 Bernoullis with bias 1 have each acquired an item. Recall that Q_1 is the probability that exactly one unit is left available for the two Bernoullis from the others. Rephrasing, Q_1 is the probability that of the remaining Bernoullis, exactly $k - 1$ take a nonzero value. Because n_1 Bernoullis take on a nonzero value with certainty, this means that among the n_r Bernoullis with bias $1 - r$, exactly $k' - 1$ take on a nonzero value where $k' = k - n_1$. Likewise, Q_2 is the probability that among the n_r Bernoullis with bias $1 - r$, exactly $k' - 2$ take on a nonzero value, and $Q_{\geq 3}$ is the probability that among the n_r Bernoullis with bias $1 - r$, at most $k' - 3$ take on a nonzero value. Let us also define Q_i for $i \geq 3$ as the probability that among the n_r Bernoullis with bias $1 - r$, exactly $k' - i$ take on a nonzero value. Observe that $k' \geq 2$ and thus we must have $Q_1, Q_2 > 0$.

We first claim that $k' \geq 3$. If not, then we have $Q_{\geq 3} = 0$, which implies $(Q_2)^2 - Q_{\geq 3}(Q_1 - Q_2) > 0$ and completes the proof. We can now compute the probabilities Q_i :

$$\begin{aligned} Q_1 &= \binom{n_r - 1}{k' - 1} (1 - r)^{k' - 1} r^{n_r - k'}, \\ Q_2 &= \binom{n_r - 1}{k' - 2} (1 - r)^{k' - 2} r^{n_r - k' + 1}, \quad \text{etc.} \end{aligned}$$

Let us denote by α the ratio between Q_2 and Q_1 :

$$\alpha = \frac{Q_2}{Q_1} = \frac{k' - 1}{n_r - k' + 1} \cdot \frac{r}{1 - r}.$$

We then observe

$$\begin{aligned} \frac{Q_3}{Q_2} &= \frac{k' - 2}{n_r - k' + 2} \cdot \frac{r}{1 - r} < \alpha \quad \text{and likewise,} \\ \frac{Q_{i+1}}{Q_i} &\leq \frac{k' - i}{n_r - k' + i} \cdot \frac{r}{1 - r} < \alpha \quad \forall i > 3 \text{ with } Q_i > 0. \end{aligned}$$

Here, the second inequality follows by noting that $\frac{k' - i}{n_r - k' + i}$ decreases with i .

We can, therefore, write $Q_{\geq 3} = \sum_{i \geq 3} Q_i < Q_1(\alpha^2 + \alpha^3 + \dots) = \alpha^2 Q_1 / (1 - \alpha)$. Putting these expressions together, we obtain

$$(Q_2)^2 - Q_{\geq 3}(Q_1 - Q_2) > \alpha^2 Q_1^2 - \alpha^2 \frac{Q_1}{1 - \alpha} (Q_1 - \alpha Q_1) = 0.$$

This completes the proof. $\square$

5.4. Concluding the Proof of Lemma 2

We are now ready to provide the proof of the main lemma of the section.

Proof of Lemma 2. By Lemma 3, $\phi_k(\mathcal{F})$ attains its lowest value when all the distributions are Bernoulli. By Lemma 5, all these Bernoulli distributions either have equal biases or are degenerate (have zero or one bias). By Lemma 6, there exists an optimal solution with no degenerate Bernoulli distributions. If there are finite Bernoulli distributions, we can always repeat the argument of Lemma 6 and obtain a solution with strictly higher objective and one more nonzero bias. As a result, there exists an optimal solution that consists of an infinite collection of Bernoulli distributions with the same bias. This establishes that the worst-case instance is the limit $n \rightarrow \infty$.

We now claim that the quantity $X_{\pi(k, \mathcal{F})}$ in this limit is a Poisson distribution. Let Y_n be the sum of n equal Bernoulli distributions with $\mu_n := \mu_n(Y_n) = \delta_n(Y_n) =: \delta_n$. Note that μ_n and δ_n are both decreasing functions of n by the argument and bounded in $[0, 1]$. Therefore, they converge to a single limit, $\mu^* = \delta^*$. Let λ_n denote $E[Y_n]$. We will show that the sequence $\{\lambda_n\}$ converges. Then, by the Poisson limit theorem, the limit of $\{Y_n\}$, namely $Y^* = X_{\pi(k, \mathcal{F})}$, is Poisson. This means that μ_n and δ_n , which are linear functions of the probability density function (pdf), also converge to $\mu(Y^*)$ and $\delta(Y^*)$, respectively. Because μ_n and δ_n are identical sequences, it holds that $\mu(Y^*) = \delta(Y^*)$, which completes the proof.

It remains to show that $\{\lambda_n\}$ converges; this does not follow immediately from the convergence of $\{\mu_n\}$ and $\{\delta_n\}$ because $\{\lambda_n\}$ may not evolve monotonically. We will show that for any $\epsilon > 0$, there exists an $N(\epsilon)$

such that for any $n_1, n_2 > N(\epsilon)$, $|\lambda_{n_1} - \lambda_{n_2}| \leq \epsilon$, from which our claim follows.

In the remainder of this proof, we use $\text{Bin}(n, \lambda)$ to denote the sum of n Bernoulli variables each with bias λ/n . We will use the following lemma that shows that two binomials with sufficiently different expectations will also have sufficiently different truncated expectations. The proof of the lemma appears at the end of this section.

Lemma 7. For any $\epsilon > 0$, any $\lambda, \lambda' < k$ with $|\lambda - \lambda'| > \epsilon$, and any $n \geq 2k$, define $\epsilon' := (1 - e^{-\epsilon/2})/4k$. Then, it holds that $|\mu(\text{Bin}(n, \lambda)) - \mu(\text{Bin}(n, \lambda'))| > 2\epsilon'$.

We now make the following observations.

1. Observe that for any $n > 0$, $\delta_n \geq 1/2$, which means that $X_n \leq k - 1$ with probability at least $1/2$, and the median is thus no more than $k - 1$. The mean and the median of a binomial differ by at most $\ln(2)$ (Hamza 1995), which implies that $\lambda_n < k$ for all n .

2. For any $\epsilon' > 0$, using the fact that $\{\mu_n\}$ converges, there exists $N_1(\epsilon')$ such that for any $n_1, n_2 > N_1(\epsilon')$, $|\mu_{n_1} - \mu_{n_2}| \leq \epsilon'$ (i.e., $|\mu(\text{Bin}(n_1, \lambda_{n_1})) - \mu(\text{Bin}(n_2, \lambda_{n_2}))| \leq \epsilon'$).

3. Using the fact that the pdf of the binomial distribution converges in ℓ_1 norm as the number of samples increases, whereas the bias stays the same, and that μ is a linear function of the pdf, we get that for any $\epsilon' > 0$, there exists $N_2(\epsilon')$ such that for every $\lambda < k$ and every $n_1, n_2 > N_2(\epsilon')$, $|\mu(\text{Bin}(n_1, \lambda)) - \mu(\text{Bin}(n_2, \lambda))| \leq \epsilon'$. (This is immediate for a fixed λ by the Poisson convergence theorem. We can then define $N_2(\epsilon')$ as the supremum over the corresponding λ -specific values. That this is finite can be seen by discretizing the set of all λ 's.)

4. Putting these together, for $n_1, n_2 \geq \max\{N_1(\epsilon'), N_2(\epsilon')\}$,

$$\begin{aligned} &|\mu(\text{Bin}(n_2, \lambda_{n_1})) - \mu(\text{Bin}(n_2, \lambda_{n_2}))| \\ &\leq |\mu(\text{Bin}(n_1, \lambda_{n_1})) - \mu(\text{Bin}(n_2, \lambda_{n_1}))| \\ &\quad + |\mu(\text{Bin}(n_1, \lambda_{n_1})) - \mu(\text{Bin}(n_2, \lambda_{n_2}))| \leq \epsilon' + \epsilon' = 2\epsilon'. \end{aligned}$$

5. Finally, given $\epsilon > 0$, we set $\epsilon' = (1 - e^{-\epsilon/2})/4k$ and $N(\epsilon) = \max\{N_1(\epsilon'), N_2(\epsilon'), 2k\}$. Lemma 7 then implies that for all $n_1, n_2 \geq N(\epsilon)$, $|\lambda_{n_1} - \lambda_{n_2}| \leq \epsilon$, which concludes the proof. $\square$

Proof of Lemma 7. Without loss of generality, let $\lambda' > \lambda$, and let $b = \lambda/n$ and $b' = \lambda'/n$ be the bias of the Bernoulli variables constituting $\text{Bin}(n, \lambda)$ and $\text{Bin}(n, \lambda')$, respectively. We will make correlated draws from the two distributions and use these to bound the difference between their truncated expectations. We first flip n coins with bias b each to produce an instantiation of $\text{Bin}(n, \lambda)$; call this variable X . We then consider each of the n coins that came up tails in the first experiment and flip them again with bias $(b' - b)/(1 - b)$ each, counting the total number of heads in the first

and second experiments together and denoting it X' ; X' is then an instantiation of $\text{Bin}(n, \lambda')$.

Now, we define $\mu := \mu(\text{Bin}(n, \lambda)) = \mathbb{E}[\min(X, k-1)]/k$ and $\mu' := \mu(\text{Bin}(n, \lambda')) = \mathbb{E}[\min(X', k-1)]/k$. Clearly, $\mu' > \mu$, and we can lower bound the difference between the two by $1/k$ times the probability of the event that $X < k-1$ and $X' > X$. Let us compute this probability.

First, we observe that $\mathbb{E}[X] = \lambda < k$, so $\Pr[X < k] \geq 1/2$. Then, conditioning on $X < k$ and recalling that $n \geq 2k$, we have at least $n/2$ coins flipped in the second experiment. The probability that all of the coins come up tails is at most $(1 - (b' - b))^{n/2} \leq (1 - \epsilon/n)^{n/2} \leq e^{-\epsilon/2}$.

Therefore, conditioning on $X < k$, the event $X' > X$ happens with probability at least $1 - e^{-\epsilon/2}$. Putting everything together, we get that $\mu' - \mu \geq (1 - e^{-\epsilon/2})/2k$. $\square$

6. Discussion

This paper provides an understanding of the performance of static pricing in multiunit prophet inequalities. We show a simple static pricing scheme that obtains a competitive ratio that adapts to the size of the supply. This is enabled by a clean revenue-utility decomposition, which allows us to move beyond the now dominant *balanced* approach to analyze prophet inequalities. We hope that this decomposition can further our understanding of prophet inequalities for more settings. In particular, it has already inspired follow-up work as our technique is a building block for multiunit prophet inequality guarantees under the random-order arrival model (Arnosti and Ma 2022).

An interesting direction that remains open is to extend our approach to the setting with multiple different items. For its single-unit version, the existing guarantees rely on the *balanced prices* approach, which cannot yield competitive ratios beyond $1/2$ (see Section 1.4 for a relevant discussion). It is interesting to provide improved guarantees for multi-item multiunit prophet inequalities.

Acknowledgments

The authors thank Seffi Naor for multiple valuable discussions that helped shape the paper's contributions, the anonymous review team at the 19th Conference on Web and Internet Economics (WINE 2021) and at *Operations Research* for their valuable feedback, and the Simons Institute for the Theory of Computing as part of the work was conducted during the Simons Semester on Data-Driven Decision Processes. A one-page abstract appeared at the 17th Conference on Web and Internet Economics (WINE 2022).

References

- Agrawal S, Devanur NR (2019) Bandits with global convex constraints and objective. *Oper. Res.* 67(5):1486–1502.
- Alaei S (2014) Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM J. Comput.* 43(2):930–972.
- Arnosti N, Ma W (2022) Tight guarantees for static threshold policies in the prophet secretary problem. *Oper. Res.*, ePub ahead of print December 29, <https://doi.org/10.1287/opre.2022.2419>.

- Babaioff M, Dughmi S, Kleinberg RD, Slivkins A (2015) Dynamic pricing with limited supply. *ACM Trans. Econom. Comput.* 3(1):4:1–4:26.
- Badanidiyuru A, Kleinberg R, Slivkins A (2018) Bandits with knapsacks. *J. ACM* 65(3):13.
- Balseiro SR, Brown DB, Chen C (2021) Dynamic pricing of relocating resources in large networks. *Management Sci.* 67(7):4075–4094.
- Banerjee S, Freund D, Lykouris T (2022) Pricing and optimization in shared vehicle systems: An approximation framework. *Oper. Res.* 70(3):1783–1805.
- Chawla S, Miller JB, Teng Y (2019) Pricing for online resource allocation: Intervals and paths. *Proc. Thirtieth Annual ACM-SIAM Sympos. Discrete Algorithms SODA 2019* (SIAM, Philadelphia), 1962–1981.
- Chawla S, Devanur NR, Holroyd AE, Karlin AR, Martin JB, Sivan B (2017) Stability of service under time-of-use pricing. *Proc. 49th Annual ACM SIGACT Sympos. Theory Comput.* (ACM, New York), 184–197.
- Dütting P, Kleinberg R (2015) Polymatroid prophet inequalities. *Algorithms—ESA 2015—23rd Annual Eur. Sympos. Proc. Lecture Notes Comput. Sci.*, vol. 9294 (Springer, Berlin, Heidelberg), 437–449.
- Dütting P, Feldman M, Kesselheim T, Lucier B (2020) Prophet inequalities made easy: Stochastic optimization by pricing nonstochastic inputs. *SIAM J. Comput.* 49(3):540–582.
- Feldman M, Gravin N, Lucier B (2014) Combinatorial auctions via posted prices. *Proc. Twenty-Sixth Annual ACM-SIAM Sympos. Discrete Algorithms* (SIAM, Philadelphia), 123–135.
- Hajiaghayi MT, Kleinberg R, Sandholm T (2007) Automated online mechanism design and prophet inequalities. *Proc. 22nd National Conf. Artificial Intelligence AAAI'07*, vol. 1 (AAAI Press, Palo Alto, CA), 58–65.
- Hamza K (1995) The smallest uniform upper bound on the distance between the mean and the median of the binomial and Poisson distributions. *Statist. Probab. Lett.* 23(1):21–25.
- Hoeffding W (1956) On the distribution of the number of successes in independent trials. *Ann. Math. Statist.* 27(3):713–721.
- Jiang J, Ma W, Zhang J (2022) Tight guarantees for multi-unit prophet inequalities and online stochastic knapsack. *Proc. Annual ACM-SIAM Sympos. Discrete Algorithms (SODA)* (SIAM, Philadelphia), 1221–1246.
- Jiang J, Ma W, Zhang J (2023) Tightness without counterexamples: A new approach and new results for prophet inequalities. Leyton-Brown K, Hartline JD, Samuelson L, eds. *Proc. 24th ACM Conf. Econom. Comput. EC 2023* (ACM, New York), 909.
- Kleinberg R, Weinberg SM (2012) Matroid prophet inequalities. *Proc. Forty-Fourth Annual ACM Sympos. Theory Comput.* (ACM, New York), 123–136.
- Krengel U, Sucheston L (1977) Semiamarts and finite values. *Bull. Amer. Math. Soc.* 83(4):745–747.
- Lucier B (2017) An economic view of prophet inequalities. *ACM SIGecom Exchanges* 16(1):24–47.
- Rubinstein A, Singla S (2017) Combinatorial prophet inequalities. *Proc. Twenty-Eighth Annual ACM-SIAM Sympos. Discrete Algorithms SODA 2017* (SIAM, Philadelphia), 1671–1687.
- Samuel-Cahn E (1984) Comparison of threshold stop rules and maximum for independent nonnegative random variables. *Ann. Probab.* 12(4):1213–1216.

Shuchi Chawla holds an endowed professorship in computer science at The University of Texas at Austin. Shuchi is a theoretical computer scientist specializing in the areas of algorithm design and the interface between economics, optimization, and algorithms.

Nikhil Devanur is a principal scientist at Amazon working on sponsored products. He is interested in what he calls automated economics, which studies the question of how technology can be used to improve the efficiency of economic systems.

Theodoris Lykouris is the Mitsui Career Development Assistant Professor and an assistant professor of operations management at the Massachusetts Institute of Technology Sloan School of Management. His research focuses on data-driven sequential decision making and spans across the areas of machine learning, dynamic optimization, and economics.