Static Pricing for Multi-unit Prophet Inequalities (Extended Abstract)

Shuchi Chawla[1], Nikhil R. Devanur[2], and Thodoris Lykouris[3](✉)

[1] The University of Texas at Austin, Austin, USA
shuchi@cs.utexas.edu
[2] Amazon Research, Sunnyvale, USA
[3] Massachusetts Institute of Technology, Cambridge, USA
lykouris@mit.edu

The *prophet inequality* problem constitutes one of the cornerstones of online decision-making. A designer knows a set of n distributions from which random variables 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 showed that a simple static threshold policy achieves the optimal competitive ratio for this problem. Samuel-Cahn's algorithm determines a threshold p such that the probability that there exists a realization exceeding the threshold is exactly $\frac{1}{2}$, and then accepts the first random variable that exceeds the threshold. This algorithm achieves a competitive ratio of $\frac{1}{2}$ 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 remained open: where the designer is allowed to accept $k > 1$ random variables for some small value of k. This is called the *multi-unit* prophet inequality. When k is relatively large, then it is known that static threshold policies can achieve a competitive ratio of $1 - \Theta\left(\sqrt{\frac{\log(k)}{k}}\right)$ which goes to 1 as $k \to \infty$, and this ratio is asymptotically tight. However, (for example,) for $k = 2$ or 3, prior to our work, the best known competitive ratio of static thresholds remained $\frac{1}{2}$. Our work addresses this gap by answering the following question: *Can a static threshold policy achieve a better competitive ratio than* $\frac{1}{2}$ *for small* $k = 2, 3, \ldots$*?*

We develop an algorithm for finding a static threshold policy for the multi-unit prophet inequality that is sensitive to the supply k. Our algorithm is simple and practical. For any fixed price p, it estimates two statistics: (1) the fraction of items expected to be sold at that price, $\mu_k(p)$, and (2) the probability that not all units will sell out before all the customers have been served, $\delta_k(p)$. We then pick the static price p at which these two quantities are equal: $\mu_k(p) = \delta_k(p)$; this generalizes Samuel-Cahn's algorithm and proof via a min-max approach and shows that the worst-case competitive ratio is attained for Poisson distributions.

M. Feldman et al. (Eds.): WINE 2021, LNCS 13112, pp. 545–546, 2022.
https://doi.org/10.1007/978-3-030-94676-0

The competitive ratio of our policy for $k = 2, \cdots, 5$ is $0.585, 0.630, 0.660$, and 0.682 respectively, and scales as $1 - \Omega(\sqrt{\log k/k})$ for large k.

The full version can be found here: https://arxiv.org/abs/2007.07990.