Prophet Inequalities for Cost Minimization

Vasilis Livanos

joint work with Ruta Mehta

University of Illinois Urbana-Champaign

ACAC 2022

August 25th, 2022

Motivation

- Want to sell an orange. We see n buyers *sequentially*.
- Buyer i has private valuation v_i. How to offer prices?
  - Option 1: Run an auction. Meh.

Motivation

- Want to sell an orange. We see n buyers *sequentially*.
- Buyer i has private valuation v_i. How to offer prices?
  - Option 1: Run an auction. Meh.
  - Option 2: Become a grocer!
- Plan:
  1. Set price T.
  2. Leave store.
  3. ???
  4. Profit.

Prophet Inequality

- Worst-case order + unknown v_i's = Can't do anything.
- Random order + unknown v_i's = *Secretary problem*.
- Worst-case order + some knowledge of v_i's = *Prophet inequality*.

Prophet Inequality

- Worst-case order + unknown v_i's = Can't do anything.
- Random order + unknown v_i's = *Secretary problem*.
- Worst-case order + some knowledge of v_i's = *Prophet inequality*.
- n random variables $X_1, \ldots, X_n \sim D_1, \ldots, D_n$ arriving in *adversarial* order.
- Step $i \implies$ observe realization x_i.
  - Accept $x_i \implies$ Game ends.
  - Reject $x_i \implies$ Step $i + 1$.
- Goal: Select $\max_i x_i$.

Let's Play

$U[2,4]$	$U[2,4]$	$U[1,5]$	$U[0,7]$

Let's Play

$U[2,4]$	$U[2,4]$	$U[1,5]$	$U[0,7]$

- $x_1 = 2.74$

Let's Play

$U[2,4]$ $U[2,4]$ $U[1,5]$ $U[0,7]$

- $x_1 = 2.74$
- $x_2 = 3.75$

Let's Play

$U[2,4]$ $U[2,4]$ $U[1,5]$ $U[0,7]$

- $x_1 = 2.74$
- $x_2 = 3.75$
- $x_3 = 2.81$

Let's Play

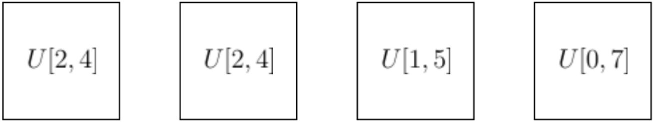

- $x_1 = 2.74$
- $x_2 = 3.75$
- $x_3 = 2.81$
- $x_4 = 5.66$

Prophet Inequality

- $\exists$ algorithm $\mathcal{A}$ s.t. $\mathbb{E}[\mathcal{A}] \geq \frac{1}{2}\,\mathbb{E}\left[\max_{i=1}^{n} X_i\right]$, and this is tight [KS77].
- Many algorithms achieve $1/2$.

$\mathcal{A}_T$: "Fixed-Threshold" Algorithm

Set threshold T based on $D_1, \ldots, D_n$. Accept first $x_i \geq T$.

- How to choose T?

Prophet Inequality

- $\exists$ algorithm $\mathcal{A}$ s.t. $\mathbb{E}[\mathcal{A}] \geq \frac{1}{2}\,\mathbb{E}\left[\max_{i=1}^{n} X_i\right]$, and this is tight [KS77].
- Many algorithms achieve $1/2$.

$\mathcal{A}_T$: "Fixed-Threshold" Algorithm

Set threshold T based on $D_1, \ldots, D_n$. Accept first $x_i \geq T$.

- How to choose T?
  1. Set $T =$ median of the distribution of $\max_{i=1}^{n} X_i$, i.e. $\Pr[\max_{i=1}^{n} X_i \geq T] = \frac{1}{2}$ [Sam84].
  2. Set $T = \frac{1}{2}\,\mathbb{E}[\max_{i=1}^{n} X_i]$ [KW12].

1/2 is Tight

$$X_1 = 1 \quad \text{w.p. } 1, \text{ and } X_2 = \begin{cases} \frac{1}{\varepsilon} & \text{w.p. } \varepsilon \\ 0 & \text{w.p. } 1-\varepsilon \end{cases}.$$

For every algorithm $\mathcal{A}$, $\mathbb{E}[\mathcal{A}] = 1$.
Prophet:

$$\mathbb{E}[\max_i X_i] = \frac{1}{\varepsilon} \cdot \varepsilon + 1 \cdot (1-\varepsilon) = 2 - \varepsilon.$$

Cost Prophet Inequality

- If *we* are buying? Same problem?
- Cost Prophet Inequality: Select $\min_i x_i$ subject to selecting at least one i.

Cost Prophet Inequality

- If *we* are buying? Same problem?
- Cost Prophet Inequality: Select $\min_i x_i$ subject to selecting at least one i.

$$X_1 = 1 \quad \text{w.p. } 1, \text{ and } X_2 = \begin{cases} 0 & \text{w.p. } 1 - 1/L \\ L & \text{w.p. } 1/L \end{cases}.$$

For every algorithm $\mathcal{A}$, $\mathbb{E}[\mathcal{A}] = 1$.
Prophet:

$$\mathbb{E}[\min_i X_i] = 1 \cdot \frac{1}{L} + 0 \cdot \left(1 - \frac{1}{L}\right) = 1/L.$$

What *can* we do?

- Focus on I.I.D. $X_1, \ldots, X_n \sim D$.
- Single threshold not good enough*, need multiple thresholds.

Theorem 1: Optimal Threshold Algorithm $\mathcal{A}$

Let $G(i)$ be $\mathcal{A}$'s expected value, when it sees $X_i, \ldots, X_n$.
The algorithm $\mathcal{A}$ which sets $\tau_i = G(i+1)$ and accepts the first i such that $X_i \le \tau_i$ is optimal.

Detour: Hazard Rate

- Need a tool to classify different distributions.

Hazard Rate (aka Failure Rate)

For a distribution D with cdf F and pdf f, the *hazard rate* is defined as

$$h(x) \triangleq \frac{f(x)}{1 - F(x)}.$$

- Intuition: $h(x) = \Pr\left[X = x \mid X \geq x\right]$ (for discrete distributions).

Detour: Hazard Rate

- Need a tool to classify different distributions.

Hazard Rate (aka Failure Rate)

For a distribution D with cdf F and pdf f, the *hazard rate* is defined as

$$h(x) \triangleq \frac{f(x)}{1 - F(x)}.$$

- Intuition: $h(x) = \Pr[X = x \mid X \geq x]$ (for discrete distributions).
- h monotonically increasing $\implies$ Monotone Hazard Rate (MHR) distribution.
- MHR distributions don't have heavy tails.
- Important in revenue maximization via virtual valuations.

Detour: Cumulative Hazard Rate

- Cumulative Hazard Rate: $H(x) = \int_0^x h(u)\,du$.
- Idea: Study distributions with polynomial H; let's call them P_H. Approximate all* other distributions via polynomials.

$$H(x) = \sum_{i=1}^{k} a_i x^{d_i}, \qquad 0 < d_1 \leq \cdots \leq d_k.$$

- d_1 controls how heavy D's tail is.
- $d_1 \geq 1 \implies$ MHR distribution.

Fixed Distribution, Optimal Algorithm

Theorem 2

For every distribution $D \in P_H$ and I.I.D. random variables drawn from D, $\exists$ a $\lambda(d_1)$-competitive cost prophet inequality, where

$$\lambda(d_1) = \frac{(1 + 1/d_1)^{1/d_1}}{\Gamma(1 + 1/d_1)}.$$

Furthermore, this is tight for $H(x) = x^{d_1}$.

- Via Stirling's approximation

$$\lambda(d_1) \approx c \cdot e^{1/d_1}.$$

Positive or Negative Result?

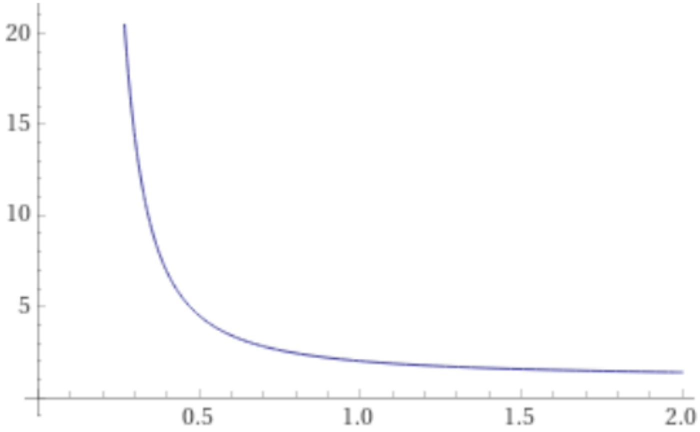

- Both! Ratio can be arbitrarily bad, but constant for every fixed D.

Special Case: MHR Distributions

- For $d_1 \geq 1 \implies$ distribution is MHR.
- Special case of "regular" distributions; exponential-like.
- $\lambda(d_1)$ decreasing in $d_1 \implies$

$$\lambda(d_1) \leq \frac{(1+1/1)^1}{\Gamma(1+1/1)} = 2,$$

  for all MHR distributions.
- 2 is tight for the exponential distribution.

Single Threshold

- Single threshold suffices for single-item classical prophet inequality!
- Impossible for cost prophet inequality.

Theorem 3

For every distribution $D \in P_H$ and I.I.D. random variables drawn from D, $\exists$ a single threshold T such that accepting the first i where $X_i \leq T$ yields a $O\left((\log n)^{1/d_1}\right)$-approximation to $\mathbb{E}[\min_i X_i]$.
Furthermore, this is tight for $H(x) = x^{d_1}$.

Summary

1. Optimal-threshold algorithm characterization for CPI.
2. Distribution-dependent constant for polynomial H.
3. Universal constant 2 for MHR distributions.
4. Poly-logarithmic ratio with a single threshold.

Open Questions

- Only use k-thresholds for $1 < k < n$. How does the ratio change?
- Get universal constant for subclass of distributions, like MHR; maybe *regular*?
- Only have sample access to $\mathcal{D}$, how does the ratio with the number of samples?
- Impossibility result does not apply in the *free order* setting. I.I.D. case is upper bound, but is a distribution-dependent constant-factor ratio possible?

QUESTIONS ?

References I

[KS77] Ulrich Krengel and Louis Sucheston. "Semiamarts and finite values". In: *Bull. Amer. Math. Soc.* 83.4 (July 1977), pp. 745–747. URL: https://projecteuclid.org:443/euclid.bams/1183538915.

[KW12] Robert Kleinberg and S. Matthew Weinberg. "Matroid Prophet Inequalities". In: *CoRR* abs/1201.4764 (2012). arXiv: 1201.4764. URL: http://arxiv.org/abs/1201.4764.

[Sam84] Ester Samuel-Cahn. "Comparison of Threshold Stop Rules and Maximum for Independent Nonnegative Random Variables". In: *The Annals of Probability* 12.4 (1984), pp. 1213–1216. ISSN: 00911798. URL: http://www.jstor.org/stable/2243359.