Combinatorial Optimization under Uncertainty and Prophet Inequalities

Vasilis Livanos

livanos3@illinois.edu

Department of Computer Science
University of Illinois Urbana-Champaign

Partially based on joint work with

Chandra Chekuri

Ruta Mehta

November 23rd, 2022

MATHEMATICAL GAMES

A fifth collection of "brain-teasers"

by Martin Gardner

Every eight months or so this department presents an assortment of short problems drawn from various mathematical fields. This is the fifth such collection. The answers to the problems will be given here next month. I welcome letters from readers who find fault with an answer, solve a problem more elegantly, or generalize a problem in some interesting way. In the past I have tried to avoid puzzles that play verbal pranks on the reader, so I think it only fair to say that several of this month's "brain-teasers" are touched with whimsy. They must be read with care; otherwise you may find the road to a solution blocked by an unwarranted assumption.

1.

Mel Stover of Winnipeg was the first to send this amusing problem—amusing because of the ease with which even the best of geometers may fail to approach it properly. Given a triangle with one obtuse angle, is it possible to cut the triangle into smaller triangles, all of them acute? (An acute triangle is a triangle with three acute angles. A right angle is of course neither acute nor obtuse.) If this cannot be done, give a proof of impossibility. If it can be done, what is the smallest number of acute triangles into which any obtuse triangle can be dissected?

The illustration at right shows a typical attempt that leads nowhere. The triangle has been divided into three acute triangles, but the fourth is obtuse, so nothing has been gained by the preceding cuts.

This delightful problem led me to ask myself: "What is the smallest number of acute triangles into which a square can be dissected?" For days I was convinced that nine was the answer; then suddenly I saw how to reduce it to eight. I wonder how many readers can discover an eight-triangle solution, or perhaps an even better one. I am unable to prove that eight is the minimum, though I strongly suspect that it is.

2.

In H. G. Wells's novel *The First Men in the Moon* our natural satellite is found to be inhabited by intelligent insect creatures who live in caverns below the surface. These creatures, let us assume, have a unit of distance that we shall call a "lunar." It was adopted because the moon's surface area, if expressed in square lunars, exactly equals the moon's volume in cubic lunars. The moon's diameter is 2,160 miles. How many miles long is a lunar?

3.

In 1958 John H. Fox, Jr., of the Minneapolis-Honeywell Regulator Co., and L. Gerald Marnie of the Massachusetts Institute of Technology devised an unusual betting game which they call Googol. It is played as follows: Ask someone to take as many slips of paper as he pleases, and on each slip write a different positive number. The numbers may range from small fractions of one to a number the size of a "googol" (1 followed by a hundred zeros) or even larger. These slips are turned face-down and shuffled over the top of a table. One at a time you turn the slips face-up. The aim is to stop turning when you come to the number that you guess to be the largest of the series. You cannot go back and pick a previously turned slip. If you turn over all the slips, then of course you must pick the last one turned.

Most people will suppose the odds

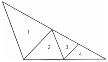

Can this triangle be cut into acute ones?

Secretary Problem

- *n unknown* values $x_1, \ldots, x_n$
- Random order
- Step i:
  1. Select x_i and stop
  2. Ignore x_i and continue

$$\Pr[\text{We select} \max_i x_i]?$$

Secretary Problem

x_1	x_2	$\cdots$	$x_{n/2}$

$x_{n/2+1}$	$\cdots$	x_{n-1}	x_n

S_1
Sampling Phase

S_2
Selection Phase

$$\left.\begin{array}{ll} \text{w.p. } 1/2, & x_1^* \in S_2 \\ \text{w.p. } 1/2, & x_2^* \in S_1 \end{array}\right\} \implies Pr[\text{We select} \max_i x_i] \geq 1/4.$$

Secretary Problem

x_1	x_2	$\dots$	$x_{n/2}$		$x_{n/2+1}$	$\dots$	x_{n-1}	x_n

S_1
Sampling Phase

S_2
Selection Phase

$$\left.\begin{array}{ll} \text{w.p. } 1/2, & x_1^* \in S_2 \\ \text{w.p. } 1/2, & x_2^* \in S_1 \end{array}\right\} \implies Pr[\text{We select} \max_i x_i] \geq 1/4.$$

- Can get $1/e$ *(optimal)* by sampling first n/e.

Prophet Inequality

What if we know *something* about the x_i's?
[Krengel, Sucheston and Garling '77]

$X_1, X_2, \ldots, X_n \sim$ (known) $\mathcal{D}_1, \mathcal{D}_2, \ldots, \mathcal{D}_n$
arrive in *adversarial* order.

- Design *stopping time* to maximize selected value.
- Compare against all-knowing *prophet*: $\mathbb{E}[\max_i X_i]$.

$\mathcal{U}[2,4]$ $\mathcal{U}[2,4]$ $\mathcal{U}[1,5]$ $\mathcal{U}[0,7]$

$\mathcal{U}[2,4]$ | $\mathcal{U}[2,4]$ | $\mathcal{U}[1,5]$ | $\mathcal{U}[0,7]$

$X_1 = 3.29$

$\mathcal{U}[2,4]$	$\mathcal{U}[2,4]$	$\mathcal{U}[1,5]$	$\mathcal{U}[0,7]$
$X_1 = 3.29$	$X_2 = 3.46$		

$\mathcal{U}[2,4]$ | $\mathcal{U}[2,4]$ | $\mathcal{U}[1,5]$ | $\mathcal{U}[0,7]$

$X_1 = 3.29$ | $X_2 = 3.46$ | $X_3 = 2.94$ |

$\mathcal{U}[2,4]$	$\mathcal{U}[2,4]$	$\mathcal{U}[1,5]$	$\mathcal{U}[0,7]$
$X_1 = 3.29$	$X_2 = 3.46$	$X_3 = 2.94$	$X_4 = 2.76$

Prophet Inequality [Krengel, Sucheston and Garling '77, '78]

$\exists$ stopping strategy that achieves $1/2 \cdot \mathbb{E}[\max_i X_i]$,
and this is tight.

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

$\mathbb{E}[\text{ALG}] = 1$ for all algorithms.

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

Prophet Inequality [Krengel, Sucheston and Garling '77, '78]

$\exists$ stopping strategy that achieves $1/2 \cdot \mathbb{E}[\max_i X_i]$,
and this is tight.

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

$\mathbb{E}[\text{ALG}] = 1$ for all algorithms.

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

- Idea: Set *threshold* T, accept first $X_i \geq T$.
  - $T : \Pr[\max_i X_i \geq T] = 1/2$ works [Samuel-Cahn '84].
  - $T = 1/2 \cdot \mathbb{E}[\max_i X_i]$ works [Kleinberg and Weinberg '12].

Why should we care?

We want to sell a banana to one of n buyers to maximize welfare.

- Option 1: Collect bids b_i, sell to highest bidder.

Why should we care?

We want to sell a banana to one of n buyers to maximize welfare.

- Option 1: Collect bids b_i, sell to highest bidder.

Why should we care?

We want to sell a banana to one of n buyers to maximize welfare.

- Option 1: Collect bids b_i, sell to highest bidder.
- Option 2: Become a grocer!

  Plan:

  1. Set price p.
  2. Leave store.
  3. ???
  4. Profit.

  Price $p \iff$ Threshold T in PI

Why should we care?

We want to sell a banana to one of n buyers to maximize welfare.

- Option 1: Collect bids b_i, sell to highest bidder.

- Option 2: Become a grocer!

  Plan:

  1. Set price p.
  2. Leave store.
  3. ???
  4. Profit.

  Price $p \iff$ Threshold T in PI

What about maximizing revenue?
Use "virtual valuations" to design T: $\phi(v) = v - \frac{1-F(v)}{f(v)}$.
[Myerson '81]

Why should we care?

- Posted prices apply when buyers arrive *online*.
- Lots of past work on this and extensions:
  - [Hajiaghayi, Kleinberg and Sandholm '07]
  - [Chawla, Hartline, Malec and Sivan '10]
  - [Alaei '11]
  - [Feldman, Gravin and Lucier '15]
  - [Dütting, Feldman, Kesselheim and Lucier '16]
  - [Correa, Foncea, Pizarro and Verdugo '19]
  - [Assadi, Kesselheim and Singla '21]

Proof

$X^* = \max_i X_i$
$p_i = \Pr[X^* = X_i] \implies \sum_i p_i = 1.$

- τ_i: $\Pr[X_i \geq \tau_i] = p_i$
- $v_i(p_i) := \mathbb{E}[X_i \mid X_i \geq \tau_i]$

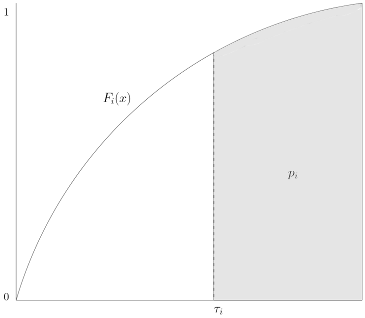

Proof

$X^* = \max_i X_i$
$p_i = \Pr[X^* = X_i] \implies \sum_i p_i = 1.$

- τ_i: $\Pr[X_i \geq \tau_i] = p_i$
- $v_i(p_i) := \mathbb{E}\left[X_i \mid X_i \geq \tau_i\right]$
- $\mathbb{E}[X^*] \leq \sum_i v_i(p_i) \cdot p_i$,

  since $X^* \sim \mathcal{D}^*$ with marginals $\boldsymbol{p}$.

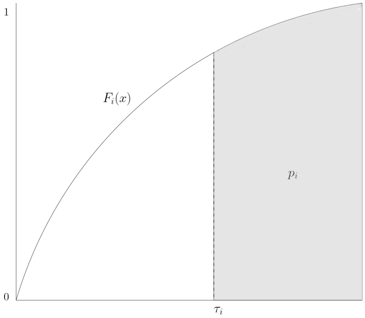

Proof

$X^* = \max_i X_i$
$p_i = \Pr[X^* = X_i] \implies \sum_i p_i = 1.$

- τ_i: $\Pr[X_i \geq \tau_i] = p_i$
- $v_i(p_i) := \mathbb{E}\left[X_i \mid X_i \geq \tau_i\right]$
- $\mathbb{E}[X^*] \leq \sum_i v_i(p_i) \cdot p_i$,
  since $X^* \sim \mathcal{D}^*$ with marginals $\boldsymbol{p}$.

$$\begin{aligned} \max \quad & \sum_i v_i(z_i) \cdot z_i \\ \text{s.t.} \quad & \sum_i z_i \leq 1 \qquad\qquad (1) \\ & 0 \leq z_i \leq p_i \quad \forall i \end{aligned}$$

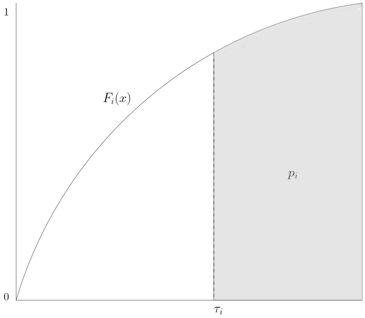

Idea

Reject every random variable X_i w.p. $1/2$.
Otherwise accept i iff $X_i \geq \tau_i$ (happens w.p. p_i).

$$\mathbb{E}[ALG] = \sum_i \Pr[\text{We reach } i] \cdot 1/2 \cdot p_i \cdot v_i(p_i)$$

Idea

Reject every random variable X_i w.p. $1/2$.
Otherwise accept i iff $X_i \geq \tau_i$ (happens w.p. p_i).

$$\mathbb{E}[ALG] = \sum_i \Pr[\text{We reach } i] \cdot 1/2 \cdot p_i \cdot v_i(p_i)$$

By a union bound,

$$\Pr[\text{We reach } i] \geq \Pr[\text{We pick nothing}] \geq 1 - \sum_i \frac{p_i}{2} \geq \frac{1}{2}.$$

- $1/4$-approximation to $\mathbb{E}[X^*]$.

Idea

Reject every random variable X_i w.p. $1/2$.
Otherwise accept i iff $X_i \geq \tau_i$ (happens w.p. p_i).

$$\mathbb{E}[ALG] = \sum_i \Pr[\text{We reach } i] \cdot 1/2 \cdot p_i \cdot v_i(p_i)$$

By a union bound,

$$\Pr[\text{We reach } i] \geq \Pr[\text{We pick nothing}] \geq 1 - \sum_i \frac{p_i}{2} \geq \frac{1}{2}.$$

- $1/4$-approximation to $\mathbb{E}[X^*]$.

Rewrite

$$\mathbb{E}[ALG] \geq \sum_i r_i \cdot q_i \cdot p_i \cdot v_i(p_i)$$

Can we ensure $r_i \cdot q_i = 1/2$?

- $r_1 = 1 \implies q_1 = 1/2.$ Then $r_{i+1} = r_i\,(1 - q_i p_i)$
- If we set $q_i = \frac{1}{2r_i} \implies r_{i+1} = r_i - \frac{p_i}{2} = 1 - \sum_{j \leq i} \frac{p_i}{2} \geq \frac{1}{2}$

How to generalize this?

$$
\begin{aligned}
\max \quad & \sum_i v_i(z_i) \cdot z_i \\
\text{s.t.} \quad & \sum_i z_i \leq 1 \\
& z_i \geq 0 \quad \forall i
\end{aligned}
\quad \Longrightarrow \quad
\begin{aligned}
\max \quad & \sum_i w_i \cdot z_i \\
\text{s.t.} \quad & \boldsymbol{z} \in \mathcal{P}(\mathcal{M}) \\
& 0 \leq z_i \leq 1 \quad \forall i
\end{aligned}
\qquad (2)
$$

- Let $\boldsymbol{x}$ be an optimal solution to (2). How should we round $\boldsymbol{x}$?

How to generalize this?

$$\begin{aligned} \max \quad & \sum_i v_i(z_i) \cdot z_i \\ \text{s.t.} \quad & \sum_i z_i \leq 1 \\ & z_i \geq 0 \quad \forall i \end{aligned} \quad \Longrightarrow \quad \begin{aligned} \max \quad & \sum_i w_i \cdot z_i \\ \text{s.t.} \quad & \boldsymbol{z} \in \mathcal{P}(\mathcal{M}) \\ & 0 \leq z_i \leq 1 \quad \forall i \end{aligned} \qquad (2)$$

- Let $\boldsymbol{x}$ be an optimal solution to (2). How should we round $\boldsymbol{x}$?

Attempt #1

Create random set R where $i \in R$ independently w.p. x_i *(active elements)*.

- 👍 $\mathbb{E}[\sum_{i \in R} w_i] = \sum_i w_i \cdot x_i$
- 👎 R might be infeasible

How to generalize this?

Attempt #2: Contention Resolution Scheme (CRS)

1. Create random set R where $i \in R$ independently w.p. x_i.
2. Drop elements from R to create feasible S.

- [Chekuri, Vondrák and Zenklusen '11].
- CRS is c-selectable if

$$\Pr[i \in S \mid i \in R] \geq c \qquad \forall i$$

- If CRS is c-selectable $\implies$ c-approximation.
- Observe R offline $\implies$ $(1 - 1/e)$-selectable CRS for single item.
- Combine in black-box way for general constraints/objectives.

Online Contention Resolution Schemes

Online Contention Resolution Scheme (OCRS)

The elements of R (active elements) are revealed to the algorithm one by one in adversarial order.

- [Alaei '11, Feldman, Svensson and Zenklusen '15].
- $\exists\, {}^1\!/_2$-selectable OCRS for single item (tight). $R = \{i \mid X_i \geq \tau_i\}$.
  Recall we guaranteed $\Pr[i \in S \mid i \in R] = r_i \cdot q_i = {}^1\!/_2$.

Why care about OCRSs?

Online Contention Resolution Schemes

Online Contention Resolution Scheme (OCRS)

The elements of R (active elements) are revealed to the algorithm one by one in adversarial order.

- [Alaei '11, Feldman, Svensson and Zenklusen '15].
- $\exists$ $1/2$-selectable OCRS for single item (tight). $R = \{i \mid X_i \geq \tau_i\}$.
  Recall we guaranteed $\Pr[i \in S \mid i \in R] = r_i \cdot q_i = 1/2$.

Why care about OCRSs?

- c-selectable OCRS for $\mathcal{C} \implies c$-approximation to Prophet Inequality with constraint $\mathcal{C}$.
- c-approximation to (slight variant of) Prophet Inequality for $\mathcal{C} \implies c$-selectable OCRS for $\mathcal{C}$.
  Can use prophet inequalities to design optimal OCRSs!

What does the adversary know?

- Offline: Nothing.
  $1/2$-OCRS for single item, $1/2$-OCRS for matroids.

Adversaries

What does the adversary know?

- Offline: Nothing.
  $1/2$-OCRS for single item, $1/2$-OCRS for matroids.
- Online: Same information as the Algorithm at every step.
  $1/2$-OCRS for single item, $1/4$-OCRS for matroids.

Adversaries

What does the adversary know?

- Offline: Nothing.
  $1/2$-OCRS for single item, $1/2$-OCRS for matroids.
- Online: Same information as the Algorithm at every step.
  $1/2$-OCRS for single item, $1/4$-OCRS for matroids.
- Almighty: All realizations and randomness of the Algorithm.
  $1/4$-OCRS for single item, $1/4$-OCRS for matroids.

Adversaries

What does the adversary know?

- Offline: Nothing.
  $1/2$-OCRS for single item, $1/2$-OCRS for matroids.
- Online: Same information as the Algorithm at every step.
  $1/2$-OCRS for single item, $1/4$-OCRS for matroids.
- Almighty: All realizations and randomness of the Algorithm.
  $1/4$-OCRS for single item, $1/4$-OCRS for matroids.

Greedy OCRS (Informal)

Decides (randomly) which elements to select *before* it sees R.

- Works against almighty adversary.

Adversaries

What does the adversary know?

- Offline: Nothing.
  $1/2$-OCRS for single item, $1/2$-OCRS for matroids.
- Online: Same information as the Algorithm.
  $1/2$-OCRS for single item, $1/4$-OCRS for matroids.
- Almighty: All realizations and randomness of the Algorithm.
  $1/4$-OCRS for single item, $1/4$-OCRS for matroids.

Greedy OCRS (Formal)

Create $\mathcal{F}_{\boldsymbol{x}} \subseteq \mathcal{I}$ before seeing R. When element i arrives, greedily select i iff $i \in R$ & $S_{i-1} + i \in \mathcal{F}_{\boldsymbol{x}}$.

What does the adversary know?

- Offline: Nothing.
  $1/2$-OCRS for single item, $1/2$-OCRS for matroids.
- Online: Same information as the Algorithm.
  $1/2$-OCRS for single item, $1/4$-OCRS for matroids.
- Almighty: All realizations and randomness of the Algorithm.
  ~~1/4~~ $\Longrightarrow$ $1/e$ -OCRS for single item, $1/4$-OCRS for matroids.

Greedy OCRS (Formal)

Create $\mathcal{F}_{\boldsymbol{x}} \subseteq \mathcal{I}$ before seeing R. When element i arrives, greedily select i iff $i \in R$ & $S_{i-1} + i \in \mathcal{F}_{\boldsymbol{x}}$.

Theorem [L. '22]

$\exists 1/e$-selectable Greedy OCRS for single items, and this is the best possible.

Recall $\boldsymbol{x}$ optimal solution to LP and $\sum_i x_i \leq 1$ (single item).

Recall $\boldsymbol{x}$ optimal solution to LP and $\sum_i x_i \leq 1$ (single item).

Idea

Create set T where $i \in T$ independently w.p. $\frac{1-e^{-x_i}}{x_i}$.
Greedily select i if $i \in R \cap T$.

- Simulates "splitting" i into many small elements.

$$
\begin{aligned}
\Pr[i \text{ is selected}] &= \Pr[i \in T] \cdot \prod_{j<i} (1 - \Pr[j \text{ is selected}]) \\
&\geq \frac{1 - e^{-x_i}}{x_i} \prod_{j<i} \left(1 - x_j \cdot \frac{1 - e^{-x_j}}{x_j}\right) \\
&= \frac{1 - e^{-x_i}}{x_i} \prod_{j<i} e^{-x_j} \\
&= \frac{1 - e^{-x_i}}{x_i} \cdot e^{-\sum_{j<i} x_j} \\
&\geq \frac{\left(1 - e^{-x_i}\right) e^{x_i - 1}}{x_i} \qquad (1)
\end{aligned}
$$

(1) is minimized for $x_i \to 0 \implies {}^1/_e$.

- Worst-case is $n \to \infty$ and $x_i \to 0 \ \forall i$.
- Idea extends to partition and transversal matroids.

Variants

What if...

- arrival order is *random*?

  Prophet Secretary: $(1 - 1/e)$-ROCRS and ≈ 0.669-PI.

  [Esfandiari, Hajiaghayi, Liaghat and Monemizadeh '15]
  [Correa, Saona and Ziliotto '20]

- arrival order is *chosen*?

  Free-Order: $(1 - 1/e)$-CRS and ≈ 0.7258-PI.

  [Bubna and Chiplunkar '22]

- $X_1, \ldots, X_n \overset{\text{i.i.d.}}{\sim} \mathcal{D}$?

  I.I.D.: $(1 - 1/e)$-CRS and ≈ 0.745-PI.

  [Hill and Kertz '82]
  [Correa, Foncea, Hoeksma, Oosterwijk and Vredeveld '21]

Extensions to Multiple Items

1. <u>k-Uniform Matroid:</u>
   - ✓ $1 - O\left(1/\sqrt{k}\right)$-OCRS
   - ✓ $1 - O\left(\sqrt{\log k/k}\right)$-Greedy OCRS

Extensions to Multiple Items

1. <u>k-Uniform Matroid:</u>
   - ✓ $1 - O\left(1/\sqrt{k}\right)$-OCRS
   - ✓ $1 - O\left(\sqrt{\log k/k}\right)$-Greedy OCRS
2. <u>General Matroid:</u>
   - ✓ $1/2$-OCRS
   - ? $1/4$-Greedy OCRS

Extensions to Multiple Items

1. *k*-Uniform Matroid:
   - ✓ $1 - O\left({}^{1}/\sqrt{k}\right)$-OCRS
   - ✓ $1 - O\left(\sqrt{{}^{\log k}/_{k}}\right)$-Greedy OCRS
2. General Matroid:
   - ✓ $^{1}/_{2}$-OCRS
   - ? $^{1}/_{4}$-Greedy OCRS
3. Bipartite Matching:
   - ? 0.349-OCRS
   - ? $^{1}/_{2e} \approx 0.184$-Greedy OCRS

Extensions to Multiple Items

1. *k-Uniform Matroid:*
   - ✓ $1 - O\left(1/\sqrt{k}\right)$-OCRS
   - ✓ $1 - O\left(\sqrt{\log k / k}\right)$-Greedy OCRS
2. General Matroid:
   - ✓ $1/2$-OCRS
   - ? $1/4$-Greedy OCRS
3. Bipartite Matching:
   - ? 0.349-OCRS
   - ? $1/2e \approx 0.184$-Greedy OCRS
4. General Matching:
   - ? 0.344-OCRS
   - ? $1/2e \approx 0.184$-Greedy OCRS

Extensions to Multiple Items

1. k-Uniform Matroid:
   - ✓ $1 - O\left(1/\sqrt{k}\right)$-OCRS
   - ✓ $1 - O\left(\sqrt{\log k/k}\right)$-Greedy OCRS
2. General Matroid:
   - ✓ $1/2$-OCRS
   - ? $1/4$-Greedy OCRS
3. Bipartite Matching:
   - ? 0.349-OCRS
   - ? $1/2e \approx 0.184$-Greedy OCRS
4. General Matching:
   - ? 0.344-OCRS
   - ? $1/2e \approx 0.184$-Greedy OCRS
5. Results extend also to submodular objective functions.

Cost Minimization

- Objective: Minimize selected value, compare against $\mathbb{E}[\min_i X_i]$.
- Forced to select an element $\implies$ upwards-closed constraint.

Cost Minimization

- Objective: Minimize selected value, compare against $\mathbb{E}[\min_i X_i]$.
- Forced to select an element $\implies$ upwards-closed constraint.
- No bounded approximation for adversarial or random order!

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

$$\frac{\mathbb{E}[ALG]}{\mathbb{E}[\min\{X_1, X_2\}]} = \frac{1}{\varepsilon}$$

- What about I.I.D.?
  Intuition:
  Set $T = 2 \cdot \mathbb{E}[\min_i X_i]$.

Cost Minimization

- Objective: Minimize selected value, compare against $\mathbb{E}[\min_i X_i]$.
- Forced to select an element $\implies$ upwards-closed constraint.
- No bounded approximation for adversarial or random order!

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

$$\frac{\mathbb{E}[ALG]}{\mathbb{E}[\min\{X_1, X_2\}]} = \frac{1}{\varepsilon}$$

- What about I.I.D.?
  ~~Intuition~~ False Intuition:
  Set $T = 2 \cdot \mathbb{E}[\min_i X_i]$.
  - Doesn't work! $\Pr[\text{We are forced to select } X_n] \to 1$.
  - Optimal single threshold $T \implies \Theta(\text{polylog } n)$-approximation.

Is Cost Minimization hopeless?

Analyze the optimal algorithm. Set τ_i, accept first $X_i \leq \tau_i$.
Intuition: $\tau_i = \mathbb{E}[\mathrm{OPTALG}_{i+1,\ldots,n}]$. How to analyze it?

Is Cost Minimization hopeless?

Analyze the optimal algorithm. Set τ_i, accept first $X_i \leq \tau_i$.
Intuition: $\tau_i = \mathbb{E}[\mathrm{OPTALG}_{i+1,\ldots,n}]$. How to analyze it?

Idea

Look at "fatness" of $\mathcal{D}$'s tail. Captured by $\mathcal{D}$'s *Hazard Rate*.

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

MHR Distribution

h is increasing.

- Important subclass, lots of past work by economists. Good guarantees in many applications (e.g. revenue maximization in auctions).

Theorem [L.-Mehta '22]

For every entire distribution, there exists an *optimal* c-approximate cost minimization prophet inequality for single items.

- c is distribution-dependent. Can be arbitrarily large.
- Use of hazard rate in prophet inequalities as *analysis tool* is new.
- For MHR distributions $\implies$ $c = 2$-approximation.

Theorem [L.-Mehta '22]

For every entire distribution, there exists an *optimal* c-approximate cost minimization prophet inequality for single items.

- c is distribution-dependent. Can be arbitrarily large.
- Use of hazard rate in prophet inequalities as *analysis tool* is new.
- For MHR distributions $\implies c = 2$-approximation.

Let $H(x) = \int_0^x h(u)\, du$ (Cumulative Hazard Rate).

Entire Distribution

$\mathcal{D}$ is *entire* if H has convergent series expansion $H(x) = \sum_{i=1}^{\infty} a_i x^{d_i}$ (where $0 < d_1 < d_2 < \dots$) for every x in the support of $\mathcal{D}$.

- E.g. uniform, exponential, Gaussian, Weibull, Rayleigh, beta, gamma

$$c(d_1) = \frac{\left(1 + {}^1\!/_{d_1}\right)^{1/d_1}}{\Gamma\left(1 + {}^1\!/_{d_1}\right)} = \Theta\left(e^{1/d_1}\right)$$

- Γ: Gamma function. $\Gamma(n+1) = n!$.
- c is tight for $\mathcal{D}$ with $H(x) = x^{d_1}$

$$c(d_1) = \frac{\left(1 + {}^1\!/_{d_1}\right)^{1/d_1}}{\Gamma\left(1 + {}^1\!/_{d_1}\right)} = \Theta\left(e^{1/d_1}\right)$$

- Γ: Gamma function. $\Gamma(n+1) = n!$.
- c is tight for $\mathcal{D}$ with $H(x) = x^{d_1}$

Why only for entire distributions?

Equal-Revenue Distribution:

$F(x) = 1 - {}^1\!/_x$. $\mathbb{E}[X] = +\infty$, but $\mathbb{E}[\min\{X_1, X_2\}] < +\infty$.

$H(x) = \log x$ and its power series converges only for $x \leq 2$.

- [Lucier '22]

Open Problems

- Extend $1/e$-selectable Greedy OCRS to general matroids.
- Tight approximations for rank-1 *prophet secretary* and *free-order prophet inequality*.
- Tight OCRSs for matchings.
- Extend cost PI to other constraints.
- Many more...

Thank You!

Questions?