Optimal Stopping Problems for Cost Minimization

Vasilis Livanos

livanos3@illinois.edu

University of Illinois Urbana-Champaign

March 10th, 2023

Overview

Prophet Inequalities

Generalizations

The Grocer's Problem

- 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.

The Grocer's Problem

- 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.

Optimal Stopping Problems

- Worst-case order + unknown v_i's = Can't do anything.
- Random order + unknown v_i's = *Secretary problem*.
- $\boxed{\text{Worst-case order + some knowledge of } v_i\text{'s}}$

Optimal Stopping Problems

- Worst-case order + unknown v_i's = Can't do anything.
- Random order + unknown v_i's = *Secretary problem*.
- $\boxed{\text{Worst-case order + some knowledge of } v_i\text{'s}}$

Prophet Inequality
[Krengel, Sucheston and Garling '77]

$X_1, X_2, \ldots, X_n \overset{\text{ind.}}{\sim}$ (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*: $\mathbb{E}[\max_i X_i]$.

Let's play!

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

Let's play!

$\mathcal{U}[2,4]$	$\mathcal{U}[2,4]$	$\mathcal{U}[1,5]$	$\mathcal{U}[0,7]$
$X_1 = 2.34$			

Let's play!

$\boxed{\mathcal{U}[2,4]}$	$\boxed{\mathcal{U}[2,4]}$	$\boxed{\mathcal{U}[1,5]}$	$\boxed{\mathcal{U}[0,7]}$
$X_1 = 2.34$	$X_2 = 3.16$		

Let's play!

$\mathcal{U}[2,4]$	$\mathcal{U}[2,4]$	$\mathcal{U}[1,5]$	$\mathcal{U}[0,7]$
$X_1 = 2.34$	$X_2 = 3.16$	$X_3 = 3.79$	

Let's play!

$\mathcal{U}[2,4]$	$\mathcal{U}[2,4]$	$\mathcal{U}[1,5]$	$\mathcal{U}[0,7]$
$X_1 = 2.34$	$X_2 = 3.16$	$X_3 = 3.79$	$X_4 = 4.39$

Prophet Inequality

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.

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

$\mathbb{E}\left[\text{ALG}\right] = 1$ 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].
  - Corresponds to item *price* in grocer's problem.

Proof for rank-1 Matroids

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

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

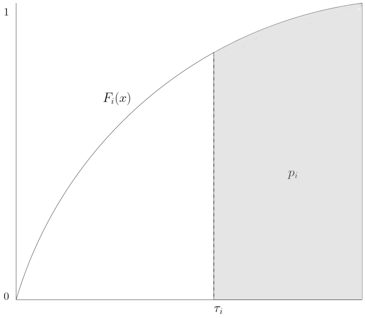

Proof for rank-1 Matroids

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).

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

Proof for rank-1 Matroids

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 $\mathbb{E}[X^*]$.

Proof for rank-1 Matroids

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 $q_i = \frac{1}{2r_i} \implies r_{i+1} = r_i - \frac{p_i}{2} = 1 - \sum_{j\leq i} \frac{p_i}{2} \geq \frac{1}{2}$

Proof for rank-1 Matroids

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 $\mathbb{E}[X^*]$.

Thus

$$\mathbb{E}[ALG] \geq \frac{1}{2} \cdot \sum_i x_i \cdot w_i.$$

Cost Minimization

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

Cost Minimization

- Objective: Minimize selected value. Prophet: $\mathbb{E}[\min_i X_i]$.
- 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. Prophet: $\mathbb{E}[\min_i X_i]$.
- 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! $\Pr[\text{We are forced to select } X_n] \to 1$.

Is Cost Minimization hopeless?

Optimal algorithm: set $\tau_i = \mathbb{E}[\mathrm{ALG}_{i+1,\ldots,n}]$, accept first $X_i \leq \tau_i$.

Is Cost Minimization hopeless?

Optimal algorithm: set $\tau_i = \mathbb{E}[\mathrm{ALG}_{i+1,\dots,n}]$, accept first $X_i \leq \tau_i$.

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.
  Good guarantees in applications (e.g. auction revenue maximization).

Theorem ([L.-Mehta '22])

For every entire distribution, there exists an optimal c-approximate Cost PI for rank-1 matroids.

- c is distribution-dependent – can be arbitrarily large.
- New use of hazard rate in PIs as *analysis tool*.
- MHR distributions $\implies$ $c = 2$-approximation.

Theorem ([L.-Mehta '22])

For every entire distribution, there exists an optimal c-approximate Cost PI for rank-1 matroids.

- c is distribution-dependent – can be arbitrarily large.
- New use of hazard rate in PIs as *analysis tool*.
- 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

PI for Cost Minimization

$$c(d_1) = \Theta\left(e^{1/d_1}\right)$$

- c is tight for $\mathcal{D}$ with $H(x) = x^{d_1}$.

Theorem ([L.-Mehta '22])

For every entire distribution, there exists a (tight) $\mathrm{O}\left(\operatorname{polylog} n\right)$ *-approximate* ***single-threshold*** *Cost PI for rank-1 matroids.*

PI for Cost Minimization

$$c(d_1) = \Theta\left(e^{1/d_1}\right)$$

- c is tight for $\mathcal{D}$ with $H(x) = x^{d_1}$.

Theorem ([L.-Mehta '22])

For every entire distribution, there exists a (tight) $\mathrm{O}(\text{polylog}\, n)$ *-approximate **single-threshold** Cost PI for rank-*1 *matroids.*

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]

Beating the Prophet?

Powerful technique. $\mathbb{E}[\mathrm{ALG}] \le \mathbb{E}[\max_{i=1}^{n} X_i]$?

Beating the Prophet?

Powerful technique. $\mathbb{E}[\mathrm{ALG}] \leq \mathbb{E}[\max_{i=1}^{n} X_i]$?

Yes; if ALG on $X_1, \ldots, X_m$ and $m > n$.

Definition (Competition Complexity)

The *competition complexity (CC)* of a distribution $\mathcal{D}$ is

$$\sup_{n \geq 1} \frac{\inf \{ m \mid \mathbb{E}[\mathrm{ALG}_m] \leq \mathbb{E}[\max_{i=1}^{n} X_i] \}}{n}$$

Extreme Value Distributions

Theorem ([Fisher–Tippett–Gnedenko '28, '43])

The minimum of n IID random variables (after renormalization) converges in distribution to one of 3 possible distributions:

- *Gumbel*
- *Weibull*
- *Fréchet*

Theorem ([L.-Saona-Verdugo unpublished])

The CC of Gumbel and Weibull distributions is 2.
The CC of Fréchet distributions is ∞.

Theorem ([Fisher–Tippett–Gnedenko '28, '43])

The minimum of n IID random variables (after renormalization) converges in distribution to one of 3 possible distributions:

- *Gumbel*
- *Weibull*
- *Fréchet*

Theorem ([L.-Saona-Verdugo unpublished])

The CC of Gumbel and Weibull distributions is 2.
The CC of Fréchet distributions is ∞.

Theorem ([L.-Saona-Verdugo unpublished])

For maximization, the CC of Gumbel and Weibull distributions is 2.
The CC of Fréchet distributions is e.

Overview

Prophet Inequalities

Generalizations

How to generalize PIs?

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

- $\boldsymbol{x}$: Optimal solution to (2). How to round $\boldsymbol{x}$?

How to generalize PIs?

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

- $\boldsymbol{x}$: Optimal solution to (2). How to 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 PIs?

Attempt #2: Contention Resolution Scheme (CRS) π

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

- [Chekuri, Vondrák and Zenklusen '11].

c-selectability

CRS is c-selectable if

$$\Pr\left[i \in \pi(R) \mid i \in R\right] \geq c \qquad \forall i.$$

- CRS is c-selectable $\implies$ c-approximation to LP.
- CRSs combine in black-box way for general constraints/objectives.

Contention Resolution Schemes (CRSs)

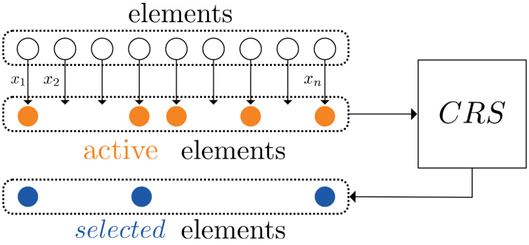

Theorem [Chekuri, Vondrák and Zenklusen '11]

There exists a $(1 - {}^1\!/_e)$-selectable CRS for matroid polytopes.

Holds if R revealed in *uniformly random* order. What about *adversarial* order?

Online Contention Resolution Scheme (OCRS)
[Alaei '11, Feldman, Svensson and Zenklusen '16]

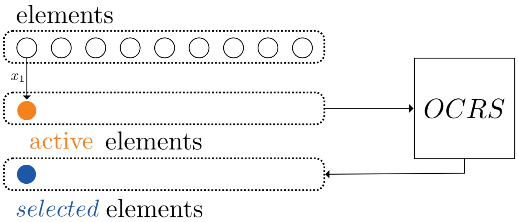

Online Contention Resolution Scheme (OCRS)
[Alaei '11, Feldman, Svensson and Zenklusen '16]

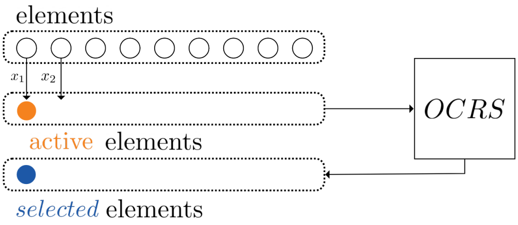

Online Contention Resolution Scheme (OCRS)
[Alaei '11, Feldman, Svensson and Zenklusen '16]

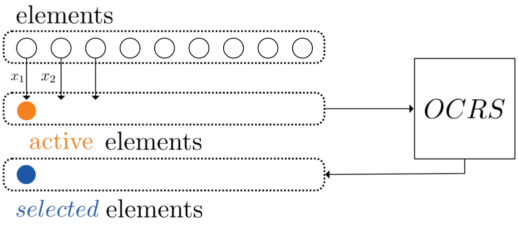

Online Contention Resolution Scheme (OCRS)
[Alaei '11, Feldman, Svensson and Zenklusen '16]

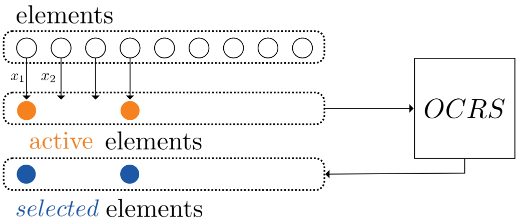

Online Contention Resolution Scheme (OCRS)

[Alaei '11, Feldman, Svensson and Zenklusen '16]

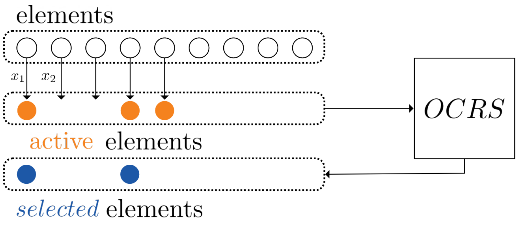

Online Contention Resolution Scheme (OCRS)
[Alaei '11, Feldman, Svensson and Zenklusen '16]

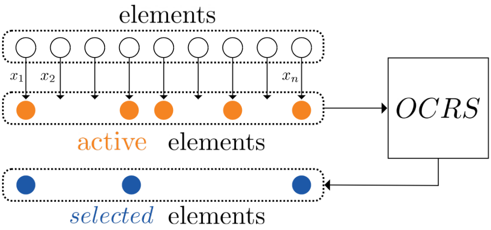

$\exists$ 1/2-selectable OCRS for rank-1 matroids (tight). [Alaei '11]
$\exists$ 1/2-selectable OCRS for matroids. [Lee, Singla '18]

Greedy OCRSs

Greedy OCRS (Informal)

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

Theorem ([L. '22])

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

- Works against more powerful adversaries.
- Idea extends to partition, uniform and transversal matroids.

Open Problems

1. For Cost PI, only use k-thresholds for $1 < k < n$. How does the ratio change?
2. Can we extend $1/e$-Greedy OCRS to general matroids?
3. What is the best c-OCRS for (bipartite / general) matchings?
   - ≥ 0.349-OCRS for bipartite, ≥ 0.344-OCRS [MacRury, Ma, Grammel '22]
   - ≤ 0.433-OCRS for bipartite, ≤ 0.4-OCRS [MacRury, Ma, Grammel '22]
   - $\geq 1/2e \approx 0.184$-Greedy OCRS [Feldman, Svensson, Zenklusen '16]

Questions?