Constrained-Order Prophet Inequalities

to appear in SODA '21

Makis Arsenis, Odysseas Drosis (EPFL), Robert Kleinberg

Cornell University
https://arxiv.org/abs/2010.09705

Prophet Inequalities

Prophet Inequalities — Example

$U[\$10, \$20]$

$U[\$1, \$50]$

$$\begin{cases} \$1,000, & \text{w.p. } 0.01 \\ \$0, & \text{w.p. } 0.99 \end{cases}$$

$U[\$0, \$10]$

[1] Clip-art source: `https://gallery.yopriceville.com/Free-Clipart-Pictures/`

Prophet Inequalities — Example

$\$17$

$U[\$1, \$50]$

$$\begin{cases} \$1,000, & \text{w.p. } 0.01 \\ \$0, & \text{w.p. } 0.99 \end{cases}$$

$\mathrm{U}[\$0, \$10]$

Prophet Inequalities — Example

$\$17$

$\$3$

$$\begin{cases} \$1,000, & \text{w.p. } 0.01 \\ \$0, & \text{w.p. } 0.99 \end{cases}$$

$U[\$0, \$10]$

Prophet Inequalities — Example

$17 $3 $0 U[$0, $10]

Prophet Inequalities — Example

$17

$3

$0

$6

Prophet Inequalities — Standard setting

- Sequence of n **independent**, **non-negative** random variables:

$$X_1, X_2, \ldots, X_n$$

Prophet Inequalities — Standard setting

- Sequence of n **independent**, **non-negative** random variables:

$$X_1, X_2, \ldots, X_n$$

- **Gambler**: Knows **n** and the **distribution** $\mathcal{D}_i$ of each X_i.
  Can inspect the r.v.'s in the **given order**.
  Accepts at most one of the values, X_τ, as **reward**.
  Acceptance decisions are **irrevocable**.

Prophet Inequalities — Standard setting

- Sequence of n **independent**, **non-negative** random variables:

$$X_1, X_2, \ldots, X_n$$

- **Gambler**: Knows **n** and the **distribution** $\mathcal{D}_i$ of each X_i.
  Can inspect the r.v.'s in the **given order**.
  Accepts at most one of the values, X_τ, as **reward**.
  Acceptance decisions are **irrevocable**.
- **Prophet**: Knows the values of **all** variables.
  Always chooses the maximum of them, i.e. $\mathbf{X_* = \max_i X_i}$.

Prophet Inequalities — Standard setting

- Sequence of n **independent**, **non-negative** random variables:

$$X_1, X_2, \ldots, X_n$$

- **Gambler**: Knows **n** and the **distribution** $\mathcal{D}_i$ of each X_i.
  Can inspect the r.v.'s in the **given order**.
  Accepts at most one of the values, X_τ, as **reward**.
  Acceptance decisions are **irrevocable**.
- **Prophet**: Knows the values of **all** variables.
  Always chooses the maximum of them, i.e. $\mathbf{X}_* = \max_\mathbf{i} \mathbf{X}_\mathbf{i}$.
- **Stopping rule** τ: Algorithm that gambler follows for a given $\mathcal{D}_1, \ldots, \mathcal{D}_n$.

Prophet Inequalities — Standard setting

- Sequence of n **independent**, **non-negative** random variables:

$$X_1, X_2, \ldots, X_n$$

- **Gambler**: Knows **n** and the **distribution** $\mathcal{D}_i$ of each X_i.
  Can inspect the r.v.'s in the **given order**.
  Accepts at most one of the values, X_τ, as **reward**.
  Acceptance decisions are **irrevocable**.
- **Prophet**: Knows the values of **all** variables.
  Always chooses the maximum of them, i.e. $\mathbf{X}_* = \max_\mathbf{i} \mathbf{X}_\mathbf{i}$.
- **Stopping rule** τ: Algorithm that gambler follows for a given $\mathcal{D}_1, \ldots, \mathcal{D}_n$.
- **Prophet Inequality** is a statement of the following form:
  For all $\mathcal{D}_1, \ldots, \mathcal{D}_n$, there exists a stopping rule τ s.t.:

$$\mathbb{E}[X_\tau] \geq r \cdot \mathbb{E}[X_*]$$

Prophet Inequalities — Standard setting

- Sequence of n **independent**, **non-negative** random variables:

$$X_1, X_2, \ldots, X_n$$

- **Gambler**: Knows **n** and the **distribution** $\mathcal{D}_i$ of each X_i.
  Can inspect the r.v.'s in the **given order**.
  Accepts at most one of the values, X_τ, as **reward**.
  Acceptance decisions are **irrevocable**.
- **Prophet**: Knows the values of **all** variables.
  Always chooses the maximum of them, i.e. $\mathbf{X}_* = \max_{\mathbf{i}} \mathbf{X}_{\mathbf{i}}$.
- **Stopping rule** τ: Algorithm that gambler follows for a given $\mathcal{D}_1, \ldots, \mathcal{D}_n$.
- **Prophet Inequality** is a statement of the following form:
  For all $\mathcal{D}_1, \ldots, \mathcal{D}_n$, there exists a stopping rule τ s.t.:

$$\mathbb{E}[X_\tau] \geq r \cdot \mathbb{E}[X_*]$$

- **Gambler-to-prophet/Competitive ratio**:

$$r = \inf_{\mathcal{D}_1, \ldots, \mathcal{D}_n} \sup_{\text{Stopping rule } \tau} \frac{\mathbb{E}[X_\tau]}{\mathbb{E}[X_*]}$$

A Note on Stopping Rules

- **Threshold stopping rule:** Gambler decides on a threshold T.
  - If Gambler reaches $X_i > T$, then Gambler accepts.
  - If Gambler reaches $X_i < T$, then Gambler rejects and proceeds.

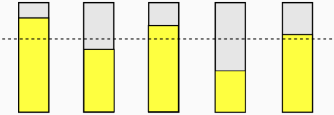

Why are we interested in Prophet Inequalities?

- Mathematicians started looking into this problem in the 70s and 80s.

Why are we interested in Prophet Inequalities?

- Mathematicians started looking into this problem in the 70s and 80s.
- Models **real-life situations** where decisions have to be made under uncertainty
  e.g. Hiring and job interviews, investing in the stock market or even choosing a life partner!

Why are we interested in Prophet Inequalities?

- Mathematicians started looking into this problem in the 70s and 80s.
- Models **real-life situations** where decisions have to be made under uncertainty
  e.g. Hiring and job interviews, investing in the stock market or even choosing a life partner!
- Regained attention by Computer Scientists in the late 2000s because of it's applications in **Mechanism Design**:

Why are we interested in Prophet Inequalities?

- Mathematicians started looking into this problem in the 70s and 80s.
- Models **real-life situations** where decisions have to be made under uncertainty
  e.g. Hiring and job interviews, investing in the stock market or even choosing a life partner!
- Regained attention by Computer Scientists in the late 2000s because of it's applications in **Mechanism Design**:
  - Selling an item to n potential buyers.

Why are we interested in Prophet Inequalities?

- Mathematicians started looking into this problem in the 70s and 80s.
- Models **real-life situations** where decisions have to be made under uncertainty
  e.g. Hiring and job interviews, investing in the stock market or even choosing a life partner!
- Regained attention by Computer Scientists in the late 2000s because of it's applications in **Mechanism Design**:
  - Selling an item to n potential buyers.
  - Buyers arrive sequentially and are offered a take-it or leave-it price for the item.

Why are we interested in Prophet Inequalities?

- Mathematicians started looking into this problem in the 70s and 80s.
- Models **real-life situations** where decisions have to be made under uncertainty
  e.g. Hiring and job interviews, investing in the stock market or even choosing a life partner!
- Regained attention by Computer Scientists in the late 2000s because of it's applications in **Mechanism Design**:
  - Selling an item to n potential buyers.
  - Buyers arrive sequentially and are offered a take-it or leave-it price for the item.
  - **Simpler** compared to e.g. an auction

Why are we interested in Prophet Inequalities?

- Mathematicians started looking into this problem in the 70s and 80s.
- Models **real-life situations** where decisions have to be made under uncertainty
  e.g. Hiring and job interviews, investing in the stock market or even choosing a life partner!
- Regained attention by Computer Scientists in the late 2000s because of it's applications in **Mechanism Design**:
  - Selling an item to n potential buyers.
  - Buyers arrive sequentially and are offered a take-it or leave-it price for the item.
  - **Simpler** compared to e.g. an auction
  - Take-it or leave-it price corresponds to the **threshold** of a stopping rule.

Why are we interested in Prophet Inequalities?

- Mathematicians started looking into this problem in the 70s and 80s.
- Models **real-life situations** where decisions have to be made under uncertainty
  e.g. Hiring and job interviews, investing in the stock market or even choosing a life partner!
- Regained attention by Computer Scientists in the late 2000s because of it's applications in **Mechanism Design**:
  - Selling an item to n potential buyers.
  - Buyers arrive sequentially and are offered a take-it or leave-it price for the item.
  - **Simpler** compared to e.g. an auction
  - Take-it or leave-it price corresponds to the **threshold** of a stopping rule.
  - Prophet inequalities provide **welfare/revenue guarantees** for Sequential Posted-Price Mechanisms.

Standard Prophet Inequality

Theorem ([Samuel-Cahn, 1984, Krengel and Sucheston, 1977])
There exists a (threshold) **stopping rule** *that satisfies a prophet inequality with factor* $\mathbf{1/2}$.

Standard Prophet Inequality

Theorem ([Samuel-Cahn, 1984, Krengel and Sucheston, 1977])
There exists a (threshold) **stopping rule** *that satisfies a prophet inequality with factor* $\mathbf{1/2}$.

Proof.

Let T be the threshold and let $\Pr[\max X_i \geq T] = p \in [0, 1]$.

Standard Prophet Inequality

Theorem ([Samuel-Cahn, 1984, Krengel and Sucheston, 1977])
There exists a (threshold) **stopping rule** *that satisfies a prophet inequality with factor* $\mathbf{1/2}$.

Proof.

Let T be the threshold and let $\Pr[\max X_i \geq T] = p \in [0, 1]$.

$$\begin{aligned}
\mathbb{E}[X_*] &= \mathbb{E}\left[\max_{i=1}^{n} X_i\right] \\
&\leq \mathbb{E}\left[T + \max_{i=1}^{n}(X_i - T)^+\right] \\
&\leq T + \sum_{i=1}^{n} \mathbb{E}\left[(X_i - T)^+\right]
\end{aligned}$$

Theorem ([Samuel-Cahn, 1984, Krengel and Sucheston, 1977])
There exists a (threshold) **stopping rule** *that satisfies a prophet inequality with factor* $\mathbf{1/2}$.

Proof.

Let T be the threshold and let $\Pr[\max X_i \geq T] = p \in [0,1]$.

$$\begin{aligned}\mathbb{E}[X_*] &= \mathbb{E}\left[\max_{i=1}^{n} X_i\right] \\ &\leq \mathbb{E}\left[T + \max_{i=1}^{n}(X_i - T)^+\right] \\ &\leq T + \sum_{i=1}^{n} \mathbb{E}\left[(X_i - T)^+\right]\end{aligned}$$

$$\begin{aligned}\mathbb{E}[X_\tau] &= \mathbb{E}\left[\sum_{i=1}^{n} X_i \cdot \mathbb{I}[\tau = i]\right] \\ &= \mathbb{E}\left[\sum_{i=1}^{n} T \cdot \mathbb{I}[\tau = i] + \sum_{i=1}^{n}(X_i - T) \cdot \mathbb{I}[\tau = i]\right] \\ &= pT + \sum_{i=1}^{n} c_i \cdot \mathbb{E}\left[(X_i - T)^+\right]\end{aligned}$$

where $c_i = \Pr[\text{No item is accepted before reaching } X_i]$.

Standard Prophet Inequality

$$\mathbb{E}[X_*] \leq T + \sum_{i=1}^{n} \mathbb{E}[(X_i - T)^+] \qquad \mathbb{E}[X_\tau] = pT + \sum_{i=1}^{n} c_i \cdot \mathbb{E}[(X_i - T)^+]$$

$c_i = \Pr[\text{No item is accepted before reaching } X_i]$

Standard Prophet Inequality

$$\mathbb{E}[X_*] \leq T + \sum_{i=1}^{n} \mathbb{E}[(X_i - T)^+] \qquad \mathbb{E}[X_\tau] = pT + \sum_{i=1}^{n} c_i \cdot \mathbb{E}[(X_i - T)^+]$$

$$c_i = \Pr[\text{No item is accepted before reaching } X_i]$$

Bound c_i:

$$c_i = \prod_{j<i} \Pr[X_j < T] \geq \prod_{j=1}^{n} \Pr[X_j < T] = 1 - p$$

Standard Prophet Inequality

$$\mathbb{E}[X_*] \leq T + \sum_{i=1}^{n} \mathbb{E}[(X_i - T)^+] \qquad \mathbb{E}[X_\tau] = pT + \sum_{i=1}^{n} c_i \cdot \mathbb{E}[(X_i - T)^+]$$

$$c_i = \Pr[\text{No item is accepted before reaching } X_i] \geq 1 - p$$

Standard Prophet Inequality

$$\mathbb{E}[X_*] \leq T + \sum_{i=1}^{n} \mathbb{E}[(X_i - T)^+] \qquad \mathbb{E}[X_\tau] = pT + \sum_{i=1}^{n} c_i \cdot \mathbb{E}[(X_i - T)^+]$$

$$c_i = \Pr[\text{No item is accepted before reaching } X_i] \geq 1 - p$$

Substitute back,

$$\mathbb{E}[\text{Gambler}] \geq pT + (1-p)\sum_{i=1}^{n} \mathbb{E}\left[(X_i - T)^+\right]$$

Standard Prophet Inequality

$$\mathbb{E}[X_*] \le T + \sum_{i=1}^{n} \mathbb{E}[(X_i - T)^+] \qquad \mathbb{E}[X_\tau] = pT + \sum_{i=1}^{n} c_i \cdot \mathbb{E}[(X_i - T)^+]$$

$$c_i = \Pr[\text{No item is accepted before reaching } X_i] \ge 1 - p$$

Substitute back,

$$\mathbb{E}[\text{Gambler}] \ge pT + (1-p) \sum_{i=1}^{n} \mathbb{E}\left[(X_i - T)^+\right]$$

Choose T s.t. $p = 1 - p$.

$$\mathbb{E}[\text{Gambler}] \overset{p=1/2}{\ge} \frac{1}{2}\left(T + \sum_{i=1}^{n} \mathbb{E}\left[(X_i - T)^+\right]\right)$$
$$\ge \frac{1}{2} \cdot \mathbb{E}[\text{Prophet}]$$

□

Standard Prophet Inequality

- Previous result is **tight** even for **general stopping rules**:

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

$$\mathbb{E}[\text{Prophet}] = \varepsilon \cdot \frac{1}{\varepsilon} + (1-\varepsilon) \cdot 1 = 2 - \varepsilon$$

$$\mathbb{E}[\text{Gambler}] = 1$$

Standard Prophet Inequality

- Previous result is **tight** even for **general stopping rules**:

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

$$\mathbb{E}[\text{Prophet}] = \varepsilon \cdot \frac{1}{\varepsilon} + (1 - \varepsilon) \cdot 1 = 2 - \varepsilon$$
$$\mathbb{E}[\text{Gambler}] = 1$$

- **Takeaway**: The reason Gambler does bad is high uncertainty far in the future.

Constrained-Order Prophet Inequalities

Constrained-Order Prophet Inequalities

We augment the prophet inequalities model to allow for order-selection:

- Π: set of permutations on $[n]$.
- Gambler can choose any $\pi \in \Pi$ and inspect the variables in that order:

$$X_{\pi(1)}, X_{\pi(2)}, \ldots, X_{\pi(n)}$$

- **Adversarial Order**: $\Pi = \{\text{id}\}$,
  i.e. gambler must inspect the variables in the order given by the adversary.
- **Free Order**: $\Pi = S_n$, the set of all permutations on n elements,
  i.e. gambler is free to choose any ordering.
- **Random Order (Prophet secretary problem)**:
  $\Pi = S_n$ but π is chosen **uniformly at random**.
- **Forward-Reverse order**: $\Pi = \{\text{id}, \text{rev}\}$.
- **General Constrained-Order**: Arbitrary Π.

Motivation

- **Adversarial Order**: Models the **uncertainty** in decision making.
- **Free Order**: Models the power that **choice** gives us in decision making under uncertainty.
- **Constrained Order**: Offers a way to understand better where the power of choice comes from.

What is known?

	Threshold Rules	General Rules
Adversarial	$1/2$ [Samuel-Cahn, 1984]	$1/2$ [Krengel and Sucheston, 1977]
Free Order	$1-\frac{1}{e}=0.632\ldots$ [Yan, 2011, Correa et al., 2017]	LB: $0.669\ldots$ [Correa et al., 2019] UB: $0.745\ldots$ [Hill and Kertz, 1982]
Random Order	$1-\frac{1}{e}$	LB: $0.669\ldots$ [Correa et al., 2019] UB: $\sqrt{3}-1=0.732\ldots$ [Correa et al., 2019]

Our Contribution

We are exploring the landscape between the two extremes: the **Adversarial** and **Free** order setting.

Our Contribution

We are exploring the landscape between the two extremes: the **Adversarial** and **Free** order setting.

Definition
For $\Pi \subseteq S_n$, define the **threshold prophet ratio** on Π as follows:

$$\mathrm{TPR}(\Pi) = \inf_{\mathcal{D}_1,\ldots,D_n} \sup_{\text{threshold stopping rule on } \Pi} \frac{\mathbb{E}[\text{Gambler}]}{\mathbb{E}[\text{Prophet}]}$$

Our Contribution

We are exploring the landscape between the two extremes: the **Adversarial** and **Free** order setting.

Definition
For $\Pi \subseteq S_n$, define the **threshold prophet ratio** on Π as follows:

$$\mathrm{TPR}(\Pi) = \inf_{\mathcal{D}_1,\dots,D_n} \sup_{\text{threshold stopping rule on } \Pi} \frac{\mathbb{E}[\text{Gambler}]}{\mathbb{E}[\text{Prophet}]}$$

Q: For a given α, what is the minimum size m of Π such that $\mathrm{TPR}(\Pi) \geq \alpha$?

Our Contribution

We are exploring the landscape between the two extremes: the **Adversarial** and **Free** order setting.

Definition
For $\Pi \subseteq S_n$, define the **threshold prophet ratio** on Π as follows:

$$\mathrm{TPR}(\Pi) = \inf_{\mathcal{D}_1,\ldots,D_n} \sup_{\text{threshold stopping rule on } \Pi} \frac{\mathbb{E}[\text{Gambler}]}{\mathbb{E}[\text{Prophet}]}$$

Q: For a given α, what is the minimum size m of Π such that $\mathrm{TPR}(\Pi) \geq \alpha$?

$\alpha \in \left[0, \frac{1}{2}\right]$	$m = 1$
$\alpha \in \left(\frac{1}{2}, \varphi^{-1}\right)$	$m = 2$
$\alpha \in \left(\varphi^{-1}, 1 - \frac{1}{e}\right)$	$m = \Theta(\log n)$
$\alpha = 1 - \frac{1}{e}$	$m = O(n^2)$

Forward-Reverse Prophet Inequality

Forward-Reverse Order Prophet Inequality

Theorem ([A-Drosis-Kleinberg, SODA '21])
In the **forward-reverse** *prophet inequality setting, there exists a* **threshold stopping rule** *with a* **gambler-to-prophet ratio** *of at least* $\varphi^{-1} = \frac{\sqrt{5}-1}{2} = \mathbf{0.618}\ldots$

Forward-Reverse Order Prophet Inequality

Theorem ([A-Drosis-Kleinberg, SODA '21])
In the **forward-reverse** *prophet inequality setting, there exists a* **threshold stopping rule** *with a* **gambler-to-prophet ratio** *of at least* $\varphi^{-1} = \frac{\sqrt{5}-1}{2} = \mathbf{0.618}\ldots$

Proof.

Pick $\pi \in \{\text{id}, \text{rev}\}$ uniformly at random.
Similarly to previous proof, set threshold T s.t. $\Pr[\max X_i \geq T] = p \in [0,1]$.
Again,

$$\mathbb{E}[\text{Prophet}] \leq T + \sum_{i=1}^{n} \mathbb{E}\left[(X_i - T)^+\right]$$

$$\mathbb{E}[\text{Gambler}] = pT + \sum_{i=1}^{n} c_i \cdot \mathbb{E}\left[(X_i - T)^+\right]$$

where $c_i = \Pr[\text{No element is selected before reaching } X_i]$.

Forward-Reverse Order Prophet Inequality

$$
\begin{aligned}
c_i &= \frac{1}{2}\left(\prod_{j<i}\Pr[X_j < T] + \prod_{j>i}\Pr[X_j > T]\right) \\
&\overset{\text{AM-GM}}{\geq} \left(\prod_{j\neq i}^{n}\Pr[X_j < T]\right)^{1/2} \\
&\geq \sqrt{1-p}
\end{aligned}
$$

Substitute back,

$$
\begin{aligned}
\mathbb{E}[\text{Gambler}] &\geq pT + \sqrt{1-p}\cdot\sum_{i=1}^{n}\mathbb{E}\left[(X_i - T)^+\right] \\
&\overset{p=\varphi^{-1}}{=} \varphi^{-1}\left(T + \sum_{i=1}^{n}\mathbb{E}\left[(X_i - T)^+\right]\right) \\
&\geq \varphi^{-1}\cdot\mathbb{E}[\text{Prophet}]
\end{aligned}
$$

□

Forward/Reverse Order Upper Bound

Lemma
When $n \geq 3$ and $\Pi = \{id, rev\}$, no threshold stopping rule can have a gambler-to-prophet ratio greater than φ^{-1}.

Proof sketch.

- For $n = 3$:

$$X_1 = \mathrm{U}[1-\varepsilon, 1], \quad X_2 = \begin{cases} \frac{2\varphi^{-1}}{\varepsilon}, & \text{w.p. } \varepsilon \\ 0, & \text{w.p. } 1-\varepsilon \end{cases}, \quad X_3 = \mathrm{U}[1-\varepsilon, 1]$$

- For $n > 3$:
  Let $i < j < k$ be arbitrary r.v. indices. Define X_i, X_j, X_k just like X_1, X_2, X_3 above and let $X_l = 0$ for all $l \notin \{i, j, k\}$.

□

Beating the Golden Ratio

Beating the golden ratio

- Two permutations suffice to go from 0.5 to $\varphi^{-1} = 0.618\ldots$.

Beating the golden ratio

- Two permutations suffice to go from 0.5 to $\varphi^{-1} = 0.618\ldots$.
- **Q:** How many permutations are needed to guarantee a ratio $> \varphi^{-1}$?

Beating the golden ratio

- Two permutations suffice to go from 0.5 to $\varphi^{-1} = 0.618\ldots$.
- **Q:** How many permutations are needed to guarantee a ratio $> \varphi^{-1}$?
- **Idea**: Require the existence of a "central element".

Beating the golden ratio

Definition
We say $j \in [n]$ is ε-**centered w.r.t.** Π (a set of permutations of $[n]$) if there exists a **probability distribution** p on $[n] \setminus \{j\}$ such that:

$$\forall \pi \in \Pi : \Pr_{i \sim p}[\pi^{-1}(i) < \pi^{-1}(j)] \geq 1/2 - \varepsilon$$
$$\Pr_{i \sim p}[\pi^{-1}(i) > \pi^{-1}(j)] \geq 1/2 - \varepsilon$$

Beating the golden ratio

Definition
We say $j \in [n]$ is ε**-centered w.r.t.** Π (a set of permutations of $[n]$) if there exists a **probability distribution** p on $[n] \setminus \{j\}$ such that:

$$\forall \pi \in \Pi : \Pr_{i \sim p}[\pi^{-1}(i) < \pi^{-1}(j)] \geq 1/2 - \varepsilon$$
$$\Pr_{i \sim p}[\pi^{-1}(i) > \pi^{-1}(j)] \geq 1/2 - \varepsilon$$

Lemma
If Π *is a set of permutations of* $[n]$ *and* j *is an* ε*-centered element w.r.t.* Π*, then* $\mathrm{TPR}(\Pi) \leq \varphi^{-1} + O(\varepsilon)$.

Beating the golden ratio

Definition
We say $j \in [n]$ is ε-**centered w.r.t.** Π (a set of permutations of $[n]$) if there exists a **probability distribution** p on $[n] \setminus \{j\}$ such that:

$$\begin{aligned} \forall \pi \in \Pi : \Pr_{i \sim p}[\pi^{-1}(i) < \pi^{-1}(j)] &\geq 1/2 - \varepsilon \\ \Pr_{i \sim p}[\pi^{-1}(i) > \pi^{-1}(j)] &\geq 1/2 - \varepsilon \end{aligned}$$

Lemma
If Π is a set of permutations of $[n]$ and j is an ε-centered element w.r.t. Π, then $\mathrm{TPR}(\Pi) \leq \varphi^{-1} + O(\varepsilon)$.

Lemma (Exact)
If $|\Pi| < \sqrt{\log n}$, then $\exists j \in [n]$ that is (0)-centered w.r.t. Π.

Lemma (Approximate)
If $|\Pi| < \log_{1/\varepsilon} n$ for $\varepsilon > 0$, then $\exists j \in [n]$ that is ε-centered w.r.t. Π.

Achieving the Optimal Threshold Ratio

Achieving the Optimal Threshold Ratio

Convention:

- Variable indices: $i \in [n]$
- Arrival position: $k \in [n]$

$$\pi : [n] \to [n], \quad \sigma = \pi^{-1} : [n] \to [n]$$

Definition
A **distribution** $\mathcal{P}$ over permutations $\Pi \subseteq S_n$ is **pairwise independent** if: $\forall i \neq j \in [n]$, $(\sigma(i), \sigma(j))$ is distributed **uniformly** over $\{(a, b) \in [n] \times [n] \mid a \neq b\}$ when $\pi \sim \mathcal{P}$.

Achieving the Optimal Threshold Ratio

Convention:

- Variable indices: $i \in [n]$
- Arrival position: $k \in [n]$

$$\pi : [n] \to [n], \quad \sigma = \pi^{-1} : [n] \to [n]$$

Definition
A **distribution** $\mathcal{P}$ over permutations $\Pi \subseteq S_n$ is **pairwise independent** if: $\forall i \neq j \in [n]$, $(\sigma(i), \sigma(j))$ is distributed **uniformly** over $\{(a, b) \in [n] \times [n] \mid a \neq b\}$ when $\pi \sim \mathcal{P}$.

Remark: Pairwise independent permutations behave like uniformly random permutations,

$$\Pr_{\pi \sim \mathcal{P}}[\sigma(i) = k] = \frac{1}{n}, \quad \forall i, k \in [n]$$

$$\Pr_{\pi \sim \mathcal{P}}[\sigma(j) < k | \sigma(i) = k] = \frac{k-1}{n-1}, \quad \forall i \neq j, k \in [n]$$

Achieving the Optimal Threshold Ratio i

Lemma
For prime n, there exists a set Π *of* $\mathbf{n(n-1)}$ **permutations** *such that the uniform distribution over* Π *is pairwise independent.*

Proof sketch: $\pi_{a,b}(k) = ak + b \pmod{n}$, $a \sim \mathrm{U}[n-1], b \sim \mathrm{U}[n]$.

Achieving the Optimal Threshold Ratio i

Theorem ([A.-Drosis-Kleinberg, SODA '21])
Let π be a **random permutation** *of $[n]$ sampled from a* **pairwise-independent distribution** *of permutations. Then, there exists a threshold T such that:*

$$\mathbb{E}[\text{Gambler}] \geq \left(1 - \frac{1}{e}\right) \cdot \mathbb{E}[\text{Prophet}]$$

Proof. (resembles [Correa et al., 2019])

Again,

$$\mathbb{E}[\text{Gambler}] = pT + \sum_{i=1}^{n} c_i \cdot \mathbb{E}\left[(X_i - T)^+\right]$$

but now,

$$c_i = \sum_{k=1}^{n} \Pr[\pi(k) = i] \prod_{l=1}^{k-1} \Pr[X_{\pi(l)} < T]$$

Achieving the Optimal Threshold Ratio ii

$$
\begin{aligned}
c_i &= \sum_{k=1}^{n} \Pr[\pi(k) = i] \prod_{l=1}^{k-1} \Pr[X_{\pi(l)} < T] \\
&= \sum_{k=1}^{n} \Pr[\pi(k) = i] \sum_{S \subset [n]} \Pr[\sigma(S) = [k-1]] \mid \pi(k) = i] \prod_{j \in S} \Pr[X_j < T] \\
&= \sum_{k=1}^{n} \Pr[\pi(k) = i] \sum_{S \subset [n]} p_{k,i}(S) \prod_{j \in S} q_j \\
&\overset{\text{AM-GM}}{\geq} \sum_{k=1}^{n} \Pr[\pi(k) = i] \prod_{S \subset [n]} \left(\prod_{j \in S} q_j \right)^{p_{k,i}(S)} \\
&= \sum_{k=1}^{n} \Pr[\pi(k) = i] \prod_{j \in [n] \setminus \{i\}} q_j^{\sum_{S \subset [n]: j \in S} p_{k,i}(S)} \\
&= \sum_{k=1}^{n} \Pr[\pi(k) = i] \prod_{j \in [n] \setminus \{i\}} q_j^{\Pr[\pi(k) < j \mid \pi(k) = i]}
\end{aligned}
$$

Achieving the Optimal Threshold Ratio

$$
\begin{aligned}
c_i &\geq \sum_{k=1}^{n} \Pr[\pi(k) = i] \prod_{j \in [n]\setminus\{i\}} q_j^{\Pr[\pi(k)<j \mid \pi(k)=i]} \\
&\geq \frac{1}{n} \sum_{k=1}^{n} \left(\prod_{j \in [n]\setminus\{i\}} q_j \right)^{\frac{k-1}{n-1}} \\
&\geq \frac{1}{n} \sum_{k=1}^{n} (1-p)^{\frac{k-1}{n-1}} = \frac{1}{n} \frac{1-(1-p)^{\frac{n}{n-1}}}{1-(1-p)^{\frac{1}{n-1}}} \overset{n\to+\infty}{\simeq} \frac{p}{-\ln(1-p)}
\end{aligned}
$$

Hence,

$$
\begin{aligned}
\mathbb{E}[\text{Gambler}] &\geq pT + \frac{p}{-\ln(1-p)} \sum_{i=1}^{n} \mathbb{E}[(X_i - T)^+] \\
&\overset{p=1-\frac{1}{e}}{=} \left(1 - \frac{1}{e}\right) \mathbb{E}[\text{Prophet}]
\end{aligned}
$$

□

Achieving the Optimal Threshold Ratio

Theorem ([A.-Drosis-Kleinberg, SODA '21])
Let σ be a **random permutation** *of $[n]$ sampled from an* **$(\varepsilon, \varepsilon^2)$-almost pairwise independent distribution** *of permutations. Then, there exists a threshold T such that:*

$$\mathbb{E}[\text{Gambler}] \geq \left(1 - \frac{1}{e} - O(\varepsilon)\right) \mathbb{E}[\text{Prophet}]$$

Definition
A distribution Π on permutations of $[n]$ is (ε, δ)-**almost pairwise independent** if for every $i \neq j$, the distribution of $\left(\left\lceil \frac{\sigma(i)}{\varepsilon n} \right\rceil, \left\lceil \frac{\sigma(j)}{\varepsilon n} \right\rceil\right)$ is δ-close (in TV-distance), to the uniform distribution on $[\frac{1}{\varepsilon}] \times [\frac{1}{\varepsilon}]$.

Lemma
For any $\varepsilon, \delta > 0$ (with $1/\varepsilon \in \mathbb{Z}$, $1/\varepsilon | n$ and $\varepsilon n \geq 2/\delta$), then there exists a set Π of $O((\frac{1}{\delta\varepsilon})^2 \log n)$ permutations such that the uniform distribution over Π is (ε, δ)-almost pairwise independent.

Conclusion

Q: For a given α, what is the minimum size m of Π such that $\text{TPR}(\Pi) \geq \alpha$?

$\alpha \in \left[0, \frac{1}{2}\right]$	$m = 1$
$\alpha \in \left(\frac{1}{2}, \varphi^{-1}\right)$	$m = 2$
$\alpha \in \left(\varphi^{-1}, 1 - \frac{1}{e}\right)$	$m = \Theta(\log n)$
$\alpha = 1 - \frac{1}{e}$	$m = O(n^2)$

Open Questions in the area

- Bridge the gaps in our theorems:

Open Questions in the area

- Bridge the gaps in our theorems:
  - $\alpha = 1 - \frac{1}{e} - \varepsilon$ vs. $\alpha = 1 - \frac{1}{e}$ ($\Theta(\log n)$ vs $O(n^2)$ permutations).

Open Questions in the area

- Bridge the gaps in our theorems:
  - $\alpha = 1 - \frac{1}{e} - \varepsilon$ vs. $\alpha = 1 - \frac{1}{e}$ ($\Theta(\log n)$ vs $O(n^2)$ permutations).
  - What's the exact barrier for beating the golden ratio?

Open Questions in the area

- Bridge the gaps in our theorems:
  - $\alpha = 1 - \frac{1}{e} - \varepsilon$ vs. $\alpha = 1 - \frac{1}{e}$ ($\Theta(\log n)$ vs $O(n^2)$ permutations).
  - What's the exact barrier for beating the golden ratio?
- What about non-threshold stopping rules?
  - The power to update the threshold can bypass some of the barriers we discussed here.

Open Questions in the area

- Bridge the gaps in our theorems:
  - $\alpha = 1 - \frac{1}{e} - \varepsilon$ vs. $\alpha = 1 - \frac{1}{e}$ ($\Theta(\log n)$ vs $O(n^2)$ permutations).
  - What's the exact barrier for beating the golden ratio?
- What about non-threshold stopping rules?
  - The power to update the threshold can bypass some of the barriers we discussed here.
  - Optimal stopping rules are difficult to analyze even for small n.

Open Questions in the area

- Bridge the gaps in our theorems:
  - $\alpha = 1 - \frac{1}{e} - \varepsilon$ vs. $\alpha = 1 - \frac{1}{e}$ ($\Theta(\log n)$ vs $O(n^2)$ permutations).
  - What's the exact barrier for beating the golden ratio?
- What about non-threshold stopping rules?
  - The power to update the threshold can bypass some of the barriers we discussed here.
  - Optimal stopping rules are difficult to analyze even for small n.
- What is the best gambler-to-prophet ratio for the free order setting? What about the random order?

Thank You!

Questions?

References

Correa, J., Foncea, P., Hoeksma, R., Oosterwijk, T., and Vredeveld, T. (2017).
Posted price mechanisms for a random stream of customers.
In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, page 169–186, New York, NY, USA. Association for Computing Machinery.

Correa, J. R., Saona, R., and Ziliotto, B. (2019).
Prophet secretary through blind strategies.
In Chan, T. M., editor, *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 1946–1961. SIAM.

Hill, T. P. and Kertz, R. P. (1982).
Comparisons of stop rule and supremum expectations of i.i.d. random variables.
Ann. Probab., 10(2):336–345.

Krengel, U. and Sucheston, L. (1977).
Semiamarts and finite values.
Bull. Amer. Math. Soc., 83:745–747.

Samuel-Cahn, E. (1984).
Comparison of threshold stop rules and maximum for independent nonnegative random variables.
Annals of Probability, 12(4):1213–1216.

Yan, Q. (2011).
Mechanism design via correlation gap.
In Randall, D., editor, *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 710–719. SIAM.