IID Prophet Inequality with Random Horizon

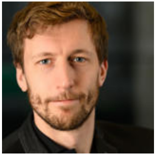  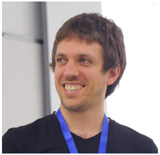

F. Mallmann-Trenn[2] G. Giambartolomei[2] **R. Saona**[1]

[1]Institute of Science and Technology Austria (ISTA)
[2]King's College London

From matchings to markets 2025

Institute of Science and Technology Austria

Mentors

Jose Correa

Krishnendu Chatterjee

Bruno Ziliotto

Miquel Oliu-Barton

IID Prophet Inequality

   

IID Prophet Inequality

$\downarrow$

10

IID Prophet Inequality

$\downarrow$ $\cancel{1}\cancel{0}$

$\downarrow$ π

IID Prophet Inequality

$\downarrow$ 98 $\downarrow$ $\cancel{10}$ $\downarrow$ $\cancel{\pi}$

IID Prophet Inequality

   

$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$
$\cancel{98}$	$\cancel{10}$	$\cancel{\pi}$	0.4

IID Prophet Inequality, with Random Horizon

Consider the following random horizon.

$$H = \begin{cases} 1 & \text{w.p. } 1/2 \\ 2 & \text{w.p. } 1/4 \\ 3 & \text{w.p. } 1/20 \\ 4 & \text{w.p. } 1/5 \end{cases}$$

Consider iid positive values $X_i \sim X$.
We would like to bound

$$\sup_{ALG} \frac{\mathbb{E}(\mathrm{ALG})}{\mathbb{E}(\mathrm{MAX})}$$

Dynamic

1. Draw a horizon H and variables $X_1, X_2, \ldots, X_H$.
2. At each step i, if $i \leq H$, then observe X_i.
3. Decide whether to take X_i or not.
   If you take it, the process end.
   If you do not take it, the process continuous.

A decision rule defines ALG.
The benchmark is the offline maximum MAX.

Assumptions

- Finite value expectation: $\mathbb{E}(X) < \infty$
- Finite horizon expectation: $\mu := \mathbb{E}(H) < \infty$
- All variables are independent of each other
- Talk simplicity: X is a continuous distribution

Under these assumptions,

$$\mathbb{E}(MAX) \leq \mathbb{E}\left(\sum_{i=1}^{H} X_i\right) \leq \mathbb{E}(H)\mathbb{E}(X) = \mu\,\mathbb{E}(X) < \infty\,.$$

Previous Results

Theorem

If $H \equiv h$, then there exists a threshold $\tau \in \mathbb{R}$ such that

$$\frac{\mathbb{E}(ALG_\tau)}{\mathbb{E}(MAX)} \geq \left(1 - \frac{1}{e}\right) \approx 0.632$$

and the bound is **tight** *considering* **single-threshold** *algorithms.*

In this proof, the threshold τ is defined by

$$\mathbb{P}(\mathrm{MAX} \leq \tau) = \frac{1}{e}$$

Previous Results 2

Define the Hazard rate of a random variable H by

$$\lambda(h) := \mathbb{P}(H = h \mid H \geq h) = \frac{\mathbb{P}(H = h)}{\mathbb{P}(H \geq h)}$$

taking the value 1 when it is not defined.

Theorem

If H is such that $\lambda(\cdot)$ is increasing, then there exists a threshold $\tau \in \mathbb{R}$ such that

$$\frac{\mathbb{E}(ALG_\tau)}{\mathbb{E}(MAX)} \geq \frac{1}{2 - 1/\mu}$$

and the bound is **tight** *considering* **all** *algorithms.*

In this proof, the threshold τ is defined by

$$\mathbb{P}(X \geq \tau) = \frac{1}{\mu}$$

Technical parenthesis: Optimal algorithm

For H with bounded support,
backward induction defines the optimal algorithm.

For H with **un**bounded support,
backward induction is **not** defined.
Because $\mathbb{E}(H) < \infty$, approximate the instance by $\min\{H, n\}$.
For time i, consider the increasing sequence of thresholds $(\tau_{i,n})_{n \geq 1}$.
The optimal algorithm is given by thresholds $\tau_i^* := \lim_{n \to \infty} \tau_{i,n}$.

Example

Recall our example

$$H = \begin{cases} 1 & \text{w.p. } 1/2 \\ 2 & \text{w.p. } 1/4 \\ 3 & \text{w.p. } 1/20 \\ 4 & \text{w.p. } 1/5 \end{cases}$$

The hazard rate is given by

$$\lambda = \left[\frac{1}{2}, \frac{1}{2}, \frac{1}{5}, 1, 1, \dots\right]$$

In particular, this example is **not covered** by previous results.

Result: Concentrated Horizon

Theorem (Concentrated horizons)

If $\sigma^2 := Var(H) < \infty$, then there exists τ such that

$$\frac{\mathbb{E}(ALG_\tau)}{\mathbb{E}(MAX)} \geq \frac{\mu^2}{\mu^2 + \sigma^2} \left(1 - \left(1 - \frac{1}{\mu} \right)^{\frac{\mu^2 + \sigma^2}{\mu}} \right) .$$

In particular, fixing σ^2 and taking $\mu \to \infty$,
we **extend** the factor the factor $(1 - 1/e)$
from fixed to low variance horizons.

Result: Reliable Horizon

Definition ($\mathcal{G}$-class)

A horizon H is in the $\mathcal{G}$-class if,
for all $t \in (0, 1)$,

$$\mathbb{E}(t^H) = \sum_{i \geq 0} \mathbb{P}(H = i) t^i \leq \frac{t}{t + (1 - t)\mu} .$$

Theorem (Better than geometric)

If H is in the $\mathcal{G}$-class, then there exists τ such that

$$\frac{\mathbb{E}(ALG_\tau)}{\mathbb{E}(MAX)} \geq \frac{1}{2 - 1/\mu} .$$

Technique

The general proof is the following.

$$
\begin{aligned}
\mathbb{E}(\mathrm{ALG}_\tau) &= \mathbb{E}\left(1 - \mathbb{P}(X < \tau)^H\right) \mathbb{E}(X \mid X \geq \tau) \\
&= \left(1 - \mathbb{E}\left(\mathbb{P}(X < \tau)^H\right)\right) \mathbb{E}(X \mid X \geq \tau) \\
&\geq \left(1 - \mathbb{E}\left(\mathbb{P}(X < \tau)^H\right)\right) \mathbb{E}(\mathrm{MAX}) && (\text{def. } \tau) \\
&\geq \left(1 - \mathbb{E}\left(\mathbb{P}(X < \tau)^G\right)\right) \mathbb{E}(\mathrm{MAX}) && (\text{def. } G) \\
&= c_G \quad \mathbb{E}(\mathrm{MAX})
\end{aligned}
$$

where c_G is a factor defined by the distribution of G and G is another random variable such that, for all $t \in (0, 1)$

$$\mathbb{E}(t^H) \leq \mathbb{E}(t^G)$$

Threshold for MAX

Lemma

Consider a continuous random variable X and a horizon H with expectation $\mu < \infty$.
Define the the threshold τ by $\mathbb{P}(X \geq \tau) = 1/\mu$.
Then,

$$\mathbb{E}(MAX) \leq \mathbb{E}(X \mid X \geq \tau)$$

Threshold for MAX, proof

Note that, for all x,

$$f_{\mathrm{MAX}}(x) \leq \mu f_X(x) \,.$$

Define the optimization problem

$$\begin{aligned} \max_f \int_0^\infty & x f(x) dx \\ \text{s.t.} \int_0^\infty & f(x) dx \leq 1 \\ & f(x) \leq \mu f_X(x) \qquad \forall x \end{aligned}$$

Threshold for MAX, proof 2

Its solution is to put as much mass in high values as possible, i.e.

$$f^*(x) = \begin{cases} \mu f_X(x) & x \geq x^* \\ 0 & \sim \end{cases}$$

and the threshold x^* is given by f^* being a density, i.e., $x^* = \tau$, so that

$$\int_0^\infty f^*(x)dx = \mu \int_\tau^\infty f_X(x)dx = \mu \mathbb{P}(X \geq \tau) = 1$$

Then,

$$\mathbb{E}(\mathrm{MAX}) \leq \int_0^\infty x f^*(x)dx = \mu \int_\tau^\infty x f_X(x)dx = \mathbb{E}(X \,|\, X \geq \tau)\,.$$

$\mathcal{G}$-class

Definition ($\mathcal{G}$-class)

A horizon H is in the $\mathcal{G}$-class if,
for all $t \in (0,1)$,

$$\mathbb{E}(t^H) = \sum_{i \geq 0} \mathbb{P}(H = i) t^i \leq \frac{t}{t + (1-t)\mu} .$$

Note that, if $G \sim \mathrm{Geo}(1/\mu)$, then

$$\mathbb{E}(t^G) = \frac{t}{t + (1-t)\mu} .$$

The **discounted cost of a failure** $\sim H$ is less than a failure $\sim G$, for all discount factors.

$\mathcal{G}$-class is big

The following classes of random horizons are **nested**.

1. Increasing Hazard Rate (IHR)
2. IHR in expectation
3. Harmonically IHR in expectation
4. Better New than Used (BNU)
5. BNU in expectation
6. Harmonically BNU in expectation
7. $\mathcal{G}$-class (more reliable than a geometric)

Reliable Horizon, proof

$$
\begin{aligned}
&\mathbb{E}(\mathrm{ALG}_\tau) \\
&= \left(1 - \mathbb{E}\left(\mathbb{P}(X < \tau)^H\right)\right) \mathbb{E}(X \mid X \geq \tau) \\
&\geq \left(1 - \mathbb{E}\left(\mathbb{P}(X < \tau)^H\right)\right) \mathbb{E}(\mathrm{MAX}) && (\text{def. } \tau) \\
&\geq \left(1 - \mathbb{E}\left(\mathbb{P}(X < \tau)^G\right)\right) \mathbb{E}(\mathrm{MAX}) && (H \in \mathcal{G}) \\
&= \left(1 - \frac{\mathbb{P}(X < \tau)}{\mathbb{P}(X < \tau) + (1 - \mathbb{P}(X < \tau))\mu}\right) \quad \mathbb{E}(\mathrm{MAX}) && (\text{def. } G) \\
&= \left(1 - \frac{1 - 1/\mu}{1 - 1/\mu + (1 - (1 - 1/\mu))\mu}\right) \quad \mathbb{E}(\mathrm{MAX}) && (\text{def. } \tau) \\
&= \left(\frac{1}{2 - 1/\mu}\right) \quad \mathbb{E}(\mathrm{MAX})
\end{aligned}
$$

Concentrated Horizon, proof

Consider a horizon H such that $\sigma^2 := \mathrm{Var}(H) < \infty$. Define

$$G \sim \frac{\mu^2 + \sigma^2}{\mu} \mathsf{Ber}\left(\frac{\mu^2}{\mu^2 + \sigma^2} \right) .$$

Then, for all $t \in (0, 1)$,

$$\mathbb{E}(t^H) \leq \mathbb{E}(t^G) .$$

Concentrated Horizon, proof 2

$$
\begin{aligned}
\mathbb{E}(\mathrm{ALG}_\tau) &= \left(1 - \mathbb{E}\left(\mathbb{P}(X < \tau)^H\right)\right) \mathbb{E}(X \mid X \geq \tau) \\
&\geq \left(1 - \mathbb{E}\left(\mathbb{P}(X < \tau)^H\right)\right) \mathbb{E}(\mathrm{MAX}) && \text{(def. } \tau\text{)} \\
&\geq \left(1 - \mathbb{E}\left(\mathbb{P}(X < \tau)^G\right)\right) \mathbb{E}(\mathrm{MAX}) && (H \in \mathcal{G}) \\
&\geq \left(1 - \mathbb{E}\left((1 - 1/\mu)^G\right)\right) \mathbb{E}(\mathrm{MAX}) && \text{(def. } \tau\text{)} \\
&= \left(\frac{\mu^2}{\mu^2 + \sigma^2}\left(1 - \left(1 - \frac{1}{\mu}\right)^{\frac{\mu^2+\sigma^2}{\mu}}\right)\right) \mathbb{E}(\mathrm{MAX}) && \text{(def. } G\text{)}
\end{aligned}
$$

Back to the example

Notice that H is such that

$$\mu = \frac{39}{20} = 1.95$$
$$\sigma^2 \approx 1.36$$

Therefore, using τ such that $\mathbb{P}(X \geq \tau) = 1/\mu$, we guarantee that

$$\frac{\mathbb{E}(\mathrm{ALG}_\tau)}{\mathbb{E}(\mathrm{MAX})} \geq \frac{\mu^2}{\mu^2 + \sigma^2} \left(1 - \left(1 - \frac{1}{\mu} \right)^{\frac{\mu^2 + \sigma^2}{\mu}} \right) \approx 0.588$$

Back to the example

Notice that H is in the $\mathcal{G}$-class, and $\mu = 1.95$.
Therefore, using the same threshold τ such that $\mathbb{P}(X \geq \tau) = 1/\mu$, we guarantee that

$$\frac{\mathbb{E}(\mathrm{ALG}_\tau)}{\mathbb{E}(\mathrm{MAX})} \geq \frac{1}{2 - 1/\mu} = \frac{39}{58} \approx 0.672$$

Future work

Instead of a fixed threshold, we can study other algorithms.

For example, a variant of the secretary algorithm achieves a constant factor in instances where fixed threshold does not.

On which class of horizons we have that simple adaptive algorithms achieve a constant factor?

Thank you!