Minimization IID Prophet Inequality via Extreme Value Theory: A Unified Approach

Vasilis Livanos

Center for Mathematical Modeling (CMM), Santiago, Chile

Joint work with

Ruta Mehta, UIUC

July 7th, 2025

Prophet Inequality

[Krengel, Sucheston, 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

Prophet Inequality

[Krengel, Sucheston, 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*: $\mathbb{E}[\max_i X_i]$
- Competitive Ratio:

$$\frac{\mathbb{E}[ALG]}{\mathbb{E}[\max_i X_i]}$$

$U[10, 11]$

$U[9, 12]$

$U[7, 14]$

$$\begin{cases} 1000 & \text{w.p. } 1/100 \\ 0 & \text{otherwise} \end{cases}$$

$U[10, 11]$

$U[9, 12]$

$U[7, 14]$

$\begin{cases} 1000 & \text{w.p. } 1/100 \\ 0 & \text{otherwise} \end{cases}$

$X_1 = 10.30$

$U[10, 11]$	$U[9, 12]$	$U[7, 14]$	$\begin{cases} 1000 & \text{w.p. } 1/100 \\ 0 & \text{otherwise} \end{cases}$

$X_1 = 10.30$

$X_2 = 9.54$

$U[10, 11]$

$U[9, 12]$

$U[7, 14]$

$\begin{cases} 1000 & \text{w.p. } 1/100 \\ 0 & \text{otherwise} \end{cases}$

$X_1 = 10.30$

$X_2 = 9.54$

$X_3 = 8.14$

$U[10, 11]$	$U[9, 12]$	$U[7, 14]$	$\begin{cases} 1000 & \text{w.p. } 1/100 \\ 0 & \text{otherwise} \end{cases}$
$X_1 = 10.30$	$X_2 = 9.54$	$X_3 = 8.14$	$X_4 = 0$

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

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

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

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

- 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 [Wittmann '95, Kleinberg and Weinberg '12]

1. When can we get better guarantees?
   Can $1/2$ be improved?

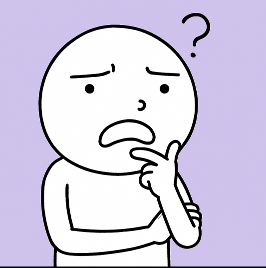

1. When can we get better guarantees? Can $1/2$ be improved?
2. What if $\mathcal{D}_1 = \mathcal{D}_2 = \cdots = \mathcal{D}_n = \mathcal{D}$?
3. Worst-case $\mathcal{D}$?

1. When can we get better guarantees? Can $1/2$ be improved?
2. What if $\mathcal{D}_1 = \mathcal{D}_2 = \cdots = \mathcal{D}_n = \mathcal{D}$?
3. Worst-case $\mathcal{D}$?

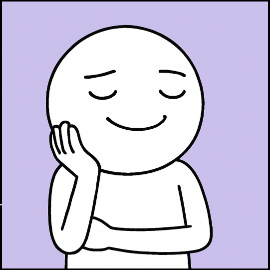

IID Prophet Inequality [Hill, Kertz '82, Correa, Foncea, Hoeksma, Oosterwijk and Vredeveld '21]

For any $\mathcal{D}$, $\exists$ threshold stopping strategy $\tau_1, \tau_2, \ldots, \tau_n$ that achieves $\beta \cdot \mathbb{E}[\max_i X_i]$, where $\beta \approx 0.745$, and this is tight

---

Worst-case $\mathcal{D}$: High variance – depends on n

Most of the mass is at 0 – low probability of getting high values

Minimization

1. Is MIN similar to MAX?
2. Required to select a value

Minimization

1. Is MIN similar to MAX?
2. Required to select a value

- No hope for universal bound: [Lucier '22]

  $\mathcal{D} : F(x) = 1 - 1/x$, with $x \in [1, +\infty)$
  (Equal-revenue distribution)

  $\mathbb{E}[X] = 1 + \int_1^\infty (1 - F(x))\, dx = +\infty$, but

  $\mathbb{E}[\min\{X_1, X_2\}] = 1 + \int_1^\infty (1 - F(x))^2\, dx < +\infty$

Large Market Setting

2. What if $\mathcal{D}$ does not depend on n?

Large Market Setting

2. What if $\mathcal{D}$ does not depend on n?

Large Market Setting: fixed common $\mathcal{D}$ and $n \to \infty$

Large Market Setting

2. What if $\mathcal{D}$ does not depend on n?

Large Market Setting: fixed common $\mathcal{D}$ and $n \to \infty$

Asymptotic Competitive Ratio (ACR)

$$ACR_{Max} = \lim_{n\to\infty} \frac{\mathbb{E}[ALG(n)]}{\mathbb{E}[\max_{i=1}^{n} X_i]} \qquad ACR_{Min} = \lim_{n\to\infty} \frac{\mathbb{E}[ALG(n)]}{\mathbb{E}[\min_{i=1}^{n} X_i]}$$

Towards A Unified Analysis

- $ACR_{Min} = \mathrm{O}\left(1\right)$ for special cases of $\mathcal{D}$
  [L., Mehta '24]

Towards A Unified Analysis

- $ACR_{Min} = \mathrm{O}(1)$ for special cases of $\mathcal{D}$ [L., Mehta '24]

Need more holistic approach

- $M_n = \max\{X_1, \dots, X_n\}$
- $m_n = \min\{X_1, \dots, X_n\}$
- Distribution of M_n, m_n as $n \to \infty$?

Towards A Unified Analysis

- $ACR_{Min} = \mathrm{O}(1)$ for special cases of $\mathcal{D}$ [L., Mehta '24]

Need more holistic approach

- $M_n = \max\{X_1, \ldots, X_n\}$
- $m_n = \min\{X_1, \ldots, X_n\}$
- Distribution of M_n, m_n as $n \to \infty$?
- $\lim_{n\to\infty} M_n = +\infty, \quad \lim_{n\to\infty} m_n = 0 \implies$ Re-scaling

Main Tool: Extreme Value Theory

Extreme Value Theorem [Fisher, Tippett '28, Gnedenko '43]

Assume there exist sequences $a_n > 0, b_n \in \mathbb{R}$ such that

$$\lim_{n \to \infty} F_{M_n}(a_n x + b_n) = G_\gamma^+(x)$$

Then,

$$G_\gamma^+(x) = \begin{cases} \exp\left(-(1+\gamma x)^{-1/\gamma}\right), & \text{if } \gamma \neq 0 \\ \exp\left(-\exp\left(-x\right)\right), & \text{if } \gamma = 0 \end{cases}$$

and we say that $\mathcal{D}$ follows EVT

Main Tool: Extreme Value Theory

Extreme Value Theorem [Fisher, Tippett '28, Gnedenko '43]

Assume there exist sequences $a_n > 0, b_n \in \mathbb{R}$ such that

$$\lim_{n\to\infty} F_{M_n}(a_n x + b_n) = G_\gamma^+(x)$$

Then,

$$G_\gamma^+(x) = \begin{cases} \exp\left(-(1+\gamma x)^{-1/\gamma}\right), & \text{if } \gamma \neq 0 \\ \exp\left(-\exp\left(-x\right)\right), & \text{if } \gamma = 0 \end{cases}$$

and we say that $\mathcal{D}$ follows EVT

- G : Extreme Value Distribution, γ : Extreme Value Index
- Three distinct G_γ^+'s:
  - $\gamma > 0$: Fréchet
  - $\gamma = 0$: Gumbel
  - $\gamma < 0$: Reverse Weibull

Main Tool: Extreme Value Theory

Extreme Value Theorem [Fisher, Tippett '28, Gnedenko '43]

Assume there exist sequences $a_n > 0, b_n \in \mathbb{R}$ such that

$$\lim_{n\to\infty} F_{M_n}(a_n x + b_n) = G_\gamma^+(x)$$

Then,

$$G_\gamma^+(x) = \begin{cases} \exp\left(-(1+\gamma x)^{-1/\gamma}\right), & \text{if } \gamma \neq 0 \\ \exp\left(-\exp\left(-x\right)\right), & \text{if } \gamma = 0 \end{cases}$$

and we say that $\mathcal{D}$ follows EVT

- G : Extreme Value Distribution, γ : Extreme Value Index
- Three distinct G_γ^+'s: $\gamma < 0, \quad \gamma = 0, \quad \gamma > 0$
- Central Limit Theorem analogue for $\textsc{Max}$
- Can get similar result for $\textsc{Min}$, but γ, G_γ^- changes

Fréchet, Gumbel and Rev. Weibull

$$G_\gamma^+(x) = \begin{cases} \exp\left(-(1+\gamma x)^{-1/\gamma}\right), & \text{if } \gamma \neq 0 \\ \exp\left(-\exp\left(-x\right)\right), & \text{if } \gamma = 0 \end{cases}$$

- $\gamma > 0$: Fréchet

$$\lim_{x\to\infty} 1 - G_\gamma^+(x) \sim \frac{1}{(\gamma x)^{1/\gamma}}$$

  Heavy tails, $\nexists$ moments of order $1/\gamma$ and above
  Examples: Cauchy, Pareto, Equal-Revenue, ...

- $\gamma = 0$: Gumbel

$$\lim_{x\to\infty} 1 - G_0^+(x) \sim e^{-x}.$$

  Light tail, exponential-like behaviour
  Examples: Gaussian, Exponential, Gamma, ...

- $\gamma < 0$: Reverse Weibull
  Necessarily bounded support, short tail
  Examples: Uniform, Beta, ...

Comparison with CLT

Central Limit Theorem

For a $\mathcal{D}$ with mean μ and variance σ^2, let $a_n = \sqrt{n\sigma^2}$ and $b_n = n\mu$. If $\mu, \sigma^2 < +\infty$, then

$$\lim_{n\to\infty} \frac{\sum_{i=1}^{n} X_i - b_n}{a_n} = Y \sim \mathcal{N}(0,1)$$

Comparison with CLT

Central Limit Theorem

For a $\mathcal{D}$ with mean μ and variance σ^2, let $a_n = \sqrt{n\sigma^2}$ and $b_n = n\mu$. If $\mu, \sigma^2 < +\infty$, then

$$\lim_{n\to\infty} \frac{\sum_{i=1}^{n} X_i - b_n}{a_n} = Y \sim \mathcal{N}(0, 1)$$

Extreme Value Theorem

For a $\mathcal{D}$ with cdf F, let $b_n = (1/(1-F))^{\leftarrow}(n)$, $a_n = 1/(nf(b_n))$. If

$$\lim_{n\to\infty} \frac{\max\{X_1, \dots, X_n\} - b_n}{a_n} = Y$$

exists, then $Y \sim G_{\gamma}^{+}(x)$

IID PI via Extreme Value Theory

Theorem

Assume there exist sequences $a_n > 0, b_n \in \mathbb{R}$ such that

$$\lim_{n\to\infty} F_{M_n}(a_n x + b_n) = G_\gamma^+(x)$$	$$\lim_{n\to\infty} F_{m_n}(a_n x + b_n) = G_\gamma^-(x)$$
for some γ	for some γ

IID PI via Extreme Value Theory

$\boxed{\Gamma(x) = (x-1)!}$

Theorem

Assume there exist sequences $a_n > 0, b_n \in \mathbb{R}$ such that

$$\lim_{n\to\infty} F_{M_n}(a_n x + b_n) = G_\gamma^+(x)$$

for some γ

$$\lim_{n\to\infty} F_{m_n}(a_n x + b_n) = G_\gamma^-(x)$$

for some γ

Then, the optimal algorithm achieves a competitive ratio, as $n \to \infty$

$$ACR_{Max} = \min\left\{\frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)}, 1\right\}$$

[Kennedy, Kertz '91]

IID PI via Extreme Value Theory

$\Gamma(x) = (x-1)!$

Theorem

Assume there exist sequences $a_n > 0, b_n \in \mathbb{R}$ such that

$$\lim_{n\to\infty} F_{M_n}(a_n x + b_n) = G_\gamma^+(x)$$

for some γ

$$\lim_{n\to\infty} F_{m_n}(a_n x + b_n) = G_\gamma^-(x)$$

for some γ

Then, the optimal algorithm achieves a competitive ratio, as $n \to \infty$

$$ACR_{Max} = \min\left\{\frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)}, 1\right\}$$

[Kennedy, Kertz '91]

$$ACR_{Min} = \max\left\{\frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)}, 1\right\}$$

[L., Mehta '25]

IID PI via Extreme Value Theory

Theorem

$\Gamma(x) = (x-1)!$

Assume there exist sequences $a_n > 0, b_n \in \mathbb{R}$ such that

$$\lim_{n\to\infty} F_{M_n}(a_n x + b_n) = G_\gamma^+(x)$$

for some γ

$$\lim_{n\to\infty} F_{m_n}(a_n x + b_n) = G_\gamma^-(x)$$

for some γ

Then, the optimal algorithm achieves a competitive ratio, as $n \to \infty$

$$ACR_{Max} = \min\left\{\frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)}, 1\right\}$$

[Kennedy, Kertz '91]

$$ACR_{Min} = \max\left\{\frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)}, 1\right\}$$

[L., Mehta '25]

- Distribution-optimal closed form
- Unified analysis of competitive ratio for both MAX and MIN

Asymptotic Competitive Ratio

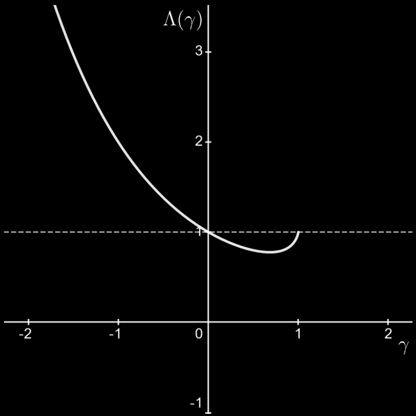

For $\gamma \to -\infty$, by Stirling's approximation

$$\frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)} \approx e^{-\gamma}$$

Asymptotic Competitive Ratio

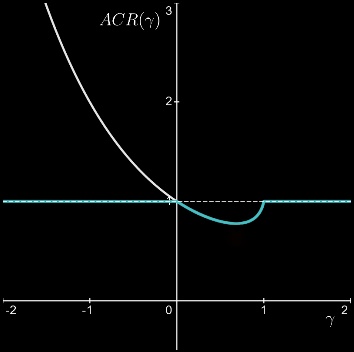

Figure: ACR(γ) for MAX

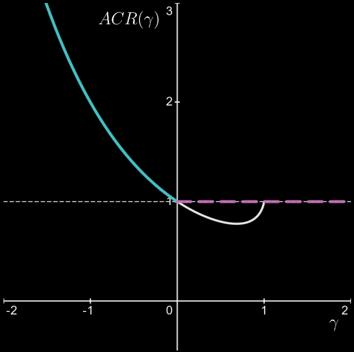

Figure: ACR(γ) for MIN

High-Level Approach

$F(t) = \Pr_{X \sim \mathcal{D}}[X \leq t]$
$F^{\leftarrow}(p)$: Inverse of F ("Quantile function")

Using EVT and heavy-machinery from theory of regularly-varying functions:

Max	Min

High-Level Approach

$F(t) = \Pr_{X\sim\mathcal{D}}[X \leq t]$
$F^{\leftarrow}(p)$: Inverse of F ("Quantile function")

Using EVT and heavy-machinery from theory of regularly-varying functions:

Max

$$\mathbb{E}[ALG(n)] \approx F^{\leftarrow}\left(1 - \frac{1-\gamma}{n}\right)$$

Min

$$\mathbb{E}[ALG(n)] \approx F^{\leftarrow}\left(\frac{1-\gamma}{n}\right)$$

High-Level Approach

$$\Gamma(x) = (x-1)!$$

$F(t) = \Pr_{X\sim\mathcal{D}}[X \le t]$
$F^{\leftarrow}(p)$: Inverse of F ("Quantile function")

Using EVT and heavy-machinery from theory of regularly-varying functions:

Max

$$\mathbb{E}[ALG(n)] \approx F^{\leftarrow}\left(1 - \frac{1-\gamma}{n}\right)$$

$$\mathbb{E}[M_n] \approx \Gamma(1-\gamma)\ F^{\leftarrow}\left(1 - \frac{1}{n}\right)$$

Min

$$\mathbb{E}[ALG(n)] \approx F^{\leftarrow}\left(\frac{1-\gamma}{n}\right)$$

$$\mathbb{E}[m_n] \approx \Gamma(1-\gamma)\ F^{\leftarrow}\left(\frac{1}{n}\right)$$

High-Level Approach

$$\Gamma(x) = (x-1)!$$

$F(t) = \Pr_{X\sim\mathcal{D}}[X \le t]$
$F^{\leftarrow}(p)$: Inverse of F ("Quantile function")

Using EVT and heavy-machinery from theory of regularly-varying functions:

Max

$$\mathbb{E}[ALG(n)] \approx F^{\leftarrow}\left(1 - \frac{1-\gamma}{n}\right)$$

$$\mathbb{E}\left[M_n\right] \approx \Gamma(1-\gamma)\ F^{\leftarrow}\left(1 - \frac{1}{n}\right)$$

$$F^{\leftarrow}\left(1 - \frac{c}{n}\right) \approx c^{-\gamma} F^{\leftarrow}\left(1 - \frac{1}{n}\right)$$

Min

$$\mathbb{E}[ALG(n)] \approx F^{\leftarrow}\left(\frac{1-\gamma}{n}\right)$$

$$\mathbb{E}\left[m_n\right] \approx \Gamma(1-\gamma)\ F^{\leftarrow}\left(\frac{1}{n}\right)$$

$$F^{\leftarrow}\left(\frac{c}{n}\right) \approx c^{-\gamma} F^{\leftarrow}\left(\frac{1}{n}\right)$$

Single-Threshold and Cardinality Constraints

Theorem

If $\mathcal{D}$ satisfies the EVT, then $\exists$ single threshold T s.t.

$$ACR_{Max}(T) = g(\gamma) = \Omega(1)$$

[Correa, Pizarro, Verdugo '21]

$$ACR_{Min}(T) = \mathrm{O}\left((\log n)^{-\gamma}\right)$$

[L., Mehta '25]

Theorem [L., Mehta '25]

If $\mathcal{D}$ satisfies the EVT, then, for $k = \Omega(\log n)$, $\exists$ a threshold T_k achieving a competitive ratio

$$ACR_{Min}(T_k) = e^{-\gamma(1-\gamma)}$$

Open Problems

- Are there $\mathcal{D}_i$ for which we can get constant approximation in the non-IID setting?
- What can you get with $1 < k < n$ thresholds?

Thank You!

Questions?