Towards a Unified Theory of IID Prophet Inequalities

Vasilis Livanos

University of Illinois Urbana-Champaign

Partly joint work with

Ruta Mehta

October 6th, 2023

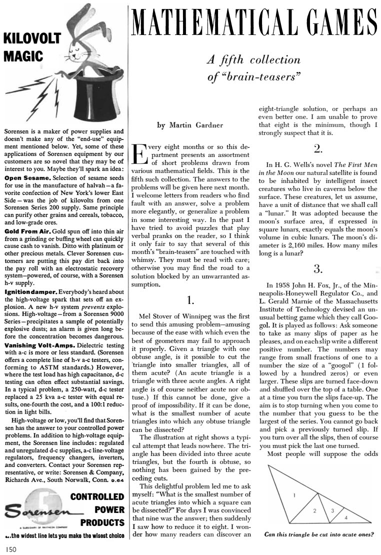

MATHEMATICAL GAMES

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?

150

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

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

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

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

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

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

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

$$1 \quad 4 \quad \dots \quad -2 \qquad\qquad 0 \quad 7 \quad \dots \quad 11$$

S_1	S_2
Sampling Phase	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

1 4 ... −2 | 0 7 ... 11

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*: $\mathbb{E}[\max_i X_i]$.

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

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

$X_1 = 2.34$

$\mathcal{U}[2,4]$	$\mathcal{U}[2,4]$	$\mathcal{U}[1,5]$	$\mathcal{U}[0,7]$
$X_1 = 2.34$	$X_2 = 3.12$		

$\mathcal{U}[2,4]$	$\mathcal{U}[2,4]$	$\mathcal{U}[1,5]$	$\mathcal{U}[0,7]$
$X_1 = 2.34$	$X_2 = 3.12$	$X_3 = 3.20$	

$\mathcal{U}[2,4]$	$\mathcal{U}[2,4]$	$\mathcal{U}[1,5]$	$\mathcal{U}[0,7]$
$X_1 = 2.34$	$X_2 = 3.12$	$X_3 = 3.20$	$X_4 = 0.87$

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.

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

$\mathbb{E}[\mathrm{ALG}] = 1$ 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?

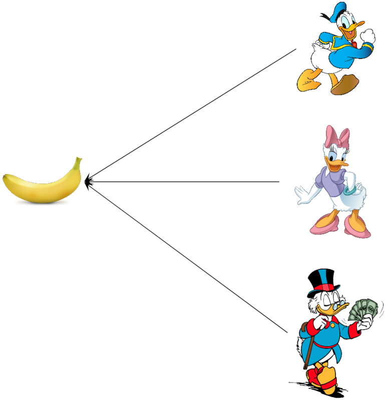

Why should we care?

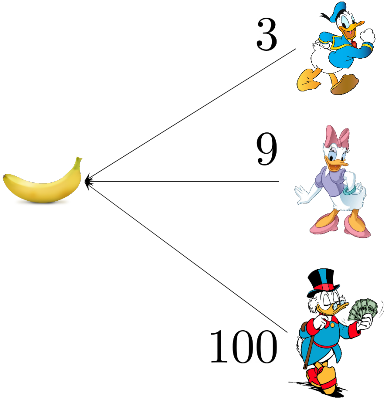

Why should we care?

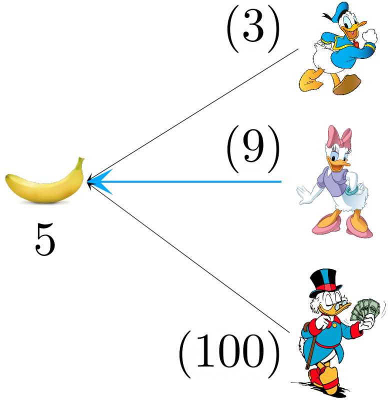

Price $p \iff$ Threshold T
in Prophet Inequality

Why should we care?

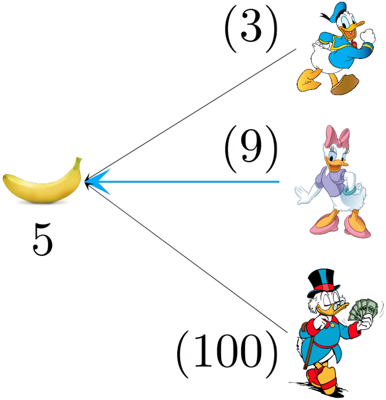

Price p $\iff$ Threshold T
in Prophet Inequality

What about maximizing revenue?
Use "virtual valuations" to design p: $\phi(v) = v - \frac{1-F(v)}{f(v)}$.
[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]
  - [Babaioff, Immorlica, Lucier and Weinberg '14]
  - [Dütting, Feldman, Kesselheim and Lucier '16]
  - [Correa, Foncea, Pizarro and Verdugo '19]
  - [Correa and Cristi '23]

IID Prophet Inequality

Fundamental questions:

1. What if $\mathcal{D}_1 = \mathcal{D}_2 = \cdots = \mathcal{D}_n = \mathcal{D}$?
   Can $1/2$ be improved?
2. Optimal (online) algorithm?
3. Worst-case $\mathcal{D}$?

IID Prophet Inequality

Fundamental questions:

1. What if $\mathcal{D}_1 = \mathcal{D}_2 = \cdots = \mathcal{D}_n = \mathcal{D}$?
   Can $1/2$ be improved?
2. Optimal (online) algorithm?
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 a high value.

Optimal Threshold DP

If we reach X_n, take it. Focus on X_{n-1}, X_n.
What should τ_{n-1} be?

Optimal Threshold DP

If we reach X_n, take it. Focus on X_{n-1}, X_n.
What should τ_{n-1} be?

$$\mathbb{E}[\mathrm{OPTALG}_{n-1,n}] = (1 - F(\tau_{n-1}))\,\mathbb{E}\left[X \mid X \geq \tau_{n-1}\right] + F(\tau_{n-1})\,\mathbb{E}[X]$$
$$\implies \tau_{n-1} = \mathbb{E}[X].$$

In general, we have $\tau_i = \mathbb{E}[\mathrm{OPTALG}_{i+1,\ldots,n}]$.

Cost Minimization

What if objective is $\min_i X_i$? Same problem?

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

Cost Minimization

What if objective is $\min_i X_i$? Same problem?

- Objective: Minimize selected value, compare against $\mathbb{E}[\min_i X_i]$.
- Forced to select an element.
- No bounded approximation for adversarial or random order (non-IID setting)!

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

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

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

Cost Minimization

What if objective is $\min_i X_i$? Same problem?

- Objective: Minimize selected value, compare against $\mathbb{E}[\min_i X_i]$.
- Forced to select an element.
- No bounded approximation for adversarial or random order (non-IID setting)!

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

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

- What about I.I.D.?

  ~~Intuition~~ False Intuition:

  Set $T = 2 \cdot \mathbb{E}[\min_i X_i]$.
  - Doesn't work! $\Pr[\text{We are forced to select } X_n] \geq c$.
  - Optimal single threshold $T \implies \Theta(\text{polylog } n)$-approximation. [L.-Mehta '22]

Is Cost Minimization hopeless?

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

Is Cost Minimization hopeless?

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

Idea

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

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

for $\textsc{Max}$

$$r(x) = \frac{f(x)}{F(x)}$$

for $\textsc{Min}$

Intuition: $h(x) = \Pr[X = x \mid X \geq x]$, $r(x) = \Pr[X = x \mid X \leq x]$ (for discrete distributions).

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]

In the Min-PI setting, for every distribution,

- If $\mathbb{E}[X] = +\infty$, the competitive ratio is infinite.
- If $\mathbb{E}[X] < +\infty$, there exists a constant c-approximate cost minimization prophet inequality, and we characterize the *optimal* c as the solution to a simple inequality.
- Closed form for c for special subclass of distributions.

Some observations:

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

Why $c = +\infty$ when $\mathbb{E}[X] = +\infty$?

Equal-Revenue (Pareto) Distribution:

$F(x) = 1 - 1/x$, with $x \in [1, +\infty)$.

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

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

- Due to [Lucier '22].

What if $\mathcal{D}$ is *independent* of n?
What is c if we fix $\mathcal{D}$ and take $n \to \infty$?

What if $\mathcal{D}$ is *independent* of n?
What is c if we fix $\mathcal{D}$ and take $n \to \infty$?

[Kennedy-Kertz '91]

For any $\mathcal{D}$, independent of n, $\exists$ threshold stopping strategy $\tau_1, \tau_2, \ldots, \tau_n$ that achieves $\lambda \cdot \mathbb{E}[\max_i X_i]$, as $n \to \infty$, where $\lambda \approx 0.776$, and this is tight.

What if $\mathcal{D}$ is *independent* of n?
What is c if we fix $\mathcal{D}$ and take $n \to \infty$?

[Kennedy-Kertz '91]

For any $\mathcal{D}$, independent of n, $\exists$ threshold stopping strategy $\tau_1, \tau_2, \ldots, \tau_n$ that achieves $\lambda \cdot \mathbb{E}[\max_i X_i]$, as $n \to \infty$, where $\lambda \approx 0.776$, and this is tight.

[Braun-Buttkus-Kesselheim '21]

For any MHR $\mathcal{D}$, independent of n, $\exists$ threshold stopping strategy $\tau_1, \tau_2, \ldots, \tau_n$ that achieves $1 \cdot \mathbb{E}[\max_i X_i]$, as $n \to \infty$.

What if $\mathcal{D}$ is *independent* of n?
What is c if we fix $\mathcal{D}$ and take $n \to \infty$?

[Kennedy-Kertz '91]

For any $\mathcal{D}$, independent of n, $\exists$ threshold stopping strategy $\tau_1, \tau_2, \ldots, \tau_n$ that achieves $\lambda \cdot \mathbb{E}[\max_i X_i]$, as $n \to \infty$, where $\lambda \approx 0.776$, and this is tight.

[Braun-Buttkus-Kesselheim '21]

For any MHR $\mathcal{D}$, independent of n, $\exists$ threshold stopping strategy $\tau_1, \tau_2, \ldots, \tau_n$ that achieves $1 \cdot \mathbb{E}[\max_i X_i]$, as $n \to \infty$.

Let $M_n = \max\{X_1, \ldots, X_n\}, m_n = \min\{X_1, \ldots, X_n\}$.

Limiting distribution of M_n or m_n as $n \to \infty$?

Clearly, $\lim_{n\to\infty} M_n = +\infty$ and $\lim_{n\to\infty} m_n = 0$, thus we need rescaling.

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,

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

where $\gamma = \lim_{x\to\infty} \left(\frac{1}{h(x)}\right)'$.

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,

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

where $\gamma = \lim_{x\to\infty} \left(\frac{1}{h(x)}\right)'$.

- Analogue of Central Limit Theorem for MAX instead of averages.
- G : Extreme Value Distribution, γ : Extreme Value Index
- Can get similar result for MIN by $G^-(x) = 1 - G^+(-x)$, but γ changes.

IID PI via Extreme Value Theory

Theorem [L. '23]

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

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

with $\gamma = \lim_{x\to\infty} \left(\frac{1}{h(x)}\right)'$.

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

with $\gamma = \lim_{x\to 0^+} \left(\frac{1}{r(x)}\right)'$.

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

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

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

IID PI via Extreme Value Theory

Theorem [L. '23]

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

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

with $\gamma = \lim_{x\to\infty} \left(\frac{1}{h(x)}\right)'$.

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

with $\gamma = \lim_{x\to 0^+} \left(\frac{1}{r(x)}\right)'$.

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

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

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

- Unified analysis of competitive ratio for both MAX and MIN.
- Recovers and generalizes previous results of [Kennedy-Kertz '91] and [Braun-Buttkus-Kesselheim '21].
- Competitive ratio has same form for both MAX and MIN!

Asymptotic Competitive Ratio

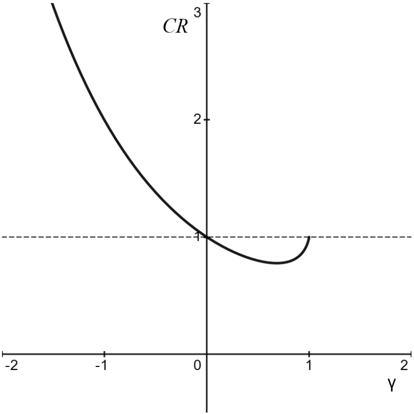

For $\gamma \to -\infty$, by Stirling's approximation

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

Asymptotic Competitive Ratio

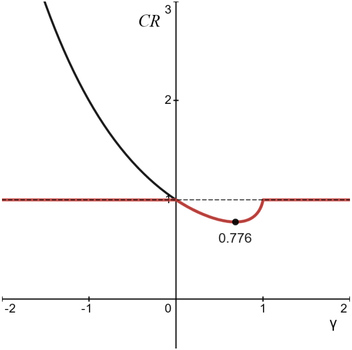

Figure: ACR(γ) for MAX

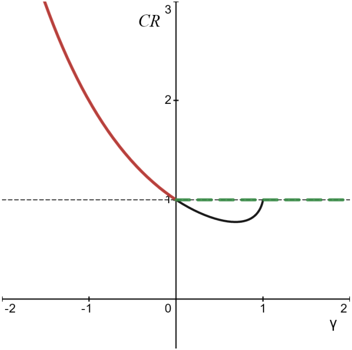

Figure: ACR(γ) for MIN

Open Problems

- Extend Min-PI to multiple selection.
- 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?

Questions?