LOOKBACK PROPHET INEQUALITIES

ZIYAD BENOMAR

ENSAE CRITEO Inria

DORIAN BAUDRY

UNIVERSITY OF OXFORD

VIANNEY PERCHET

ENSAE CRITEO Inria

Prophet Inequality

Setting: Known probability distributions $F_1, \ldots, F_n$

Realisations $X_1 \sim F_1, \ldots, X_n \sim F_n$ observed sequentially

If algorithm ALG stops at time $\tau \in [n]$, its payoff is

$$\text{ALG}(X_1, \ldots, X_n) = X_\tau$$

Prophet Inequality

Setting: Known probability distributions $F_1, \ldots, F_n$

Realisations $X_1 \sim F_1, \ldots, X_n \sim F_n$ observed sequentially

If algorithm ALG stops at time $\tau \in [n]$, its payoff is

$$\text{ALG}(X_1, \ldots, X_n) = X_\tau$$

Prophet Inequality

Setting: Known probability distributions $F_1, \ldots, F_n$

Realisations $X_1 \sim F_1, \ldots, X_n \sim F_n$ observed sequentially

If algorithm ALG stops at time $\tau \in [n]$, its payoff is

$$\text{ALG}(X_1, \ldots, X_n) = X_\tau$$

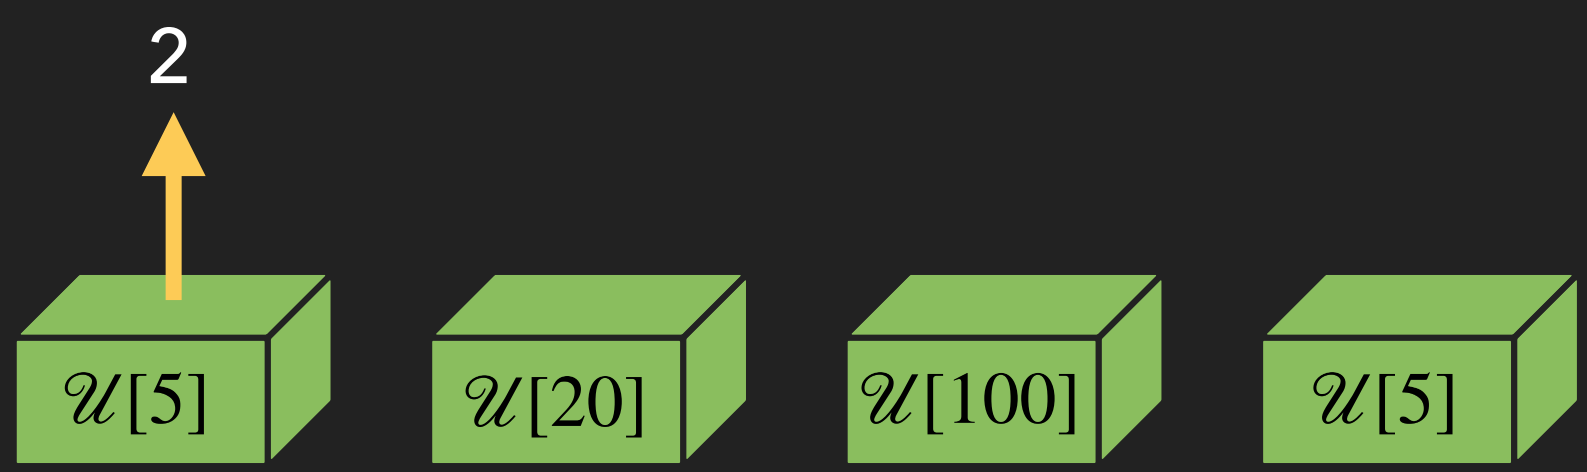

Prophet Inequality

Setting: Known probability distributions $F_1, \ldots, F_n$

Realisations $X_1 \sim F_1, \ldots, X_n \sim F_n$ observed sequentially

If algorithm ALG stops at time $\tau \in [n]$, its payoff is

$$\text{ALG}(X_1, \ldots, X_n) = X_\tau$$

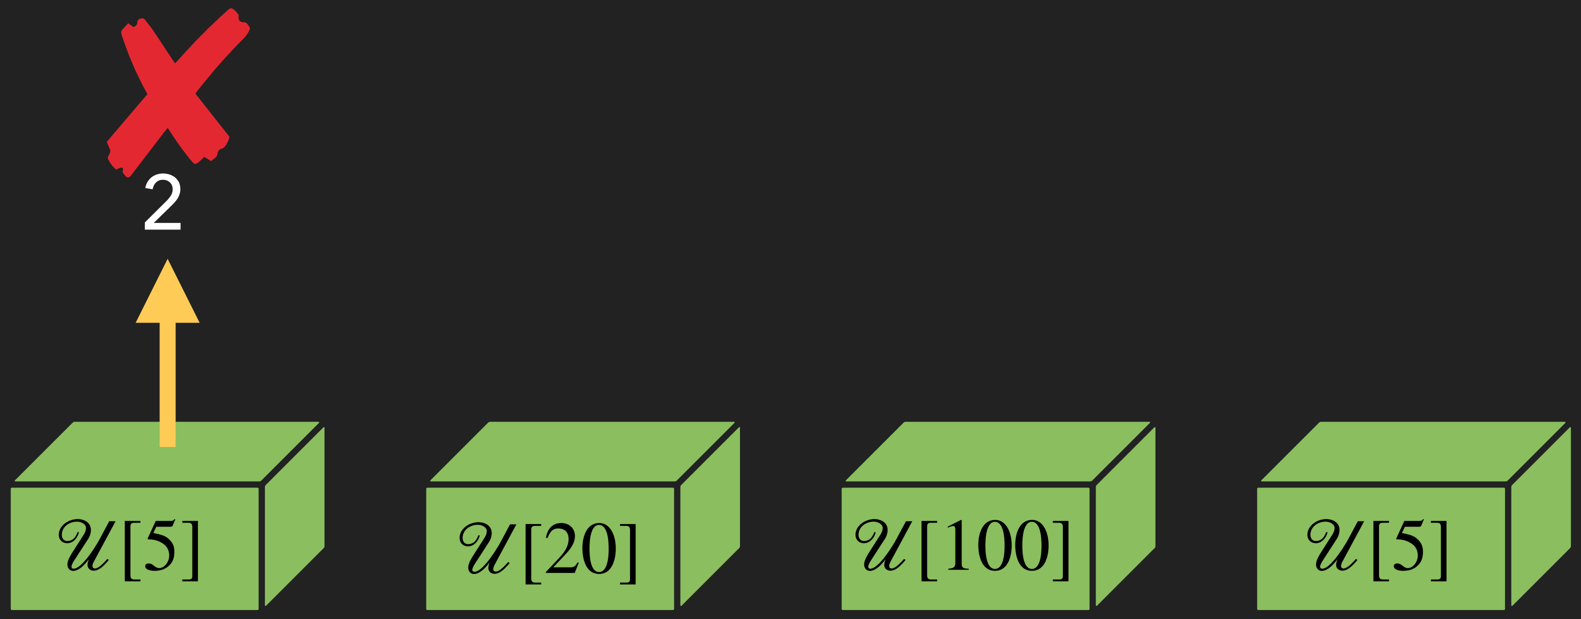

Prophet Inequality

Setting: Known probability distributions $F_1, \ldots, F_n$

Realisations $X_1 \sim F_1, \ldots, X_n \sim F_n$ observed sequentially

If algorithm ALG stops at time $\tau \in [n]$, its payoff is

$$\text{ALG}(X_1, \ldots, X_n) = X_\tau$$

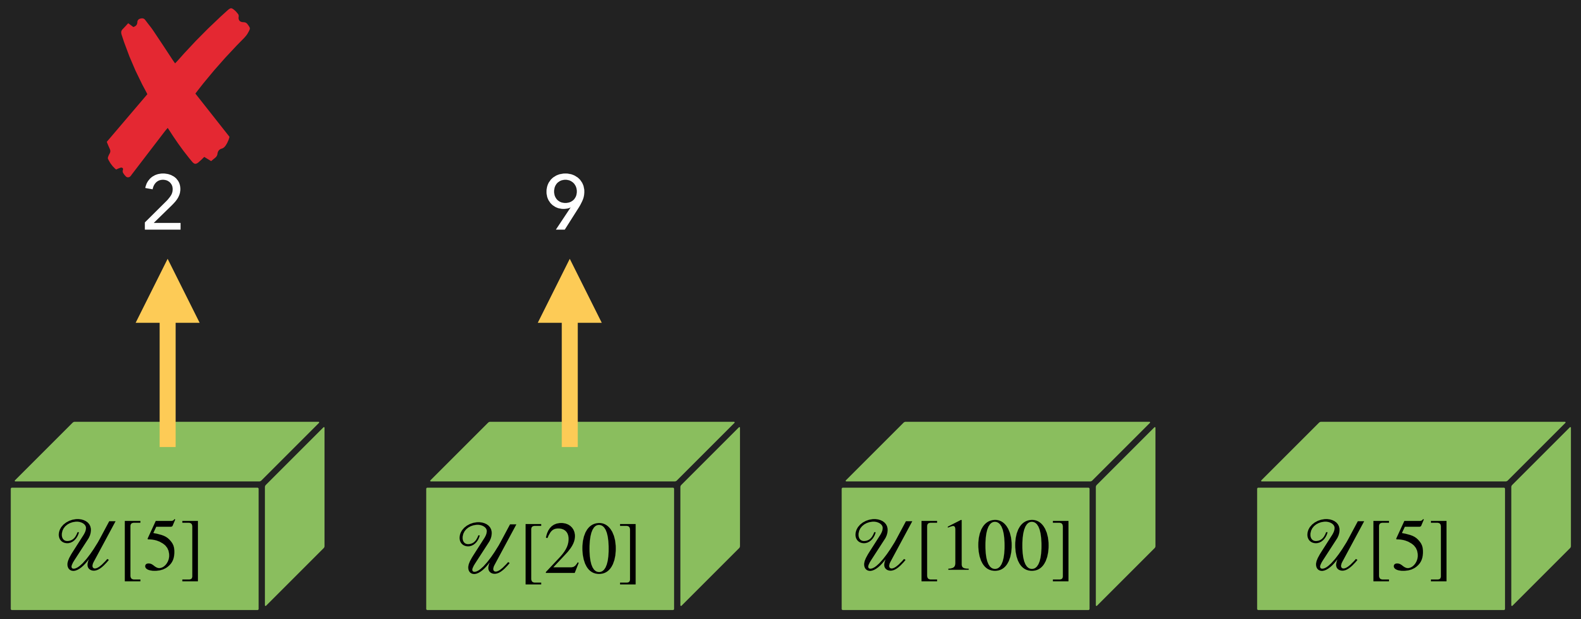

Prophet Inequality

Setting: Known probability distributions $F_1, \ldots, F_n$

Realisations $X_1 \sim F_1, \ldots, X_n \sim F_n$ observed sequentially

If algorithm ALG stops at time $\tau \in [n]$, its payoff is

$$\text{ALG}(X_1, \ldots, X_n) = X_\tau$$

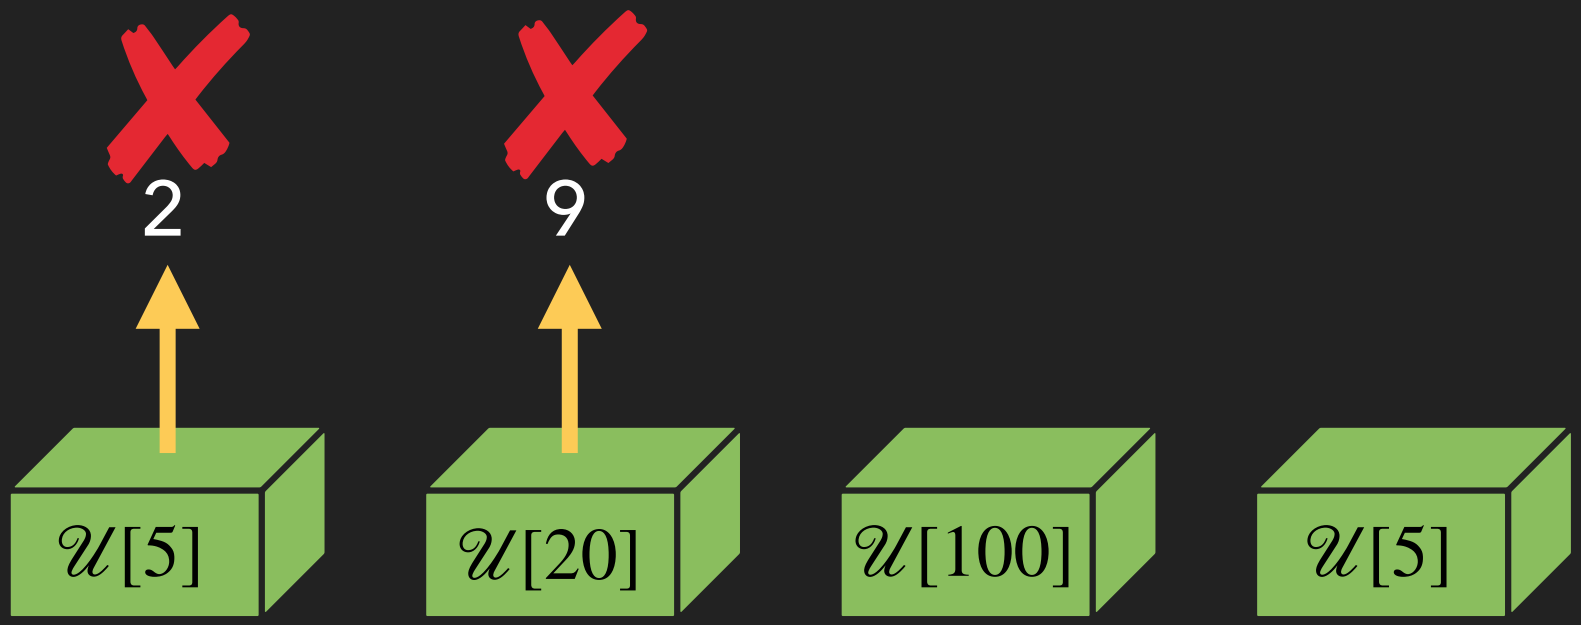

Prophet Inequality

Setting: Known probability distributions $F_1, \ldots, F_n$

Realisations $X_1 \sim F_1, \ldots, X_n \sim F_n$ observed sequentially

If algorithm ALG stops at time $\tau \in [n]$, its payoff is

$$\text{ALG}(X_1, \ldots, X_n) = X_\tau$$

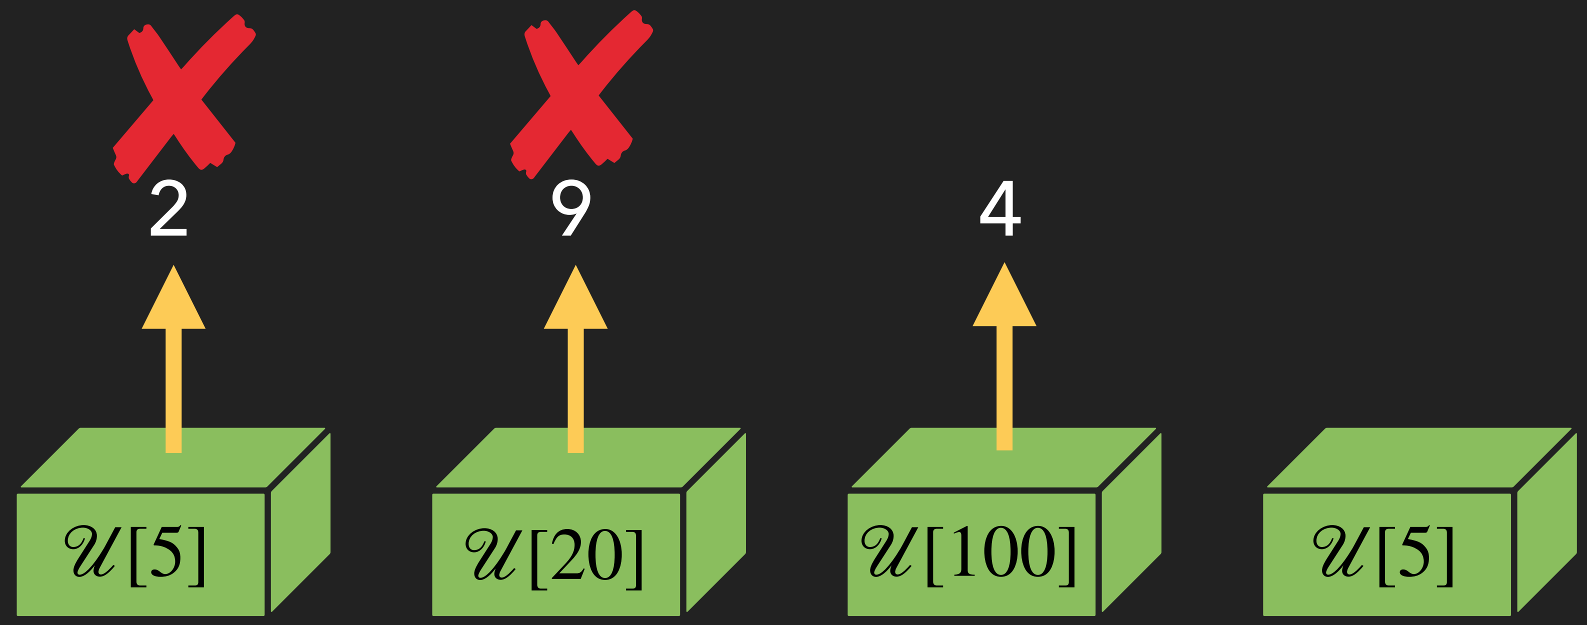

Prophet Inequality

Setting: Known probability distributions $F_1, \ldots, F_n$

Realisations $X_1 \sim F_1, \ldots, X_n \sim F_n$ observed sequentially

If algorithm ALG stops at time $\tau \in [n]$, its payoff is

$$\text{ALG}(X_1, \ldots, X_n) = X_\tau$$

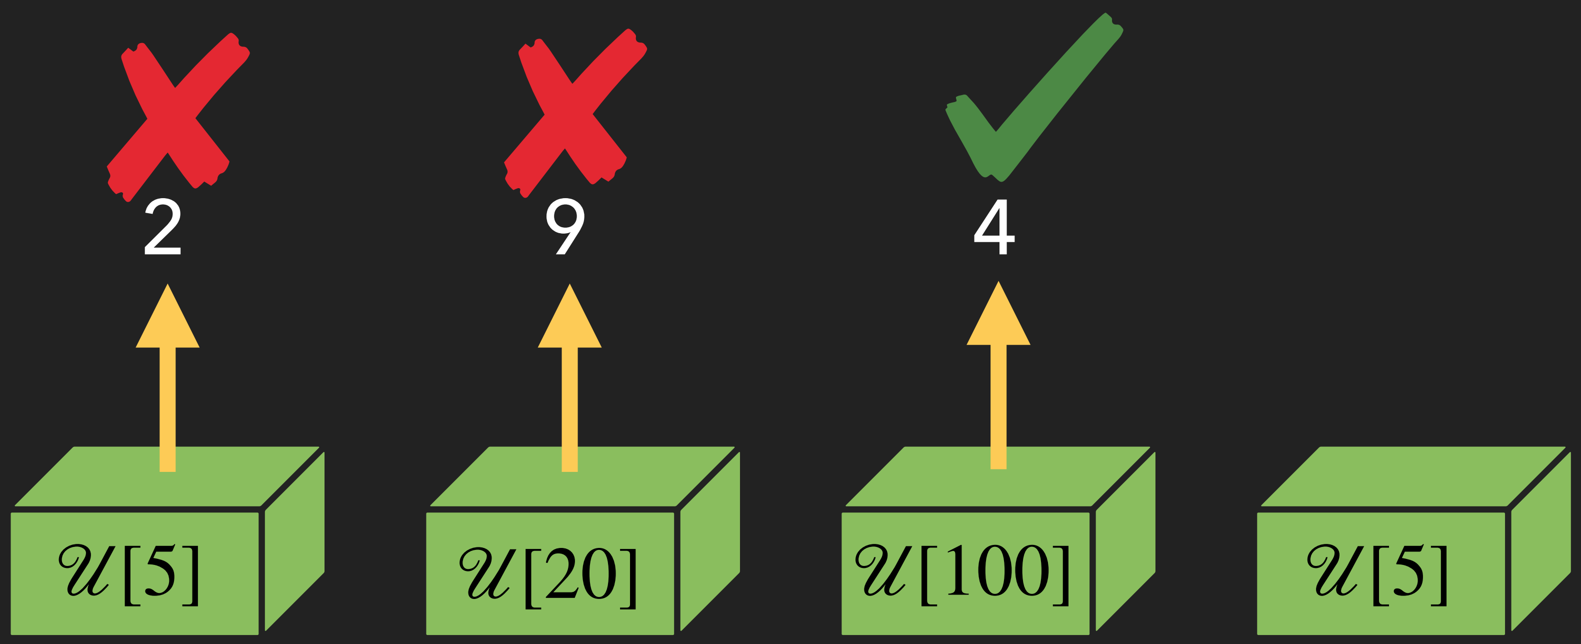

$$\text{ALG}(X_1, \ldots, X_5) = 4$$

Irrevocable decisions!

Prophet Inequality

Setting: Known probability distributions $F_1, \ldots, F_n$

Realisations $X_1 \sim F_1, \ldots, X_n \sim F_n$ observed sequentially

If algorithm ALG stops at time $\tau \in [n]$, its payoff is

$$\text{ALG}(X_1, \ldots, X_n) = X_\tau$$

Definition: The competitive ratio of ALG is

$$\text{CR(ALG)} = \inf_{F_1, \ldots, F_n} \frac{\mathbb{E}[\text{ALG}(X_1, \ldots, X_n)]}{\mathbb{E}[\max(X_1, \ldots, X_n)]}$$

Prophet Inequality

Setting: Known probability distributions $F_1, \ldots, F_n$

Realisations $X_1 \sim F_1, \ldots, X_n \sim F_n$ observed sequentially

If algorithm ALG stops at time $\tau \in [n]$, its payoff is

$$\text{ALG}(X_1, \ldots, X_n) = X_\tau$$

Definition: The competitive ratio of ALG is

$$\text{CR(ALG)} = \inf_{F_1,\ldots,F_n} \frac{\mathbb{E}[\text{ALG}(X_1, \ldots, X_n)]}{\mathbb{E}[\max(X_1, \ldots, X_n)]}$$

Theorem: The best possible competitive ratio is 1/2.

ALG: Select the first value exceeding θ, with $\mathbb{P}(\max(X_1, \ldots, X_n) \geq \theta) = \frac{1}{2}$.

Lookback Prophet Inequality

Weakness: The model is too pessimistic,
rejections are often reversible in real-world scenarios

Lookback Prophet Inequality

Weakness: The model is too pessimistic,

rejections are often reversible in real-world scenarios

Definition: Decay functions $\mathscr{D} = (D_1, D_2, \ldots)$

If ALG stops at time τ, its reward is

$$\mathsf{ALG}^{\mathscr{D}}(X_1, \ldots, X_n) = \max(X_\tau\,, D_1(X_{\tau-1})\,, D_2(X_{\tau-2}), \ldots)$$

Lookback Prophet Inequality

Weakness: The model is too pessimistic,
rejections are often reversible in real-world scenarios

Definition: Decay functions $\mathscr{D} = (D_1, D_2, \ldots)$

If ALG stops at time τ, its reward is

$$\text{ALG}^{\mathscr{D}}(X_1, \ldots, X_n) = \max(X_\tau, D_1(X_{\tau-1}), D_2(X_{\tau-2}), \ldots)$$

Lookback Prophet Inequality

Weakness: The model is too pessimistic,
rejections are often reversible in real-world scenarios

Definition: Decay functions $\mathscr{D} = (D_1, D_2, \ldots)$

If ALG stops at time τ, its reward is

$$\text{ALG}^{\mathscr{D}}(X_1, \ldots, X_n) = \max(X_\tau\,, D_1(X_{\tau-1})\,, D_2(X_{\tau-2}), \ldots)$$

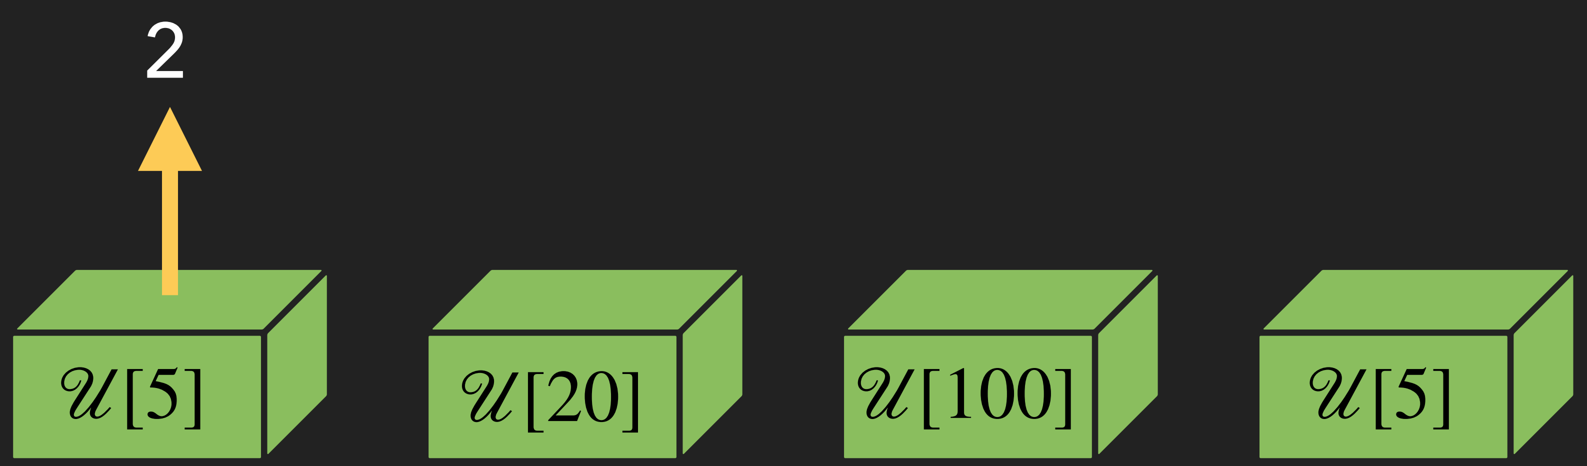

Lookback Prophet Inequality

Weakness: The model is too pessimistic, rejections are often reversible in real-world scenarios

Definition: Decay functions $\mathscr{D} = (D_1, D_2, \ldots)$

If ALG stops at time τ, its reward is

$$\text{ALG}^{\mathscr{D}}(X_1, \ldots, X_n) = \max(X_\tau, D_1(X_{\tau-1}), D_2(X_{\tau-2}), \ldots)$$

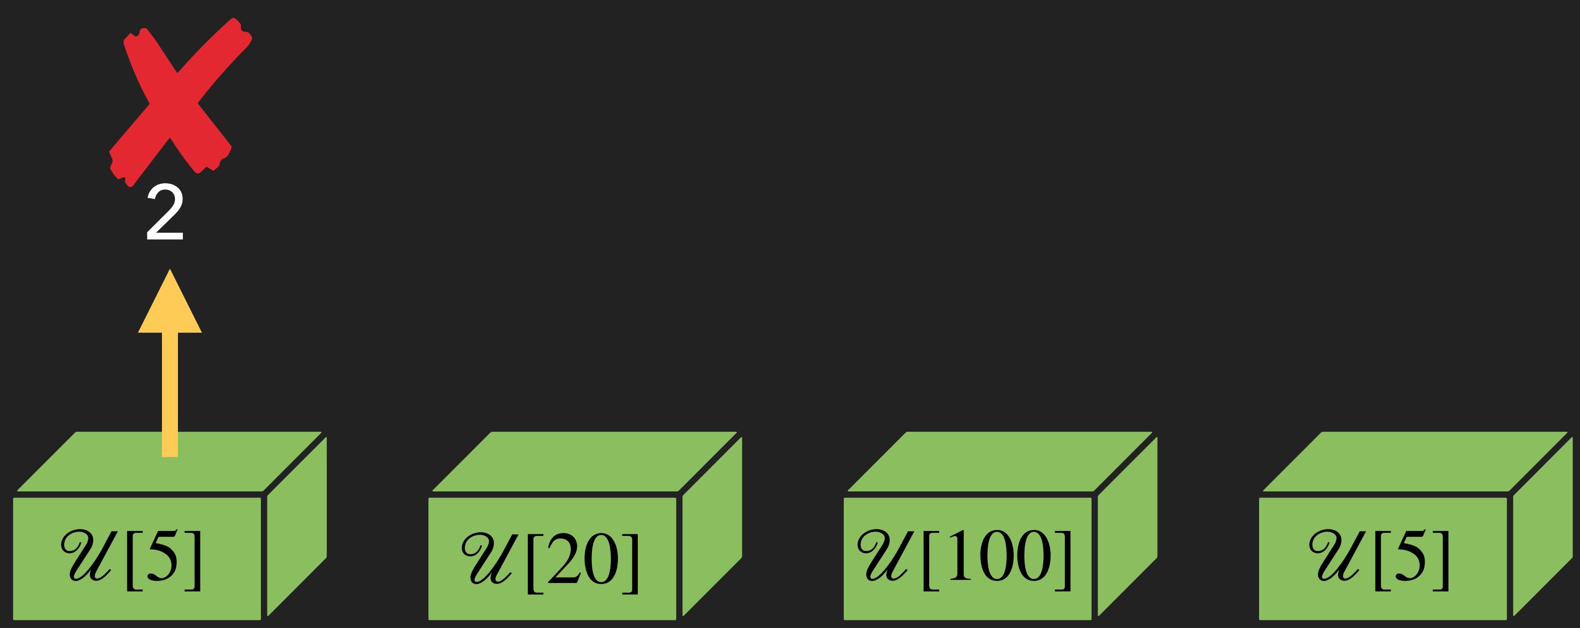

Lookback Prophet Inequality

Weakness: The model is too pessimistic, rejections are often reversible in real-world scenarios

Definition: Decay functions $\mathscr{D} = (D_1, D_2, \ldots)$

If ALG stops at time τ, its reward is

$$\text{ALG}^{\mathscr{D}}(X_1, \ldots, X_n) = \max(X_\tau, D_1(X_{\tau-1}), D_2(X_{\tau-2}), \ldots)$$

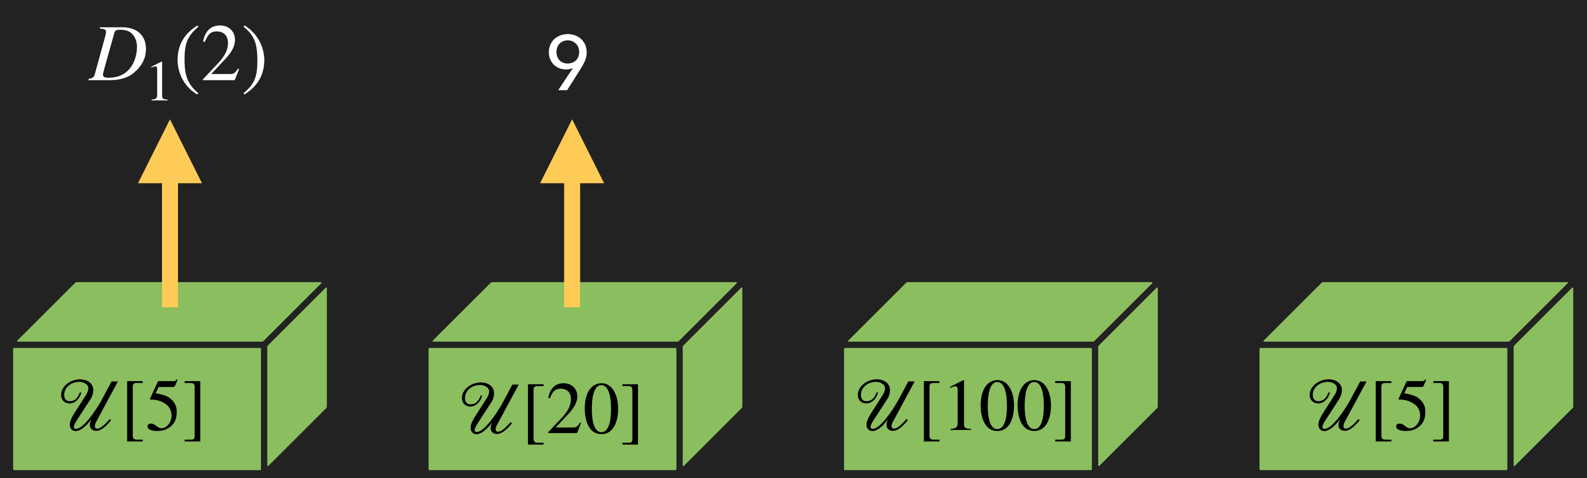

Lookback Prophet Inequality

Weakness: The model is too pessimistic,
rejections are often reversible in real-world scenarios

Definition: Decay functions $\mathscr{D} = (D_1, D_2, \dots)$

If ALG stops at time τ, its reward is

$$\text{ALG}^{\mathscr{D}}(X_1, \dots, X_n) = \max(X_\tau, D_1(X_{\tau-1}), D_2(X_{\tau-2}), \dots)$$

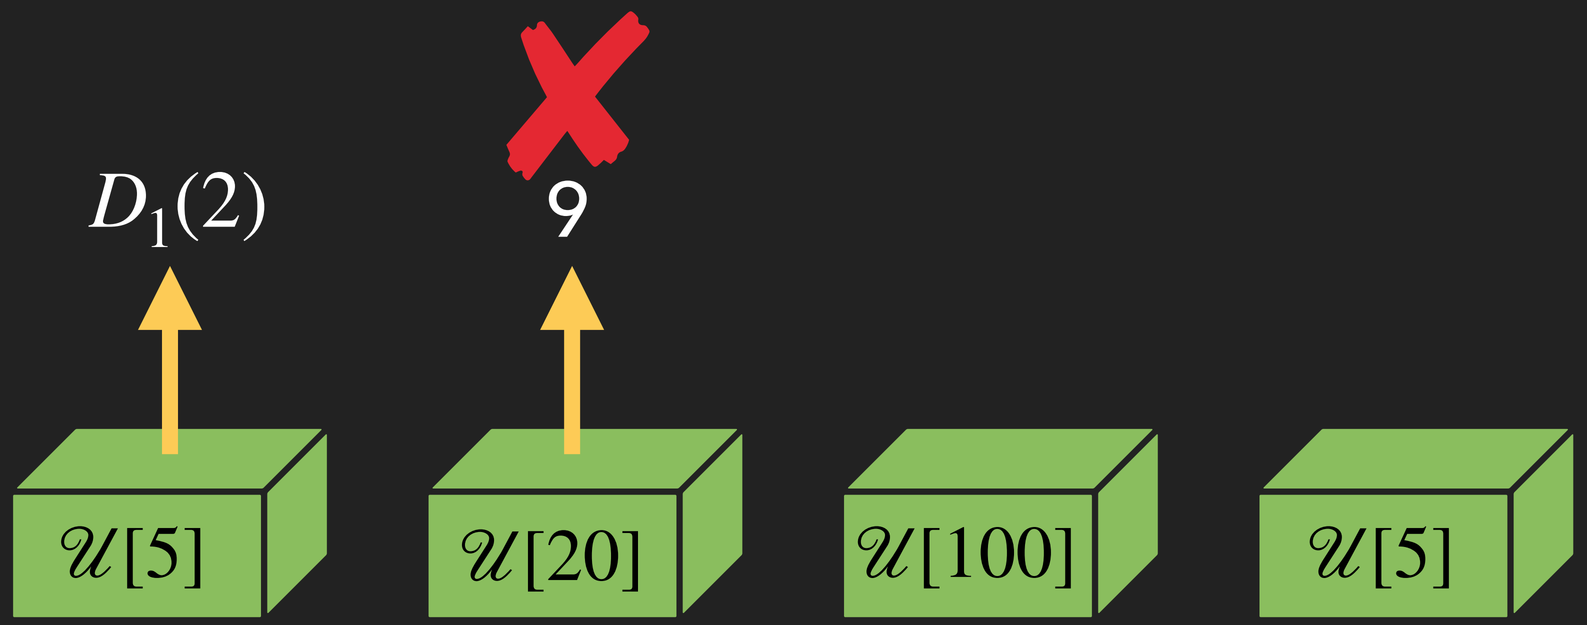

Lookback Prophet Inequality

Weakness: The model is too pessimistic,
rejections are often reversible in real-world scenarios

Definition: Decay functions $\mathscr{D} = (D_1, D_2, \ldots)$

If ALG stops at time τ, its reward is

$$\text{ALG}^{\mathscr{D}}(X_1, \ldots, X_n) = \max(X_\tau, D_1(X_{\tau-1}), D_2(X_{\tau-2}), \ldots)$$

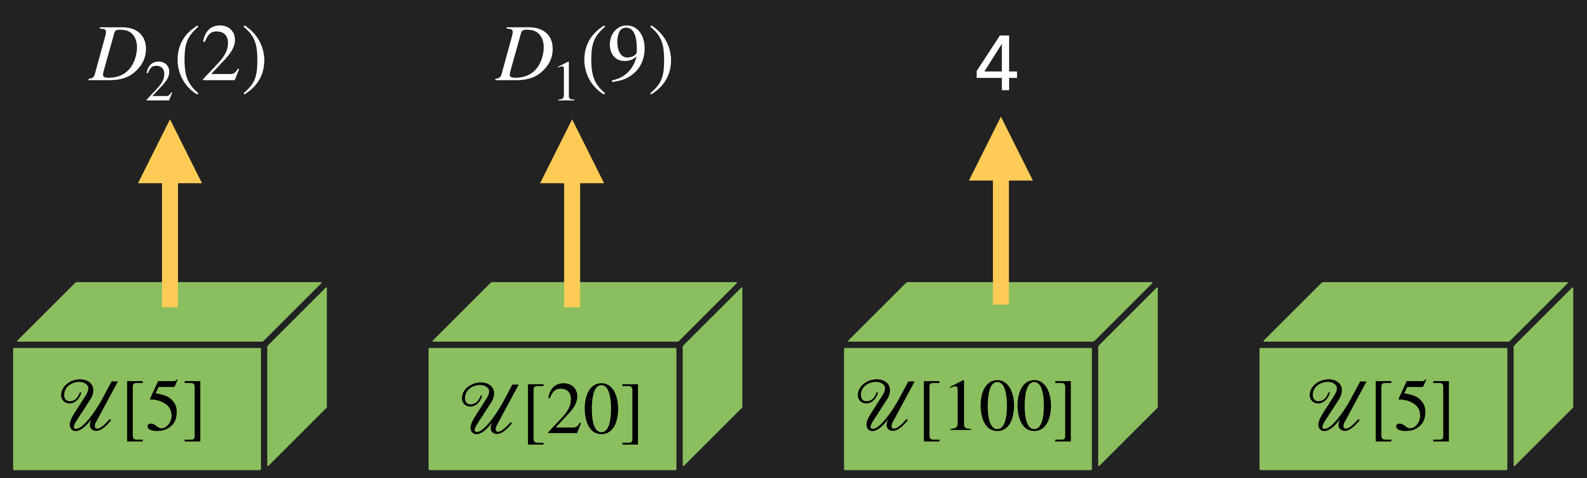

Lookback Prophet Inequality

Weakness: The model is too pessimistic,
rejections are often reversible in real-world scenarios

Definition: Decay functions $\mathscr{D} = (D_1, D_2, \ldots)$

If ALG stops at time τ, its reward is

$$\text{ALG}^{\mathscr{D}}(X_1, \ldots, X_n) = \max(X_\tau\,,\, D_1(X_{\tau-1})\,,\, D_2(X_{\tau-2}), \ldots)$$

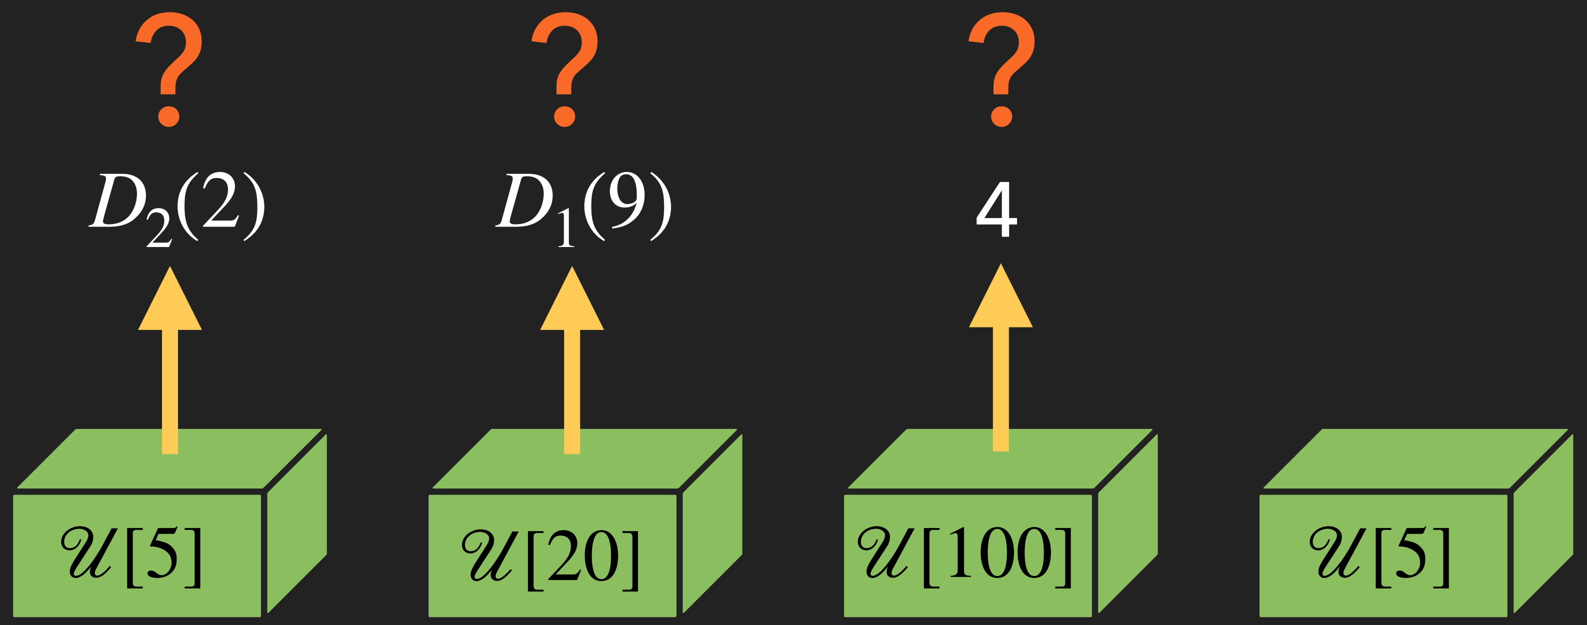

Lookback Prophet Inequality

Weakness: The model is too pessimistic,
rejections are often reversible in real-world scenarios

Definition: Decay functions $\mathscr{D} = (D_1, D_2, \ldots)$

If ALG stops at time τ, its reward is

$$\text{ALG}^{\mathscr{D}}(X_1, \ldots, X_n) = \max(X_\tau\,, D_1(X_{\tau-1})\,, D_2(X_{\tau-2}), \ldots)$$

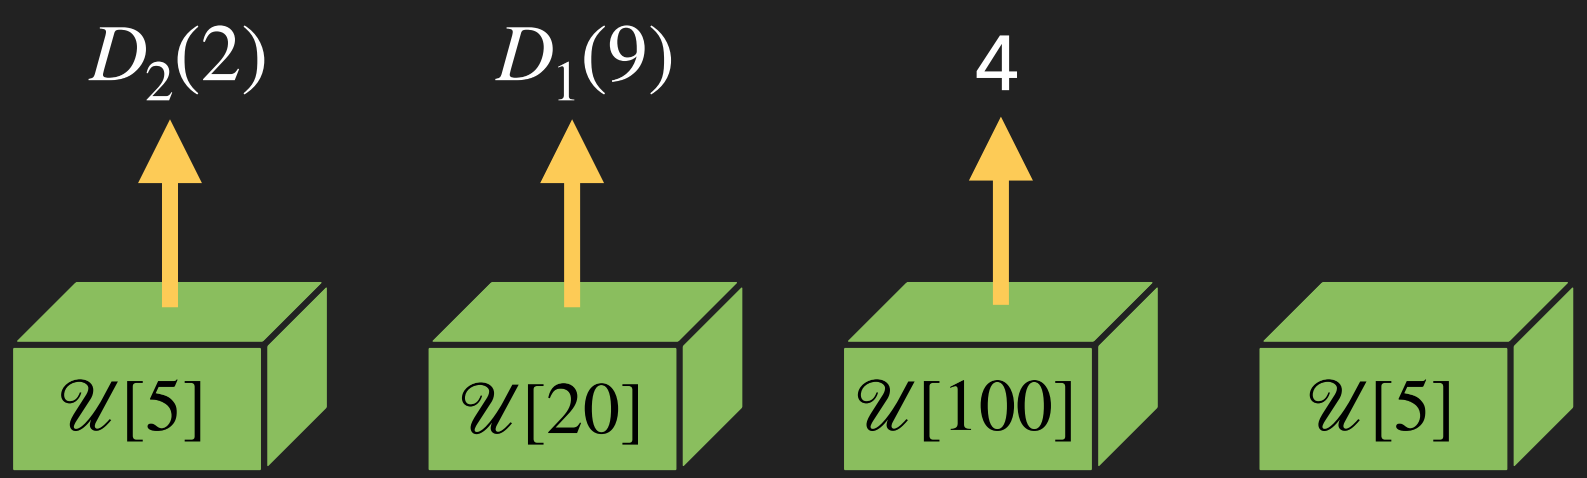

$$\text{ALG}^{\mathscr{D}}(X_1, \ldots, X_5) = \max(4\,, D_1(9)\,, D_2(2))$$

Lookback Prophet Inequality

Weakness: The model is too pessimistic,
rejections are often reversible in real-world scenarios

Definition: Decay functions $\mathscr{D} = (D_1, D_2, \ldots)$

If ALG stops at time τ, its reward is

$$\text{ALG}^{\mathscr{D}}(X_1, \ldots, X_n) = \max(X_\tau\,,\, D_1(X_{\tau-1})\,,\, D_2(X_{\tau-2}), \ldots)$$

Assumptions:

- $\forall i, x : \quad D_i(x) \in [0,x]$
- $\forall x : \quad i \mapsto D_i(x)$ is non-increasing
- $\forall i : \quad x \mapsto D_i(x)$ is non-decreasing

Lookback Prophet Inequality

Weakness: The model is too pessimistic,
rejections are often reversible in real-world scenarios

Definition: Decay functions $\mathscr{D} = (D_1, D_2, \ldots)$

If ALG stops at time τ, its reward is

$$\text{ALG}^{\mathscr{D}}(X_1, \ldots, X_n) = \max(X_\tau\,,\, D_1(X_{\tau-1})\,,\, D_2(X_{\tau-2}), \ldots)$$

Examples:

- $D_i(x) = 0$ (standard prophet inequality)
- $D_i(x) = \gamma^i x$, with $\gamma \in [0,1]$
- $D_i(x) = x - c_i$, with $c_1 \leq c_2 \leq \ldots$
- $D_i(x) \sim \mathscr{B}(p_i) \cdot x$, with $p_1 \geq p_2 \geq \ldots$ (our results extend to random functions)

Step 1: Reduction to the case of identical decay functions

$$D_1(x) \geq D_2(x) \geq \ldots \geq 0 \quad \Longrightarrow D_\infty(x) := \lim_{i \to \infty} D_i(x) \text{ exists}$$

Step 1: Reduction to the case of identical decay functions

$$D_1(x) \geq D_2(x) \geq \ldots \geq 0 \quad \Longrightarrow D_\infty(x) := \lim_{i \to \infty} D_i(x) \text{ exists}$$

Step 1: Reduction to the case of identical decay functions

$$D_1(x) \geq D_2(x) \geq \ldots \geq 0 \quad \Longrightarrow D_\infty(x) := \lim_{i\to\infty} D_i(x) \text{ exists}$$

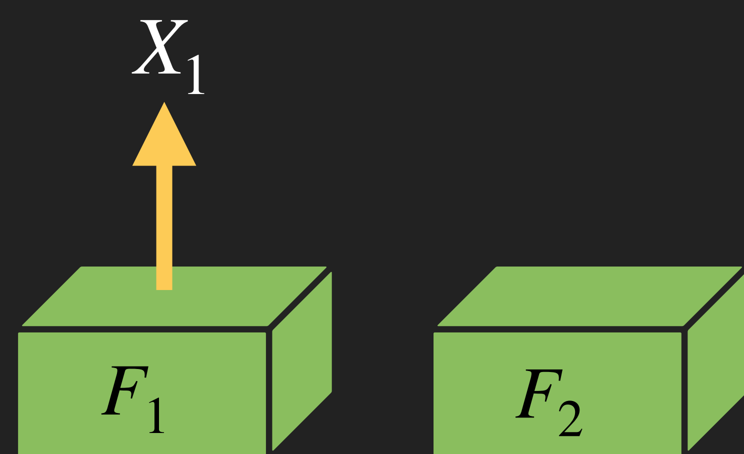

Step 1: Reduction to the case of identical decay functions

$$D_1(x) \geq D_2(x) \geq \ldots \geq 0 \quad \Longrightarrow D_\infty(x) := \lim_{i\to\infty} D_i(x) \text{ exists}$$

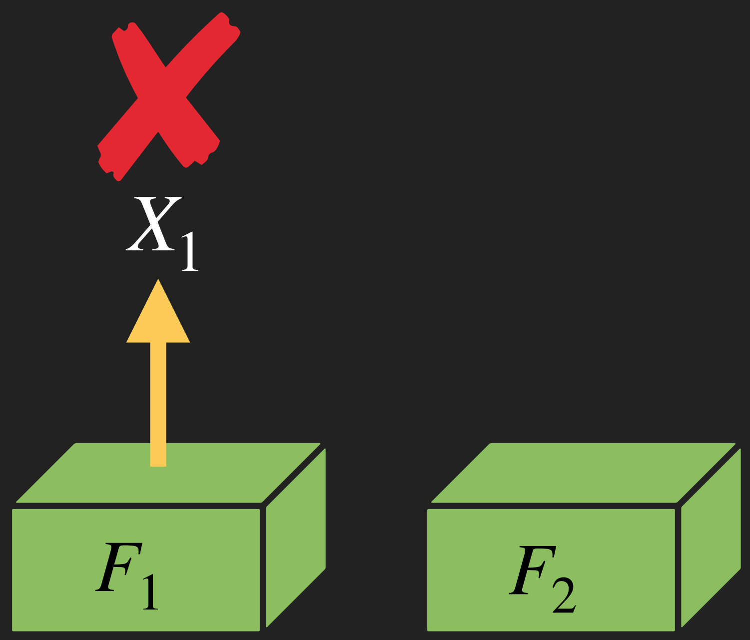

Step 1: Reduction to the case of identical decay functions

$$D_1(x) \geq D_2(x) \geq \ldots \geq 0 \quad \Longrightarrow D_\infty(x) := \lim_{i \to \infty} D_i(x) \text{ exists}$$

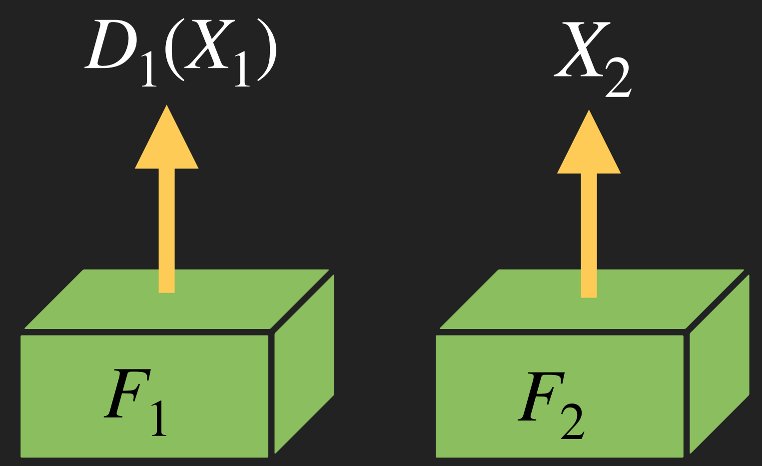

Instance $J_0 = (F_1, F_2) \quad \Longrightarrow \quad \text{ALG}^{\mathcal{D}}(J_0) = \max(X_2, D_1(X_1))$

Step 1: Reduction to the case of identical decay functions

$$D_1(x) \geq D_2(x) \geq \ldots \geq 0 \quad \Longrightarrow D_\infty(x) := \lim_{i \to \infty} D_i(x) \text{ exists}$$

Instance $J_\infty = (F_1, 0, \ldots, 0, F_2) \quad \Longrightarrow \quad \text{ALG}^{\mathcal{D}}(J_\infty) = \max(X_2, D_\infty(X_1))$

Step 1: Reduction to the case of identical decay functions

$$D_1(x) \geq D_2(x) \geq \ldots \geq 0 \quad \Longrightarrow D_\infty(x) := \lim_{i \to \infty} D_i(x) \text{ exists}$$

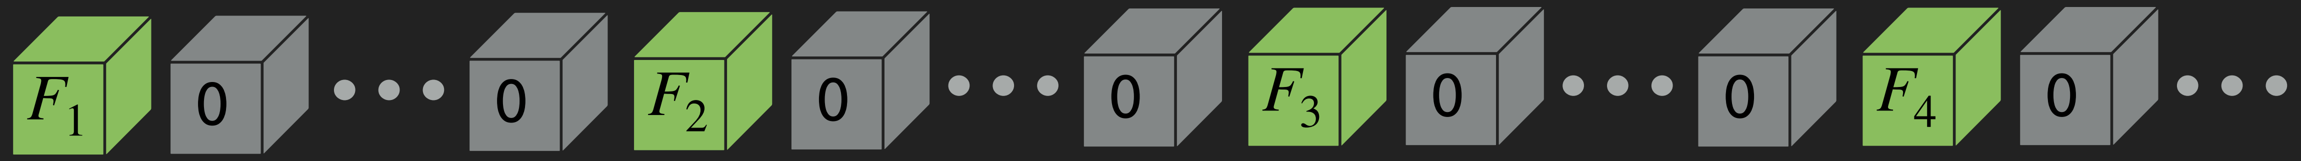

$$\text{ALG}^{\mathscr{D}}(X_1, \ldots, X_n) = \max(X_\tau, D_\infty(X_{\tau-1}), D_\infty(X_{\tau-2}), \ldots, D_\infty(X_1))$$

$$= \text{ALG}^{D_\infty}(X_1, \ldots, X_n)$$

Step 1: Reduction to the case of identical decay functions

$$D_1(x) \geq D_2(x) \geq \ldots \geq 0 \quad \Longrightarrow D_\infty(x) := \lim_{i \to \infty} D_i(x) \text{ exists}$$

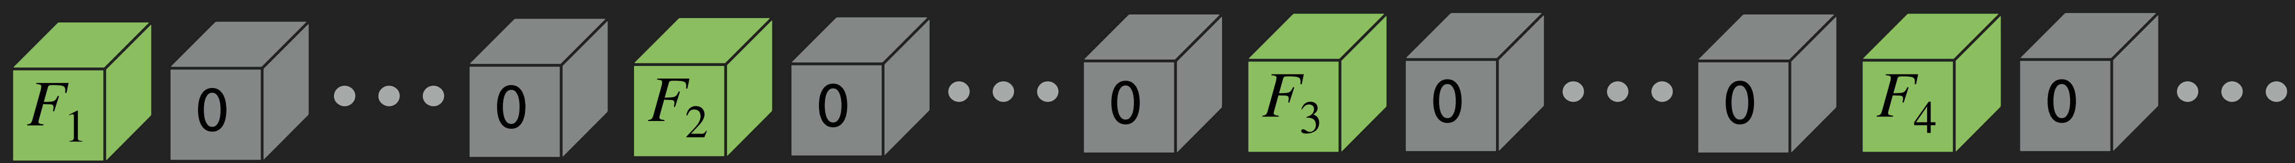

$$\text{ALG}^{\mathscr{D}}(X_1, \ldots, X_n) = \max(X_\tau, D_\infty(X_{\tau-1}), D_\infty(X_{\tau-2}), \ldots, D_\infty(X_1))$$

$$= \text{ALG}^{D_\infty}(X_1, \ldots, X_n)$$

$\Longrightarrow$ In worst-case analysis, we can assume that $D_i = D_\infty$ for all $i \geq 1$

Step 2: Reduction from $D_\infty(x)$ to γx

$$D_\infty(x) = \lim_{i\to\infty} D_i(x) \in [0,x] \quad \Longrightarrow \quad \gamma := \inf_{x>0} \frac{D_\infty(x)}{x} \in [0,1]$$

Step 2: Reduction from $D_\infty(x)$ to γx

$$D_\infty(x) = \lim_{i\to\infty} D_i(x) \in [0,x] \quad \Longrightarrow \quad \gamma := \inf_{x>0} \frac{D_\infty(x)}{x} \in [0,1]$$

Notation:

- $\mathscr{D}$-Prophet inequality: ALG$^{\mathscr{D}}(X_1, \ldots, X_n) = \max(X_\tau, D_1(X_{\tau-1}), D_2(X_{\tau-2}), \ldots)$
- D_∞-Prophet inequality: ALG$^{D_\infty}(X_1, \ldots, X_n) = \max(X_\tau, D_\infty(X_{\tau-1}), D_\infty(X_{\tau-2}), \ldots)$
- γ-Prophet inequality: ALG$^{\gamma}(X_1, \ldots, X_n) = \max(X_\tau, \gamma X_{\tau-1}, \gamma X_{\tau-2}, \ldots)$

Step 2: Reduction from $D_\infty(x)$ to γx

$$D_\infty(x) = \lim_{i\to\infty} D_i(x) \in [0,x] \quad \Longrightarrow \quad \gamma := \inf_{x>0} \frac{D_\infty(x)}{x} \in [0,1]$$

Notation:

- $\mathscr{D}$-Prophet inequality: $\text{ALG}^{\mathscr{D}}(X_1, \ldots, X_n) = \max(X_\tau, D_1(X_{\tau-1}), D_2(X_{\tau-2}), \ldots)$
- D_∞-Prophet inequality: $\text{ALG}^{D_\infty}(X_1, \ldots, X_n) = \max(X_\tau, D_\infty(X_{\tau-1}), D_\infty(X_{\tau-2}), \ldots)$
- γ-Prophet inequality: $\text{ALG}^{\gamma}(X_1, \ldots, X_n) = \max(X_\tau, \gamma X_{\tau-1}, \gamma X_{\tau-2}, \ldots)$

Theorem: Let $0 < a < b$, if $X_1, \ldots, X_n$ have support in $\{0, a, b\}$, then for any algorithm ALG

$$\text{CR}^{D_\infty}(\text{ALG}) \leq \frac{\mathbb{E}[\text{ALG}^{D_\infty}(X_1, \ldots, X_n)]}{\mathbb{E}[\max(X_1, \ldots, X_n)]} \leq \sup_{\text{A:algo}} \frac{\mathbb{E}[\text{A}^{\gamma}(X_1, \ldots, X_n)]}{\mathbb{E}[\max(X_1, \ldots, X_n)]}$$

Consequence of the reduction

Let $\gamma_{\mathscr{D}} := \inf_{x>0} \frac{\inf_i D_i(x)}{x} \in [0,1]$

Consequence of the reduction

Let $\gamma_{\mathscr{D}} := \inf_{x>0} \frac{\inf_i D_i(x)}{x} \in [0,1]$

Upper bounds in the γ-prophet inequality, proved with random variables having a support in $\{0, a, b\}$ remain true in the $\mathscr{D}$-prophet inequality

Consequence of the reduction

Let $\gamma_{\mathscr{D}} := \inf_{x>0} \frac{\inf_i D_i(x)}{x} \in [0,1]$

Upper bounds in the γ-prophet inequality, proved with random variables having a support in $\{0, a, b\}$ remain true in the $\mathscr{D}$-prophet inequality

Lower bounds in the γ-prophet inequality remain true in the $\mathscr{D}$-prophet inequality.

$$(D_i(x) \geq D_\infty(x) \geq \gamma x)$$

Main results

Let $\gamma_{\mathscr{D}} := \inf_{x>0} \frac{\inf_i D_i(x)}{x} \in [0,1]$

Upper bound: The competitive ratio of any algorithm in the $\mathscr{D}$-prophet inequality is at most $\frac{1}{2-\gamma}$.

Lower bound: Let ALG the algorithm that select the first value exceeding θ, with $\mathbb{P}(\max(X_1, \ldots, X_n) \geq \theta) = \frac{1}{2-\gamma}$, then

$$\text{CR(ALG)} = \frac{1}{2-\gamma} \ .$$

Main results

Let $\gamma_{\mathscr{D}} := \inf_{x>0} \frac{\inf_i D_i(x)}{x} \in [0,1]$

Upper bound: The competitive ratio of any algorithm in the $\mathscr{D}$-prophet inequality is at most $\frac{1}{2-\gamma}$.

Lower bound: Let ALG the algorithm that select the first value exceeding θ, with $\mathbb{P}(\max(X_1, \ldots, X_n) \geq \theta) = \frac{1}{2-\gamma}$, then

$$\text{CR(ALG)} = \frac{1}{2-\gamma} \ .$$

Problem solved!

Variants of the prophet inequality

Random order: (prophet secretary) The distributions $F_1, \ldots, F_n$ are adversarial, but the values are observed in a uniformly random order

Variants of the prophet inequality

Random order: (prophet secretary) The distributions $F_1, \ldots, F_n$ are adversarial, but the values are observed in a uniformly random order

IID model: The distributions are identical $F_1 = \ldots = F_n = F$

Variants of the prophet inequality

Random order: (prophet secretary) The distributions $F_1, \ldots, F_n$ are adversarial, but the values are observed in a uniformly random order

IID model: The distributions are identical $F_1 = \ldots = F_n = F$

Main results

- The same reductions of the decay functions to $D_\infty(x)$ then γx remain true in both models (but more technical)
- Upper bounds: depending on γ
- Lower bounds: single threshold algorithms

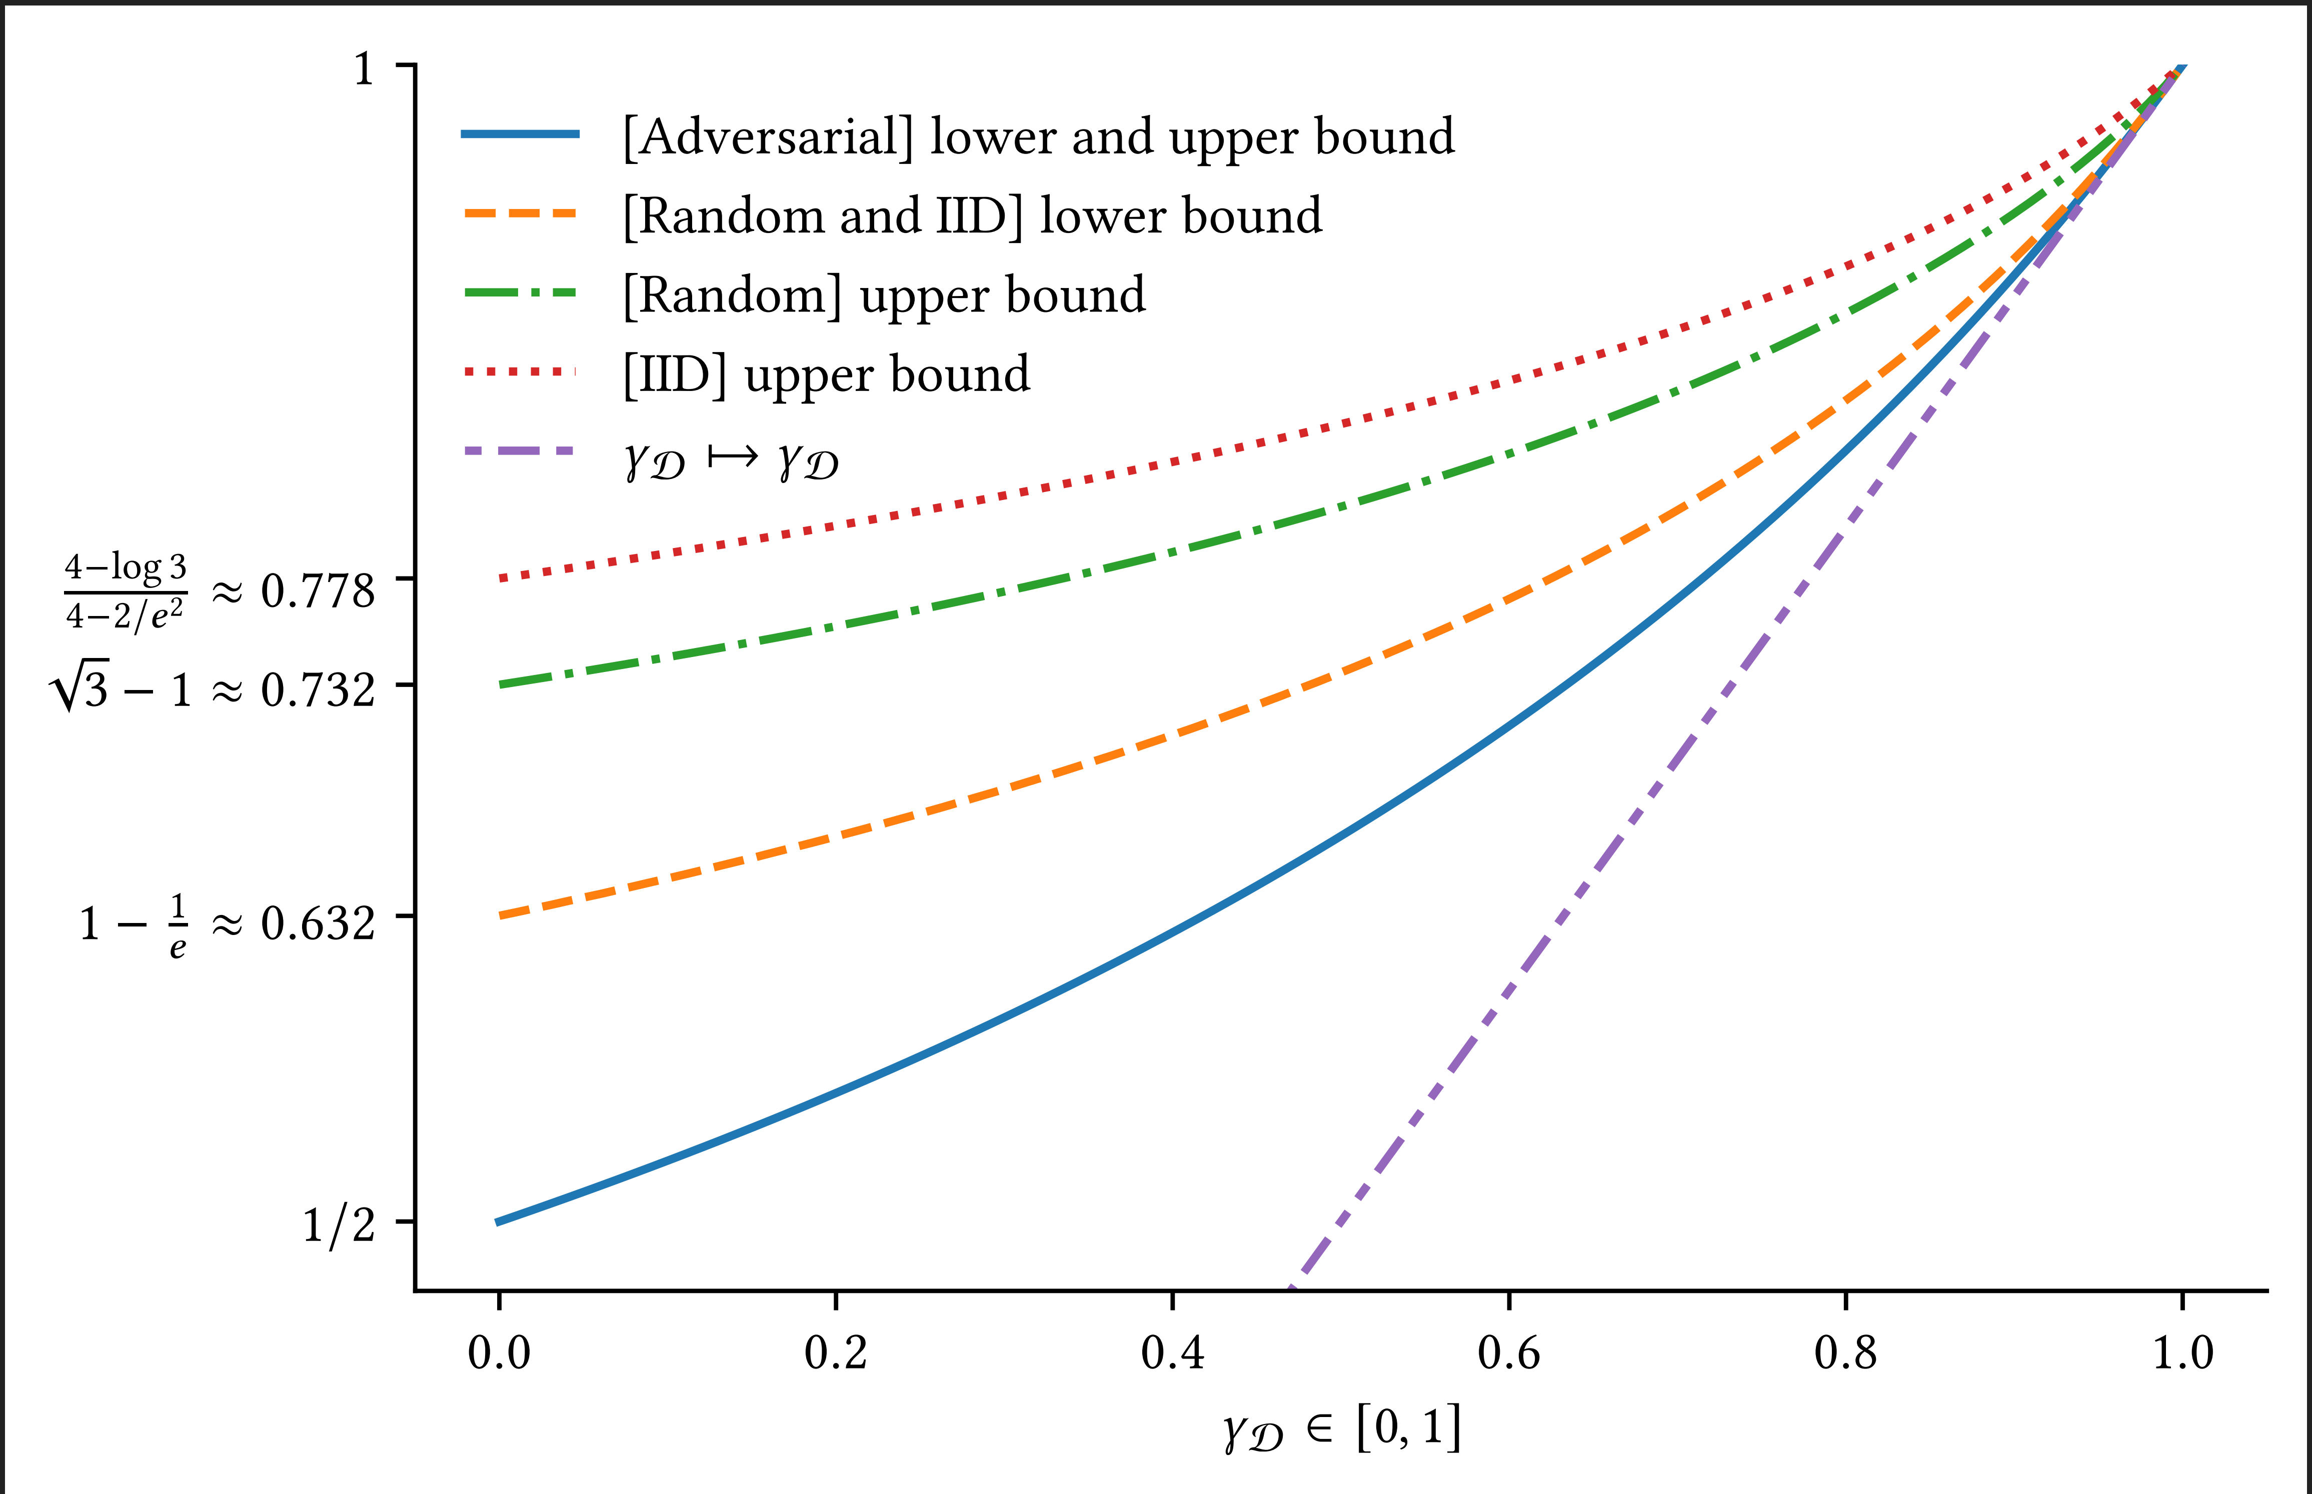

[Adversarial] lower and upper bound
[Random and IID] lower bound
[Random] upper bound
[IID] upper bound
$\gamma_{\mathcal{D}} \mapsto \gamma_{\mathcal{D}}$
1
$\frac{4-\log 3}{4-2/e^2} \approx 0.778$
$\sqrt{3}-1 \approx 0.732$
$1-\frac{1}{e} \approx 0.632$
1/2
0.0
0.2
0.4
0.6
0.8
1.0
$\gamma_{\mathcal{D}} \in [0, 1]$

Future work: in the random order and IID models

- Improve the upper bounds
- Analyse more general algorithms