Intrinsic Robustness of Prophet Inequality to Strategic Reward Signaling

Wei Tang[1], Haifeng Xu[2], Ruimin Zhang[2], Derek Zhu[2]

[1]Chinese University of Hong Kong

[2]The University of Chicago

香港中文大學
The Chinese University of Hong Kong

THE UNIVERSITY OF CHICAGO

Prophet Inequality

A gambler's Problem

N boxes containing unknown rewards $X_i \sim H_i$

H_1 H_2 H_3 $\cdots$ H_N

Prophet Inequality

A gambler's Problem

N boxes containing unknown rewards $X_i \sim H_i$

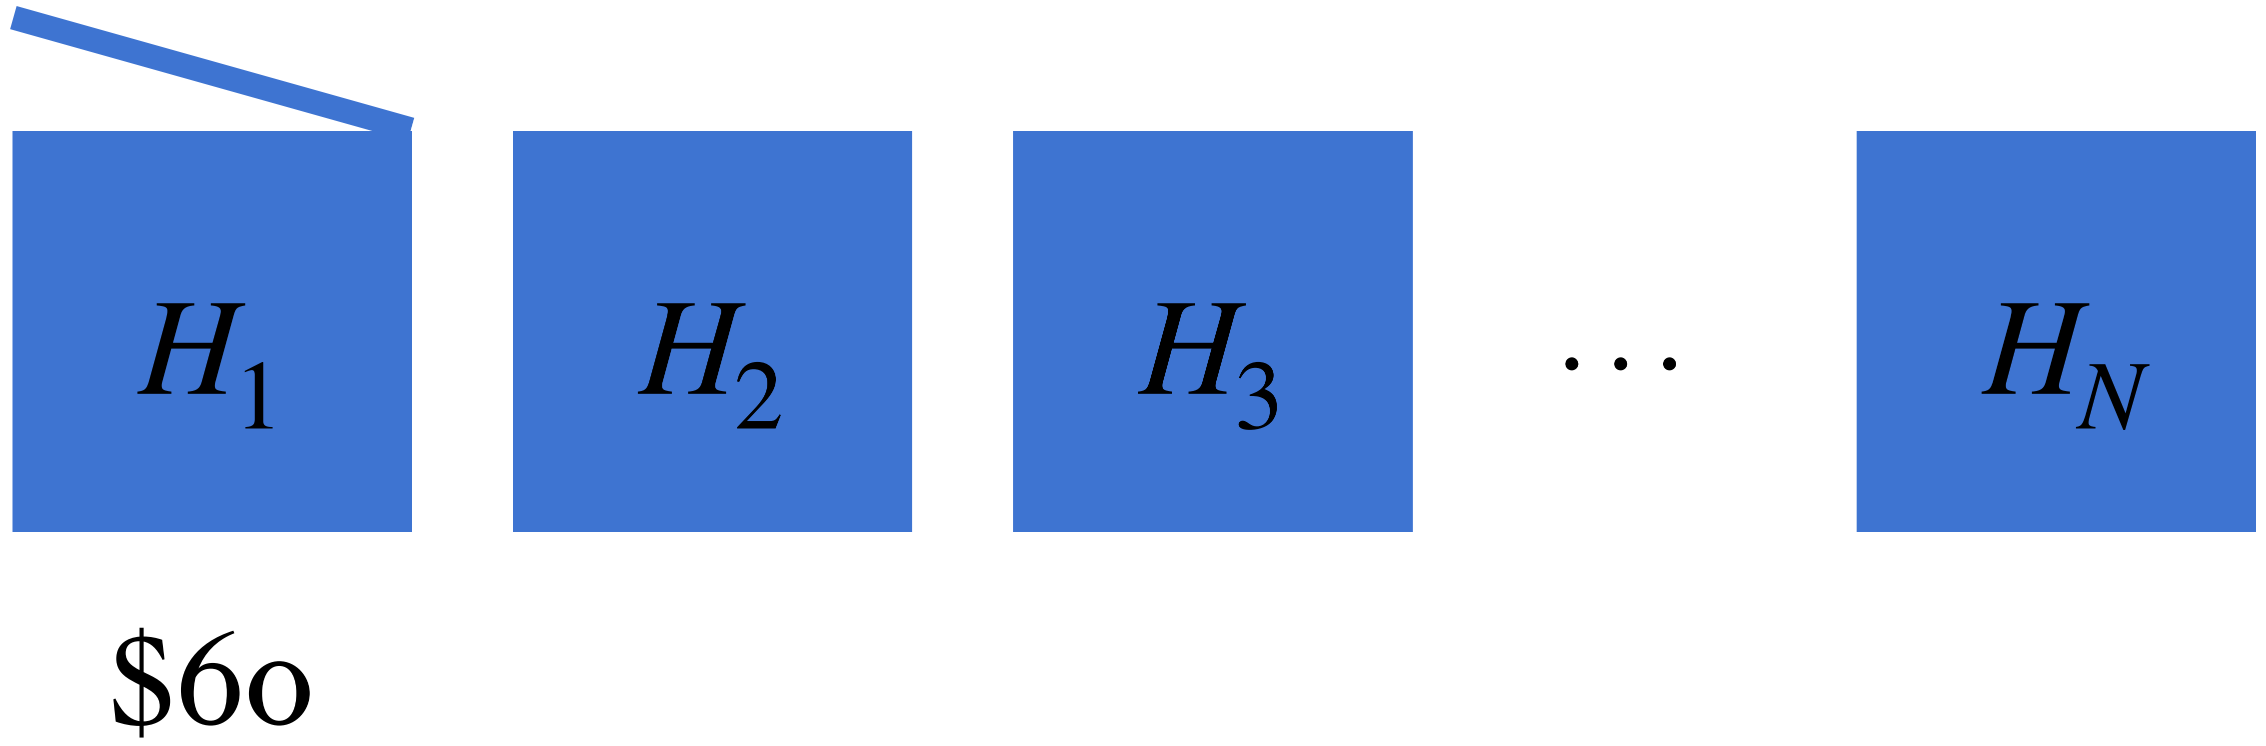

Prophet Inequality

A gambler's Problem

N boxes containing unknown rewards $X_i \sim H_i$

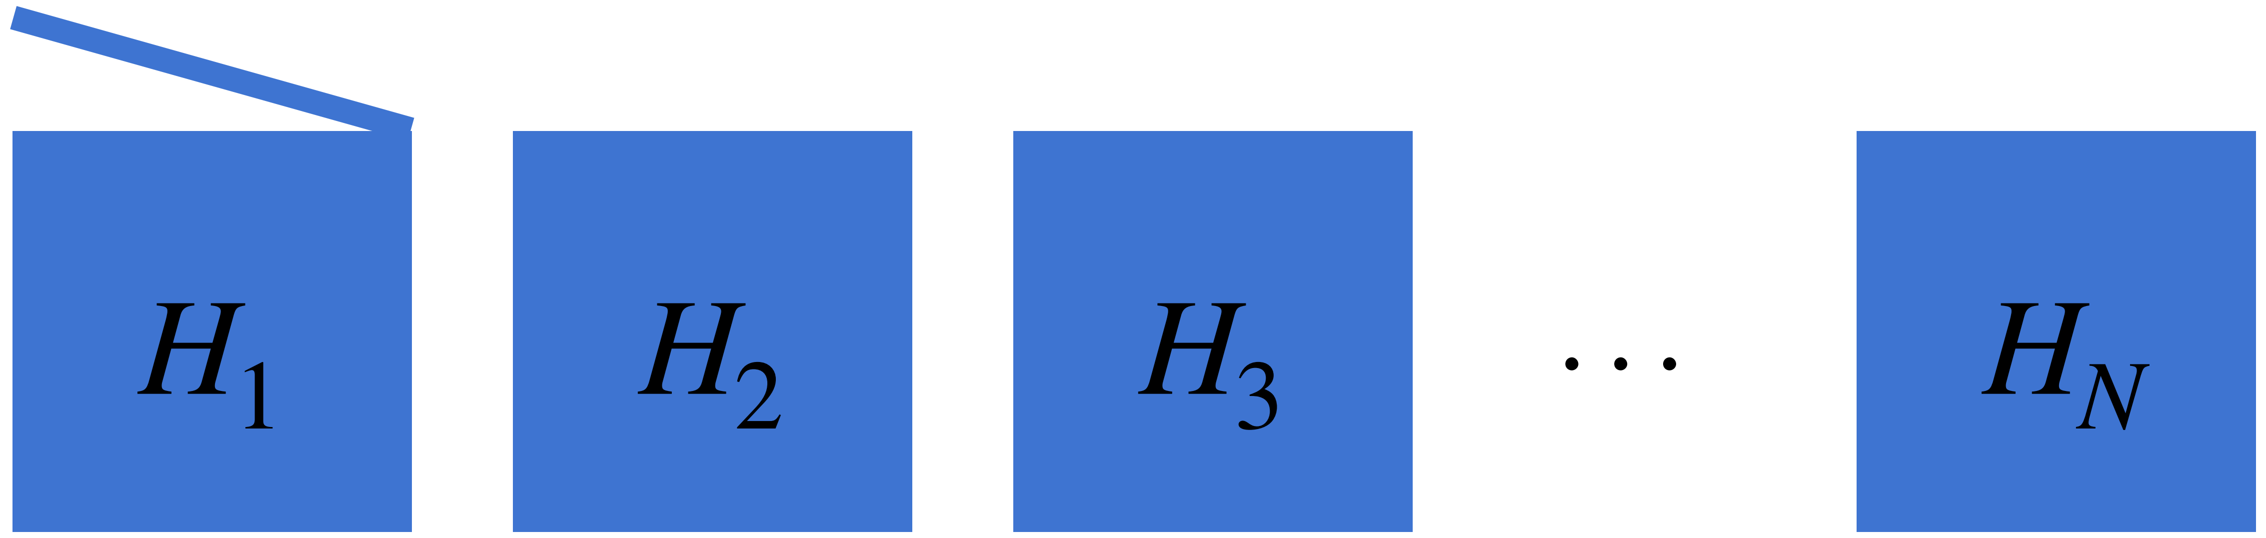

$60

Stop: win $60, and the game ends

Continue: discard $60, game continues to next box

Prophet Inequality

A gambler's Problem

N boxes containing unknown rewards $X_i \sim H_i$

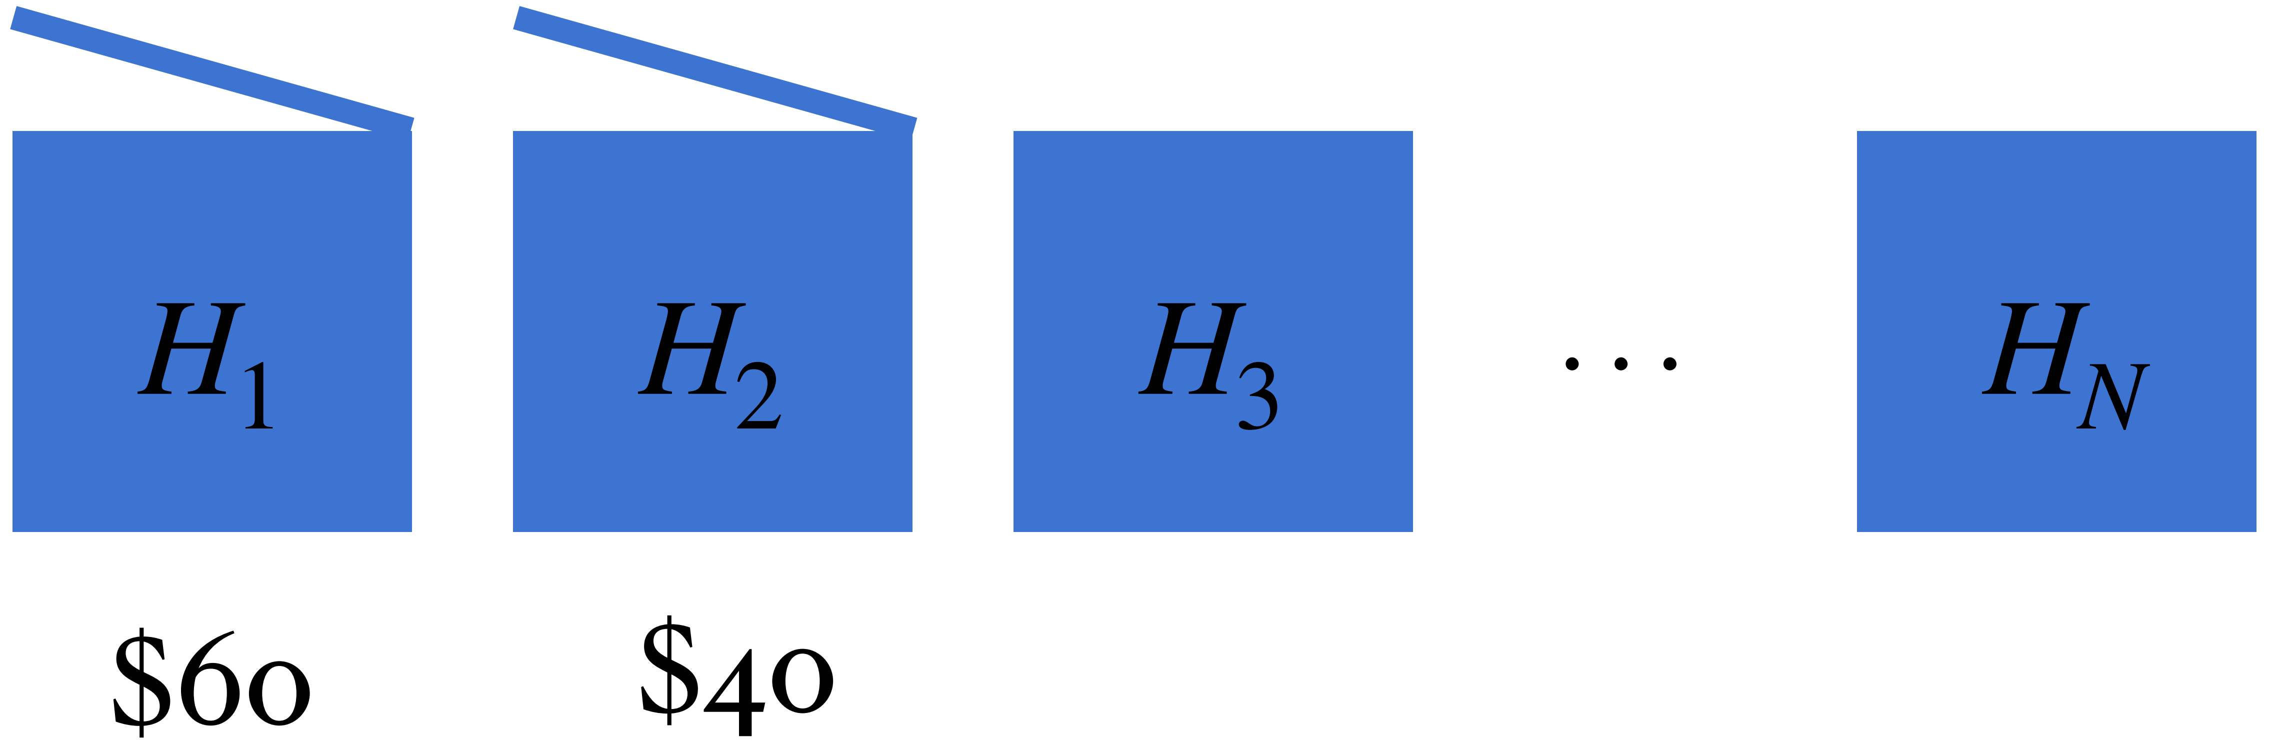

Stop: win \$40, and the game ends

Continue: discard \$40, game continues to next box

Prophet Inequality

A gambler's Problem

N boxes containing unknown rewards $X_i \sim H_i$

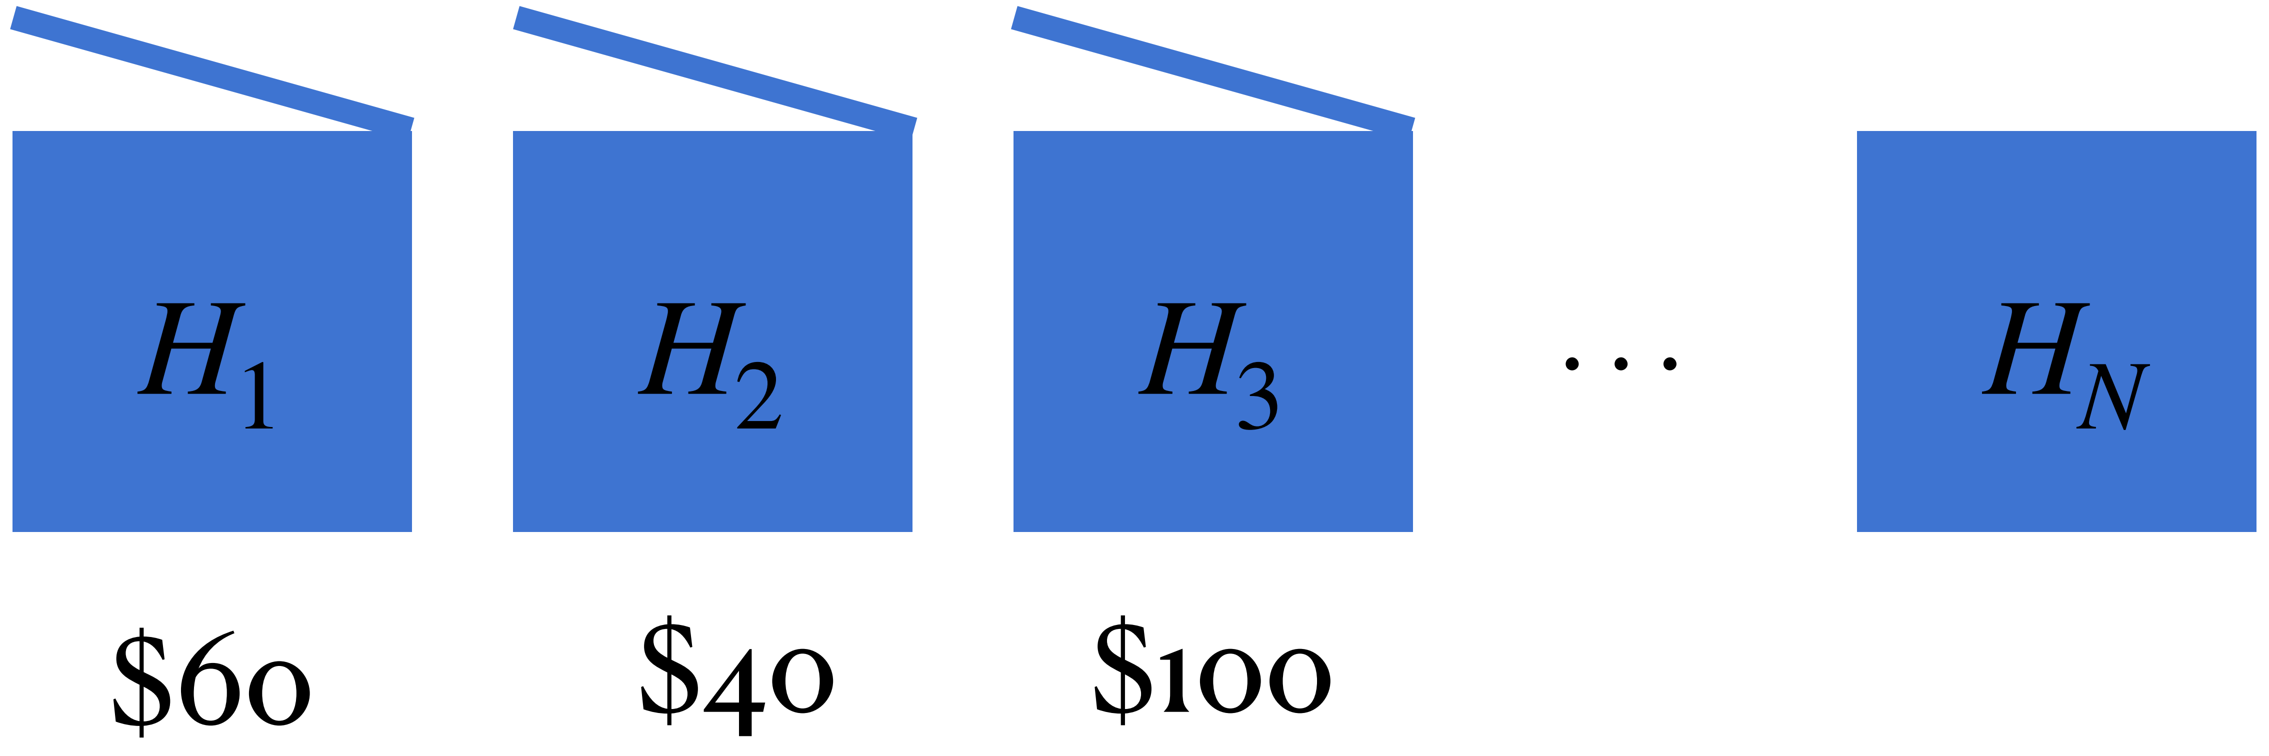

Stop: win $100, and the game ends

Continue: discard $100, game continues to next box

Prophet Inequality

A gambler's Problem

N boxes containing unknown rewards $X_i \sim H_i$

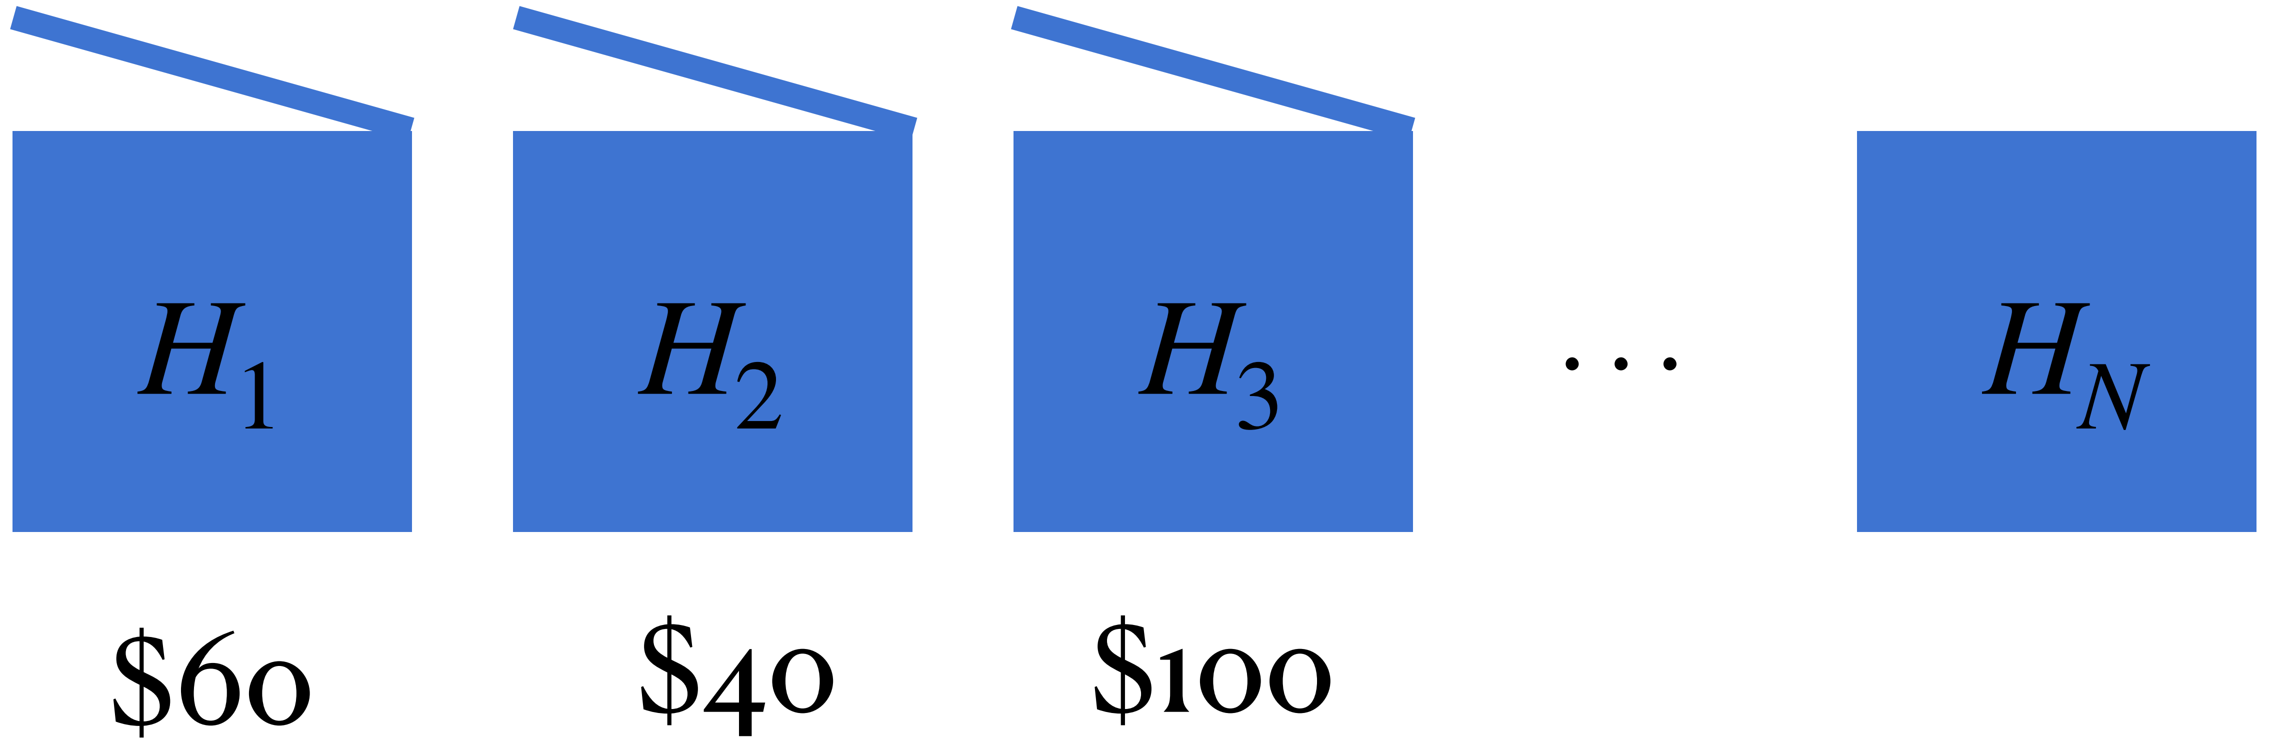

Stop: win $100, and the game ends

Continue: discard $100, game continues to next box

Goal: win as much reward as possible

Prophet Inequality

- Given N independent distributions: $H_1, H_2, \ldots, H_N$
- At step i, $X_i \sim H_i$ is revealed
- The searcher make irrevocable accept/reject decision for each X_i
- Goal: maximize the accepted value in expectation
- Benchmark: prophet's payoff $\text{OPT} = \mathbb{E}\{\max_i X_i\}$

$$\text{Competitive Ratio} = \max_I \frac{\mathbb{E}_{\text{randomness in I}}[OPT(I)]}{\mathbb{E}_{\text{randomness in I, ALG}}[ALG(I)]}$$

Theorem [Krengel, Sucheston, Garling '77]: There exist a strategy for the searcher such that

$$\mathbb{E}\{\text{reward}\} \geq \frac{1}{2}\mathbb{E}\{\max_i X_i\}$$

Theorem [Samuel-Cahn '84], [Kleinberg Weinberg 12]: There exist fixed threshold policies for the searcher such that

$$\mathbb{E}\{\text{reward}\} \geq \frac{1}{2}\mathbb{E}\{\max_i X_i\}$$

Tightness of 1/2:

$X_1 = 1$	w.p.	1
$X_2 = 0$	w.p.	$1 - p$
$X_2 = 1/p$	w.p.	p

- Find threshold t such that $\Pr(\exists i \text{ with } X_i \geq t) = 1/2$
- Pick the first element that exceeds t

- Alternatively: $t = 1/2 \cdot \text{OPT}$
- Also: any t between these two

Theorem [Krengel, Sucheston, Garling '77]: There exist a strategy for the searcher such that

$$\mathbb{E}\{\text{reward}\} \geq \frac{1}{2}\mathbb{E}\{\max_i X_i\}$$

Tightness of 1/2:

$$X_1 = 1 \quad \text{w.p.} \quad 1$$
$$X_2 = 0 \quad \text{w.p.} \quad 1 - p$$
$$X_2 = 1/p \quad \text{w.p.} \quad p$$

Theorem [Samuel-Cahn '84], [Kleinberg Weinberg 12]: There exist fixed threshold policies for the searcher such that

$$\mathbb{E}\{\text{reward}\} \geq \frac{1}{2}\mathbb{E}\{\max_i X_i\}$$

Robust?

- Find threshold t such that $\Pr(\exists i \text{ with } X_i \geq t) = 1/2$
- Pick the first element that exceeds t

- Alternatively: $t = 1/2 \cdot \text{OPT}$
- Also: any t between these two

Prophet Inequality with Strategic Reward Signaling

- Given N independent distributions: $H_1, H_2, \ldots, H_N$
  - Each H_i is associated with a strategic player i
- At step i, player i strategically disclose information about $X_i \sim H_i$
- The searcher make irrevocable accept/reject decision for each X_i
- Goal: maximize the accepted value in expectation
- Benchmark: prophet's payoff $\mathrm{OPT} = \mathbb{E}\{\max_i X_i\}$

Robustness

Definition 1 (α- robust stopping policy): A stopping policy p is α-robust if

1. It achieves α-approximation to OPT when players are strategically signaling their rewards
2. It remains a $1/2$-approximation in the standard non-strategic setting.

$$\mathrm{OPT} = \mathbb{E}[\max_i X_i] \text{ where } X_i \sim H_i$$

Player's Optimal Signaling Scheme

Proposition 1: Given a threshold stopping policy with threshold T, for each player i:

- If $T \leq \lambda_i$, then player i's optimal information revealing strategy is the no information strategy;
- If $T > \lambda_i$, then player i's optimal information revealing strategy is threshold signaling and determined by a cutoff t_i that satisfies

$$T = \mathbb{E}[X_i \,|\, X_i \geq t_i] = \int_{t_i}^{\infty} x dH_i(x) / (1 - H_i(t_i))$$

That is, player i's optimal signaling scheme sends one of two signals:

$$X_i \geq t_i \text{ or } X_i < t_i$$

First Main Result

Theorem 1: For any distributions $H_1, H_2, \ldots, H_N$, a threshold stopping policy with threshold $T = 1/2 \cdot OPT$ is $\frac{1 - 1/e}{2}$ - robust, and this is tight among a class of thresholds.

Second Main Result

achieving 1/2-robustness for special distributions

Theorem 2: For IID distributions $H_1 = H_2 = \cdots = H_N$, a threshold stopping policy with threshold

$$T = \sum_i \mathbb{E}_{H_i}(X_i - T)^+ \text{ is } 1/2\text{-robust, and this is tight.}$$

Theorem 3: If $H_1, H_2, \ldots, H_N$ satisfy certain regularity assumptions, then a threshold stopping policy with threshold T that satisfies $2 \cdot T_{KW} \leq T \leq T_{SC}$ is $1/2$ - robust

Thank you!