Prophet Inequalities via the Expected Competitive Ratio

Alexandros Tsigonias-Dimitriadis, Universidad de Chile

Joint work with Tomer Ezra, Stefano Leonardi, Rebecca Reiffenhäuser (Sapienza University of Rome), and Matteo Russo (Georgia Tech)

UNIVERSIDAD DE CHILE

Prophet Inequality

Prophet Inequality

$U[4,7]$

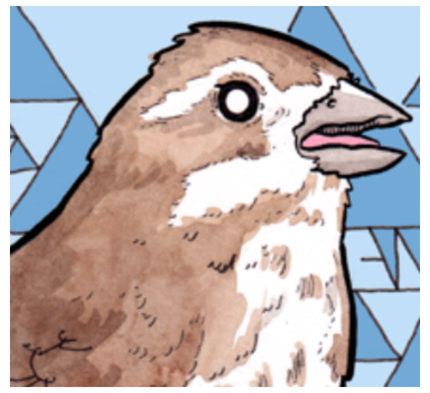

$U[2,9]$

$U[6,8]$

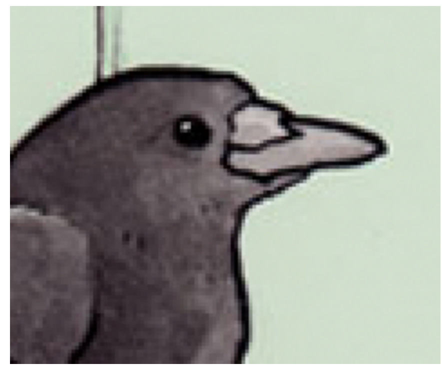

$U[2,4]$

Prophet Inequality

$U[4,7]$

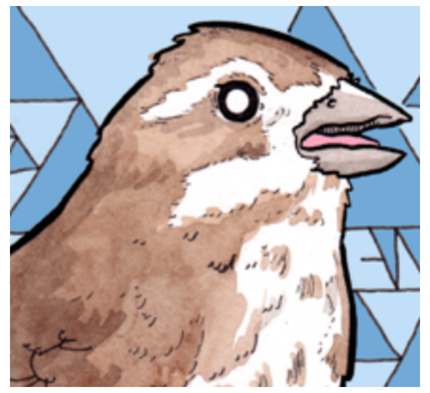

$U[2,9]$

$U[6,8]$

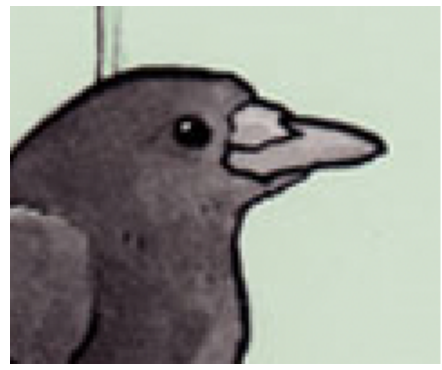

3.2

Prophet Inequality

$U[4,7]$

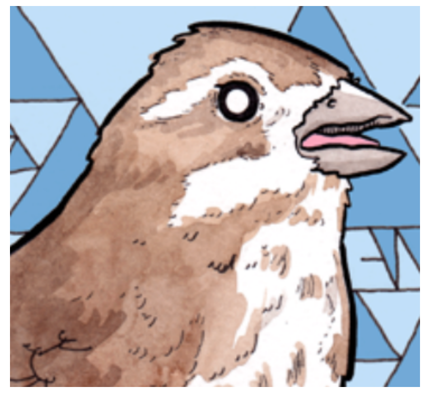

$U[2,9]$

$U[6,8]$

Prophet Inequality

$U[4,7]$

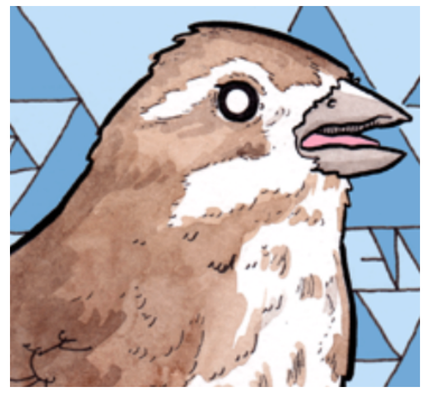

$U[2,9]$

6.3

Prophet Inequality

5.8

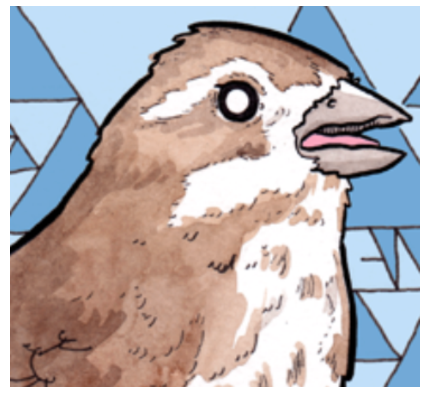

7

6.3

Prophet Inequality

5.8

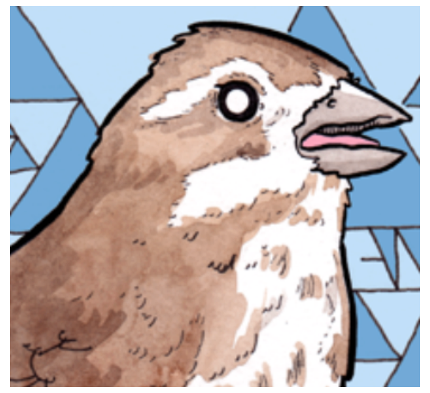

7

6.3

Thm [Krengel, Sucheston, Garling '77]: There exists a strategy for which $\mathbb{E}[\text{Seller}] \geq 1/2 \cdot \mathbb{E}[\text{Prophet}]$.

Prophet Inequality

5.8

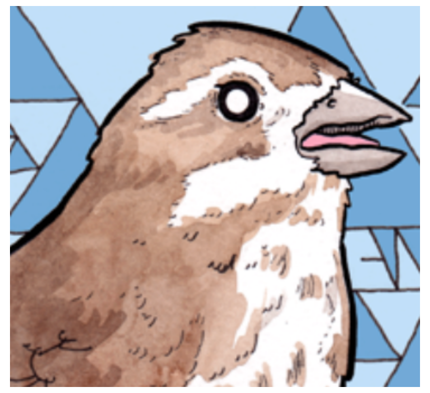

7

6.3

Thm [Krengel, Sucheston, Garling '77]: There exists a strategy for which $\mathbb{E}[\text{Seller}] \geq 1/2 \cdot \mathbb{E}[\text{Prophet}]$.

We cannot do better than $1/2$!

Prophet Inequality

5.8

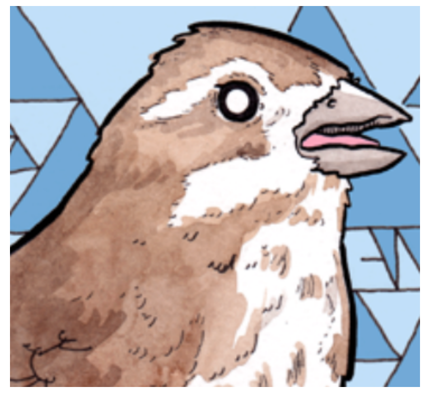

7

6.3

Thm [Krengel, Sucheston, Garling '77]: There exists a strategy for which $\mathbb{E}[\text{Seller}] \geq 1/2 \cdot \mathbb{E}[\text{Prophet}]$.

We cannot do better than $1/2$!

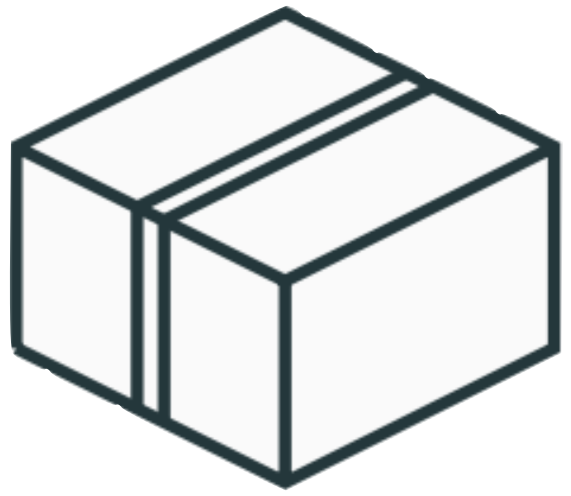

$w_1 = 1$

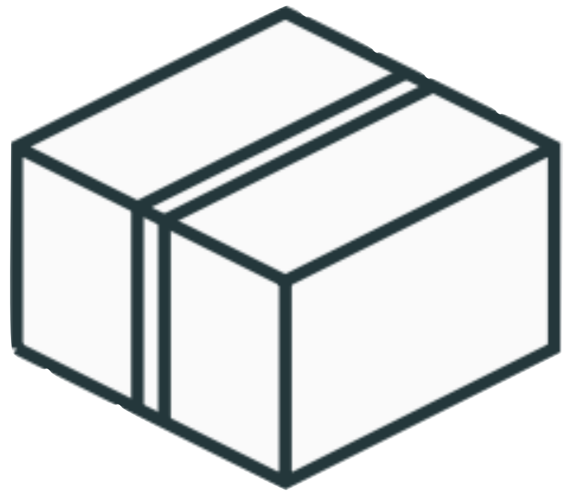

$$w_2 = \begin{cases} \varepsilon, \text{ w.p. } 1 - \varepsilon \\ 1/\varepsilon, \text{ w.p. } \varepsilon \end{cases}$$

Prophet Inequality

5.8

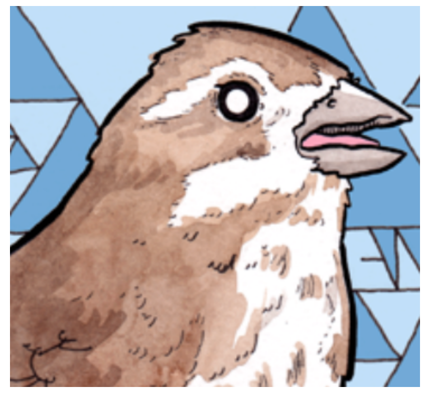

7

6.3

Thm [Krengel, Sucheston, Garling '77]: There exists a strategy for which $\mathbb{E}[\text{Seller}] \geq 1/2 \cdot \mathbb{E}[\text{Prophet}]$.

We cannot do better than $1/2$!

Any ALG gets ≤ 1.

The prophet gets $\mathbb{E}\left[\max\{w_1, w_2\}\right] \approx 2$.

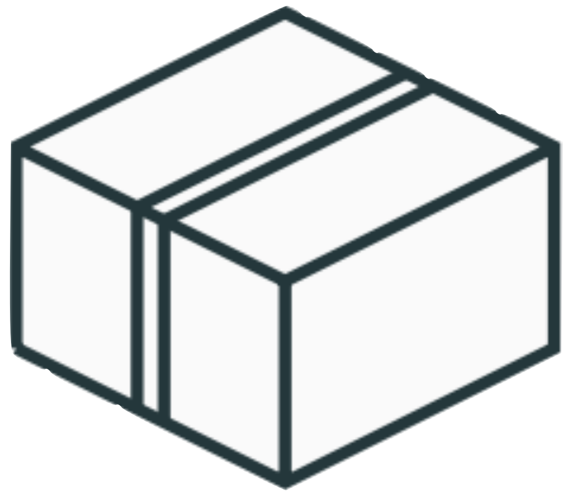

$w_1 = 1$

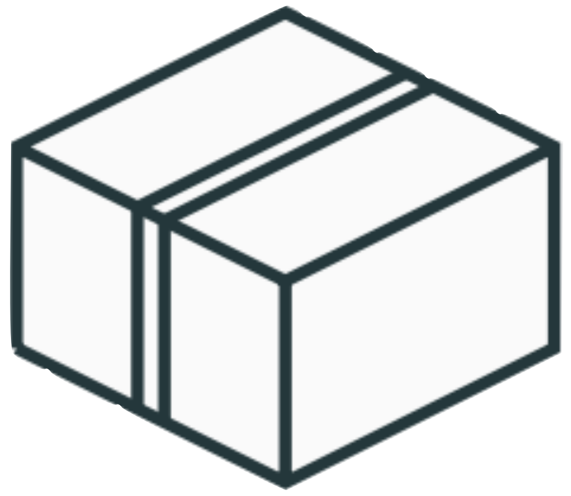

$$w_2 = \begin{cases} \varepsilon, \text{ w.p. } 1 - \varepsilon \\ 1/\varepsilon, \text{ w.p. } \varepsilon \end{cases}$$

Prophet Inequality

5.8

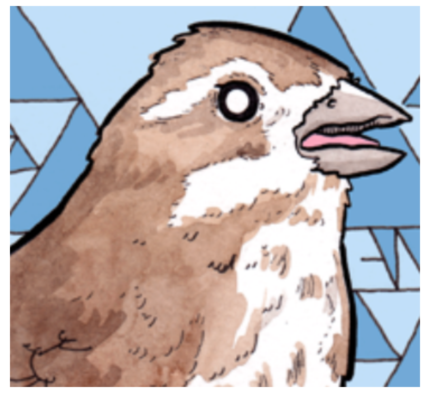

7

6.3

Thm [Krengel, Sucheston, Garling '77]: There exists a strategy for which $\mathbb{E}[\text{Seller}] \geq 1/2 \cdot \mathbb{E}[\text{Prophet}]$.

We cannot do better than $1/2$!

Any ALG gets ≤ 1.

The prophet gets $\mathbb{E}\left[\max\{w_1, w_2\}\right] \approx 2$.

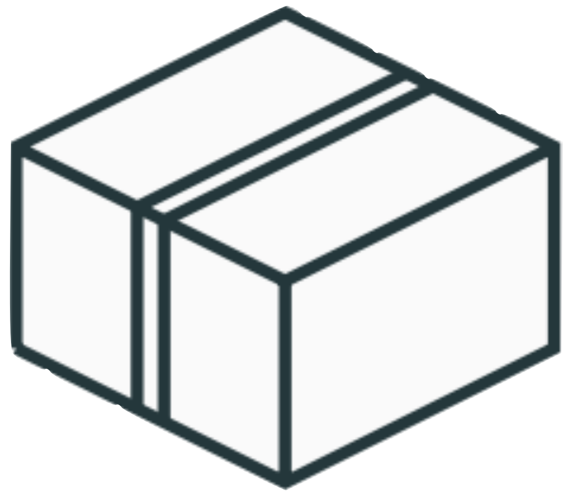

$w_1 = 1$

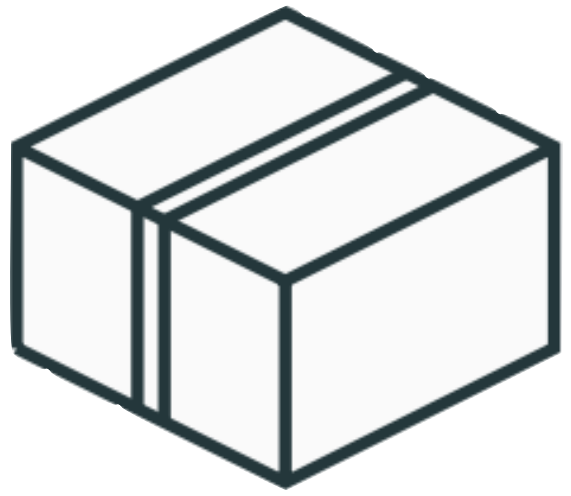

$$w_2 = \begin{cases} \varepsilon, \text{ w.p. } 1-\varepsilon \\ 1/\varepsilon, \text{ w.p. } \varepsilon \end{cases}$$

Thm In fact, it is a fixed threshold strategy! [Samuel-Cahn '84; Kleinberg, Weinberg '12]

Bayesian online selection problems

Bayesian online selection problems

- Elements arrive one by one. For each elem j, $w_j \sim D_j$ is revealed upon arrival.

Bayesian online selection problems

- Elements arrive one by one. For each elem j, $w_j \sim D_j$ is revealed upon arrival.
- ALG (or decision-maker) decides to pick each elem immediately and irrevocably.

Bayesian online selection problems

- Elements arrive one by one. For each elem j, $w_j \sim D_j$ is revealed upon arrival.
- ALG (or decision-maker) decides to pick each elem immediately and irrevocably.
- **Goal**: Pick set of elms that max a weight function s.t. (downward-closed) constraints $\mathbf{F}$.

Bayesian online selection problems

- Elements arrive one by one. For each elem j, $w_j \sim D_j$ is revealed upon arrival.

- ALG (or decision-maker) decides to pick each elem immediately and irrevocably.

- **Goal**: Pick set of elms that max a weight function s.t. (downward-closed) constraints $\mathbf{F}$.

  Whenever $\mathcal{S} \in \mathbf{F}$, then $\mathcal{Q} \in \mathbf{F}$ if $\mathcal{Q} \subseteq \mathcal{S}$.

Bayesian online selection problems

- Elements arrive one by one. For each elem j, $w_j \sim D_j$ is revealed upon arrival.
- ALG (or decision-maker) decides to pick each elem immediately and irrevocably.
- **Goal**: Pick set of elms that max a weight function s.t. (downward-closed) constraints $\mathbf{F}$.

Whenever $\mathcal{S} \in \mathbf{F}$, then $\mathcal{Q} \in \mathbf{F}$ if $\mathcal{Q} \subseteq \mathcal{S}$.

Matching

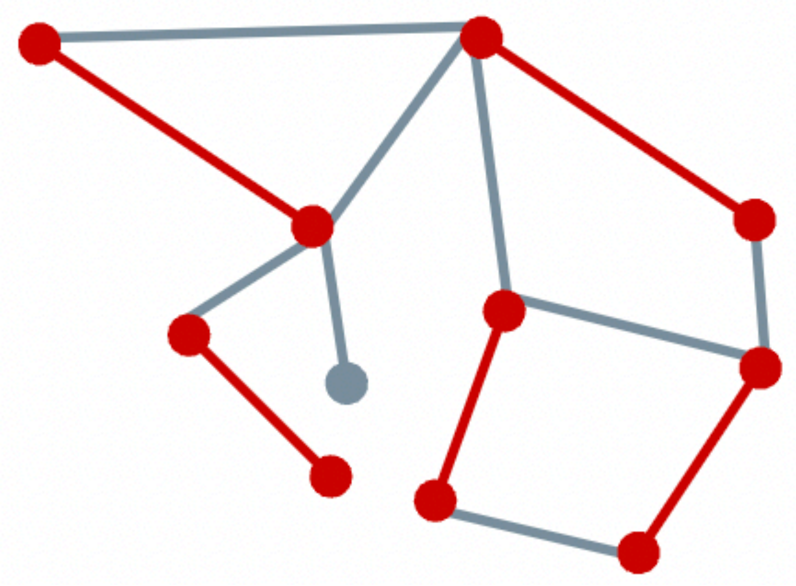

Knapsack

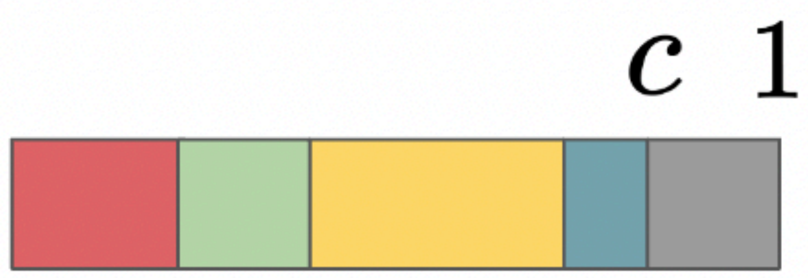

Matroid

$$\mathcal{M} = (E, \mathcal{I})$$

Prophet inequalities literature

Prophet inequalities literature

A very active area of research with lots of open questions! Some well-studied directions are:

Prophet inequalities literature

A very active area of research with lots of open questions! Some well-studied directions are:

- Arrival order of the elements
  [Hill, Kertz '82], [Yan '11], [Ehsani, Hajiaghayi, Kesselheim, Singla '18], [Correa, Saona, Ziliotto '21], [Correa, Foncea, Hoeksma, Oosterwijk, Vredeveld '21]

- Combinatorial settings
  [Alaei '11], [Kleinberg, Weinberg '12], [Gravin, Feldman, Lucier '15], [Dütting, Feldman, Kesselheim, Lucier '17], [Rubinstein, Singla '17], [Ezra, Feldman, Gravin, Tang '20], [Feldman, Svensson, Zenklusen '21], [Jiang, Ma, Zhang '22]

- Samples from unknown distributions
  [Azar, Kleinberg, Weinberg '14], [Correa, Dütting, Fischer, Schewior '19], [Rubinstein, Wang, Weinberg '20], [Correa, Cristi, Epstein, Soto '20], [Kaplan, Naori, Raz '20], [Caramanis, Dütting, Faw, Fusco, Lazos, Leonardi, Papadigenopoulos, Pountourakis, Reiffenhäuser '22]

- Connections to posted price mechanisms
  [Hajiaghayi, Kleinberg, Sandholm '07], [Chawla, Hartline, Malec, Sivan '10], , [Dütting, Feldman, Kesselheim, Lucier '17], [Correa, Pizarro, Verdugo '19]

Performance measures

Performance measures

Traditionally in PIs we use the **ratio of expectations** $\mathrm{RoE} := \dfrac{\mathbb{E}[\mathrm{ALG}]}{\mathbb{E}[\mathrm{OPT}]}$.

Performance measures

Traditionally in PIs we use the **ratio of expectations** $\mathrm{RoE} := \dfrac{\mathbb{E}[\mathrm{ALG}]}{\mathbb{E}[\mathrm{OPT}]}$.

For single choice, the **probability of selecting the max** $\mathrm{PbM} := \Pr[\mathrm{ALG} = \mathrm{OPT}]$ has been studied (e.g., [Esfandiari, HajiAghayi, Mitzenmacher, Lucier '20], [Nuti '22]).

Performance measures

Traditionally in PIs we use the **ratio of expectations** $\mathrm{RoE} := \dfrac{\mathbb{E}[\mathrm{ALG}]}{\mathbb{E}[\mathrm{OPT}]}$.

For single choice, the **probability of selecting the max** $\mathrm{PbM} := \Pr[\mathrm{ALG} = \mathrm{OPT}]$ has been studied (e.g., [Esfandiari, HajiAghayi, Mitzenmacher, Lucier '20], [Nuti '22]).

Recently, an alternative benchmark has been the *optimal online algorithm* (e.g., [Anari, Niazadeh, Saberi, Shameli '19], [Papadimitrou, Pollner, Saberi, Wajc '21]).

Performance measures

Traditionally in PIs we use the **ratio of expectations** $\mathsf{RoE} := \dfrac{\mathbb{E}[\mathsf{ALG}]}{\mathbb{E}[\mathsf{OPT}]}$.

For single choice, the **probability of selecting the max** $\mathsf{PbM} := \Pr[\mathsf{ALG} = \mathsf{OPT}]$ has been studied (e.g., [Esfandiari, HajiAghayi, Mitzenmacher, Lucier '20], [Nuti '22]).

Recently, an alternative benchmark has been the *optimal online algorithm* (e.g., [Anari, Niazadeh, Saberi, Shameli '19], [Papadimitrou, Pollner, Saberi, Wajc '21]).

This work: We initialize the study of the **expected ratio** $\mathsf{EoR} := \mathbb{E}\left[\dfrac{\mathsf{ALG}}{\mathsf{OPT}}\right]$.

Motivation

Motivation

1. RoE can perform poorly wrt the ex-post outcome, does not capture **risk aversion** (i.e., avoid the possibility of extremely bad outcomes).

Motivation

1. RoE can perform poorly wrt the ex-post outcome, does not capture **risk aversion** (i.e., avoid the possibility of extremely bad outcomes).

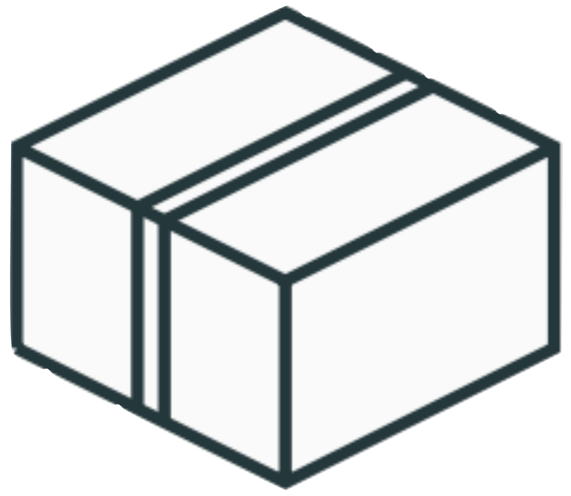

$w_1 = 1$

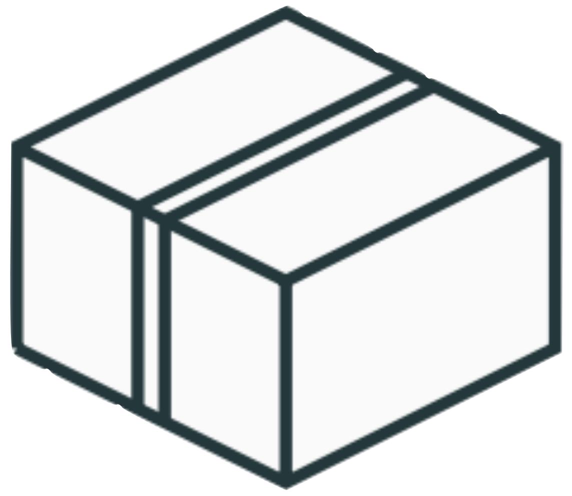

$$w_2 = \begin{cases} 0, \text{ w.p. } 1 - \varepsilon \\ \dfrac{1 + 2\varepsilon}{\varepsilon}, \text{ w.p. } \varepsilon \end{cases}$$

Motivation

1. RoE can perform poorly wrt the ex-post outcome, does not capture **risk aversion** (i.e., avoid the possibility of extremely bad outcomes).

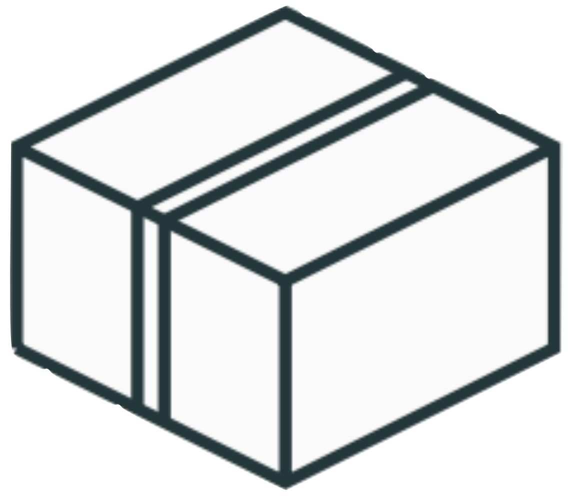

$w_1 = 1$

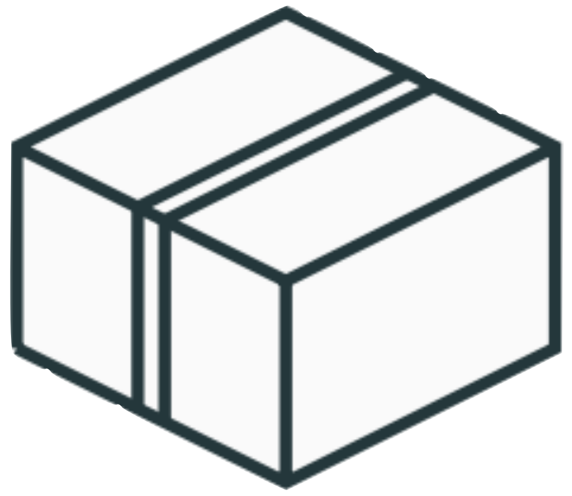

$$w_2 = \begin{cases} 0, \text{ w.p. } 1-\varepsilon \\ \frac{1+2\varepsilon}{\varepsilon}, \text{ w.p. } \varepsilon \end{cases}$$

Best ALG for **RoE** : Always choose the 2nd box

RoE $= \frac{1}{2}$

EoR $= \varepsilon$

Motivation

1. RoE can perform poorly wrt the ex-post outcome, does not capture **risk aversion** (i.e., avoid the possibility of extremely bad outcomes).

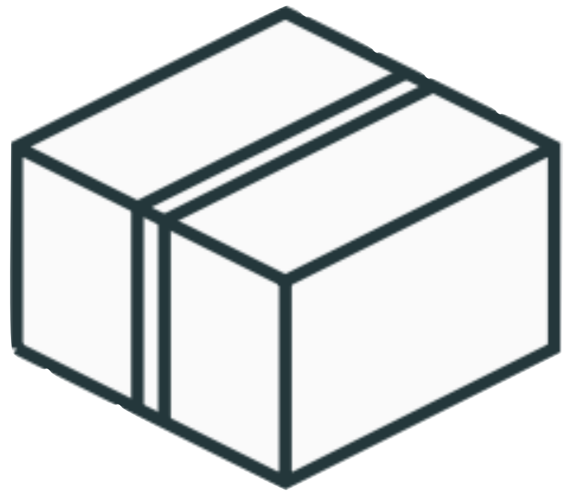

$w_1 = 1$

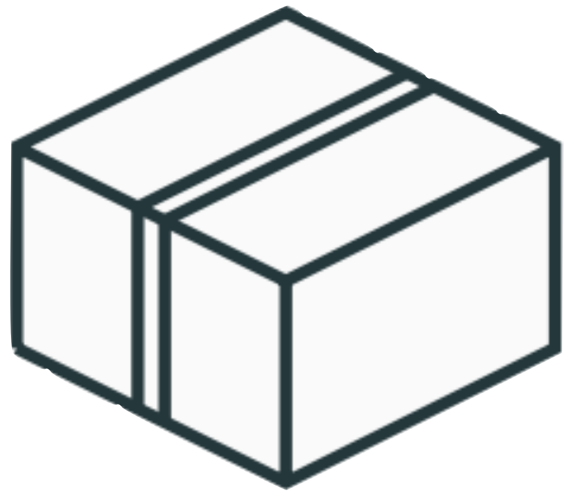

$$w_2 = \begin{cases} 0, \text{ w.p. } 1-\varepsilon \\ \dfrac{1+2\varepsilon}{\varepsilon}, \text{ w.p. } \varepsilon \end{cases}$$

Best ALG for **RoE** : Always choose the 2nd box

$$\textbf{RoE} = \frac{1}{2} \qquad \textbf{EoR} = \varepsilon$$

2. Find right generalization of PbM in **combinatorial** settings.

Motivation

1. RoE can perform poorly wrt the ex-post outcome, does not capture **risk aversion** (i.e., avoid the possibility of extremely bad outcomes).

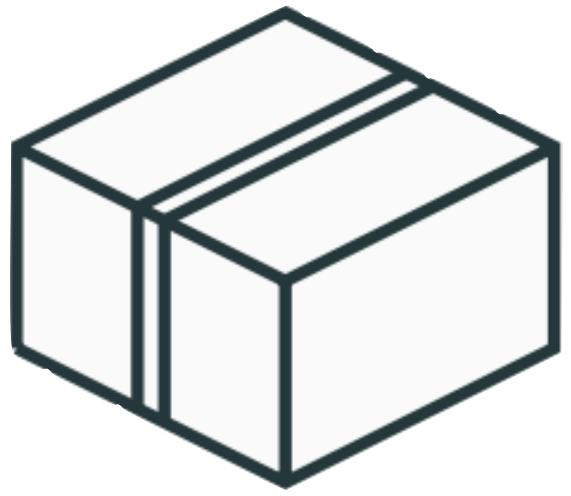

$w_1 = 1$

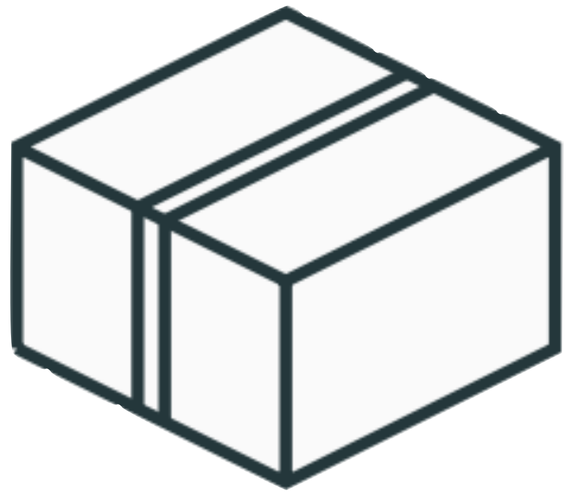

$$w_2 = \begin{cases} 0, \text{ w.p. } 1-\varepsilon \\ \dfrac{1+2\varepsilon}{\varepsilon}, \text{ w.p. } \varepsilon \end{cases}$$

Best ALG for **RoE** : Always choose the 2nd box

$$\textbf{RoE} = \frac{1}{2} \qquad \textbf{EoR} = \varepsilon$$

2. Find right generalization of PbM in **combinatorial** settings.

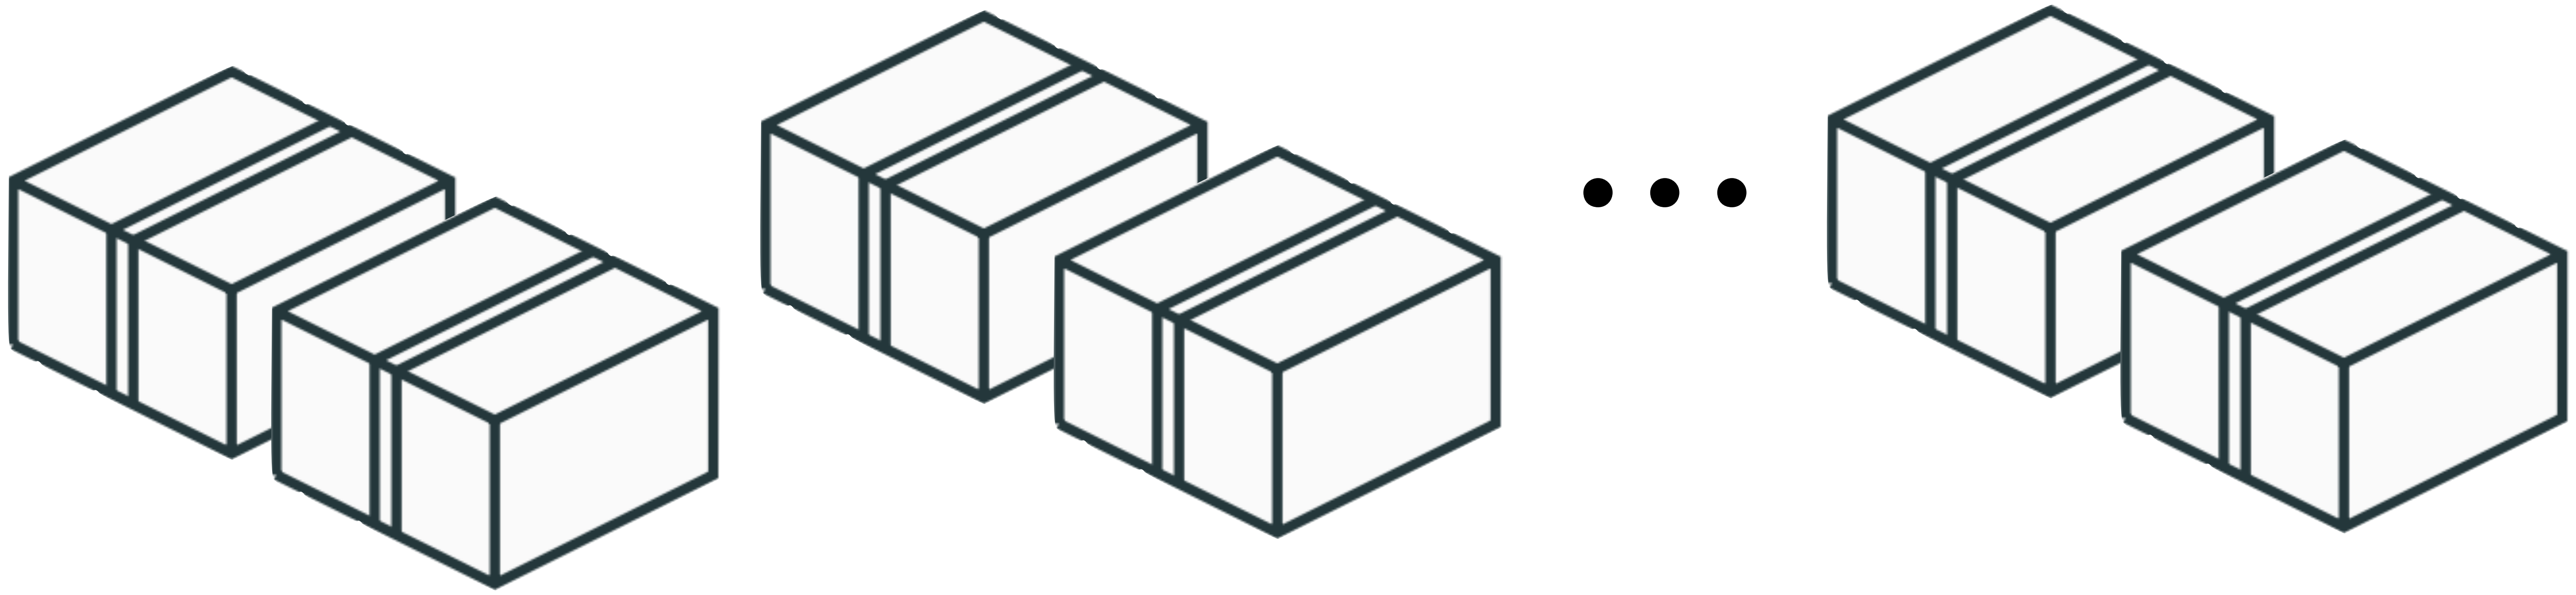

Motivation

1. RoE can perform poorly wrt the ex-post outcome, does not capture **risk aversion** (i.e., avoid the possibility of extremely bad outcomes).

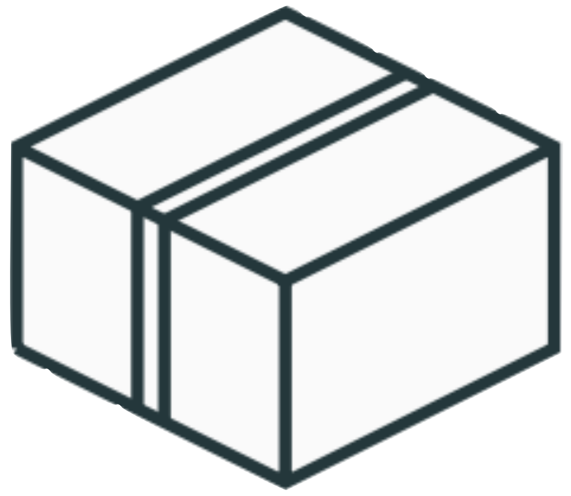

$w_1 = 1$

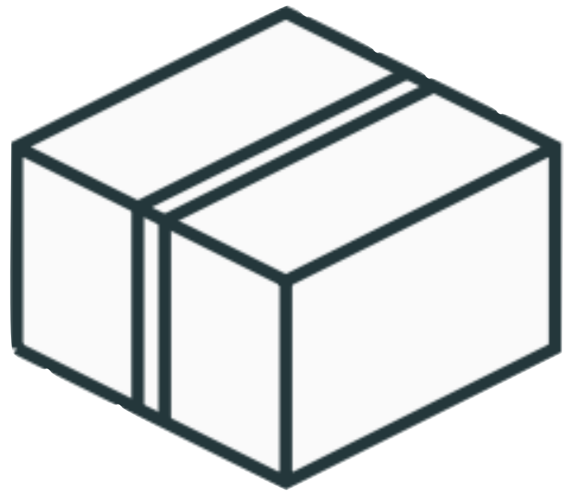

$$w_2 = \begin{cases} 0, \text{ w.p. } 1-\varepsilon \\ \dfrac{1+2\varepsilon}{\varepsilon}, \text{ w.p. } \varepsilon \end{cases}$$

Best ALG for **RoE** : Always choose the 2nd box

$$\textbf{RoE} = \frac{1}{2} \qquad \textbf{EoR} = \varepsilon$$

2. Find right generalization of PbM in **combinatorial** settings.

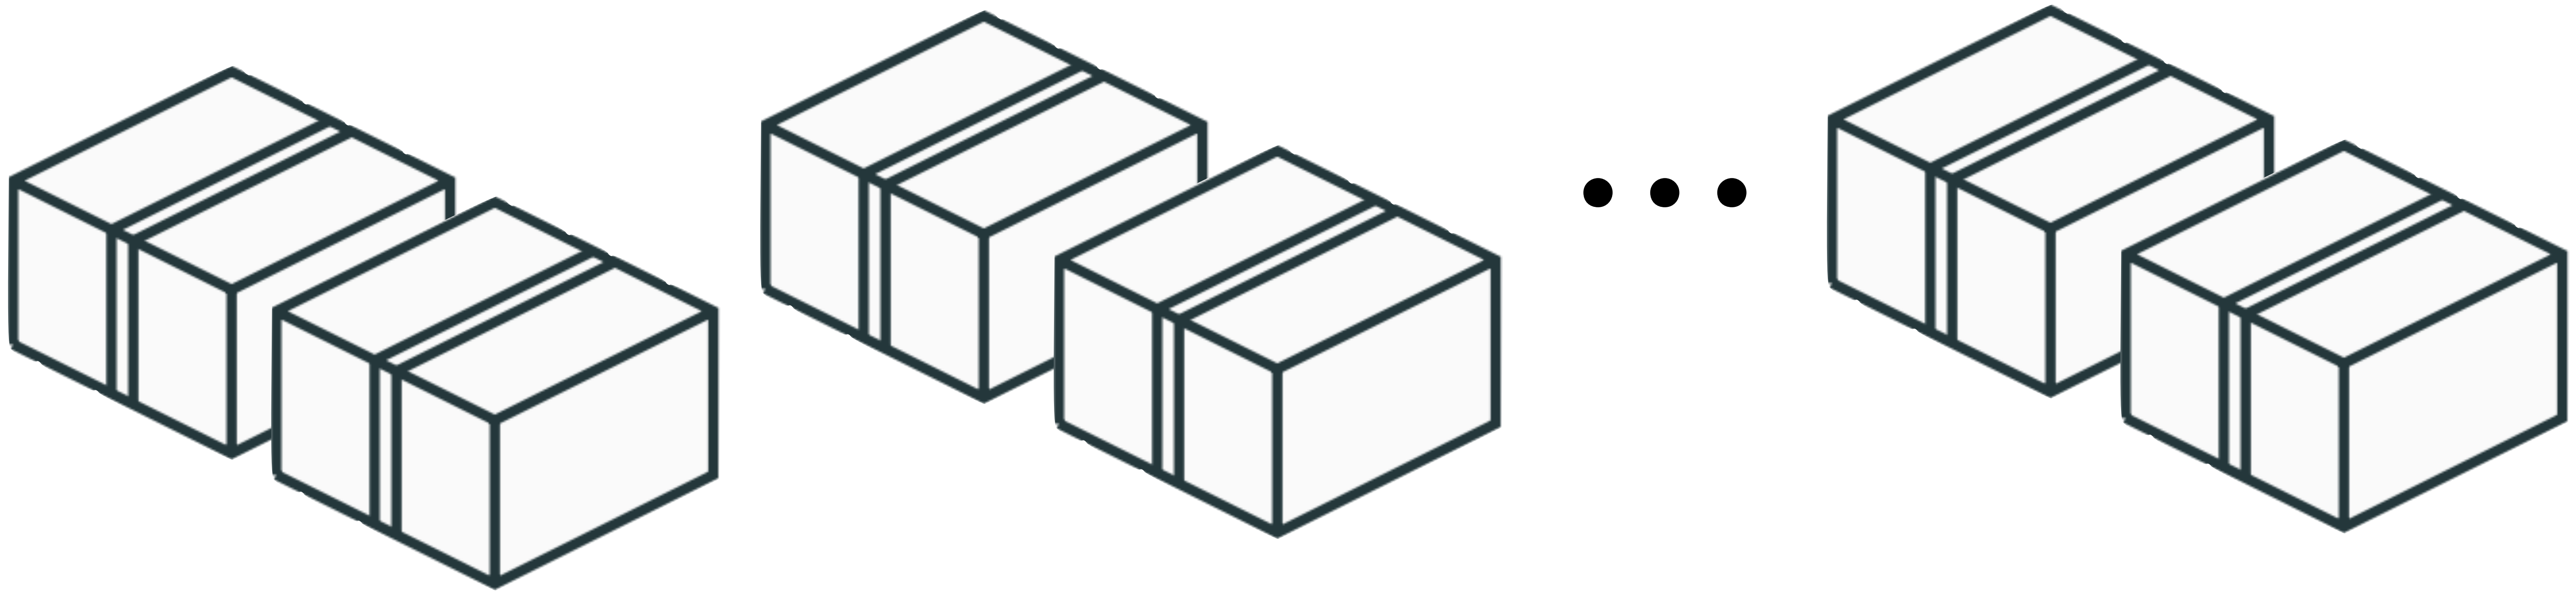

For each pair:

$$w_{1,i} = 1 \qquad w_{2,i} = \begin{cases} 0, \text{ w.p. } 1/2 \\ 2, \text{ w.p. } 1/2 \end{cases}$$

Motivation

1. RoE can perform poorly wrt the ex-post outcome, does not capture **risk aversion** (i.e., avoid the possibility of extremely bad outcomes).

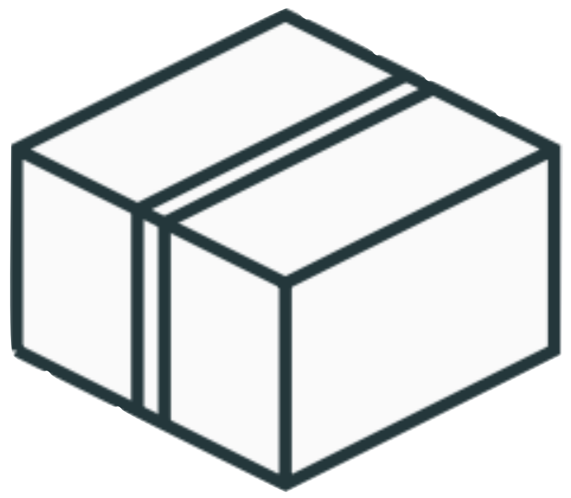

$w_1 = 1$

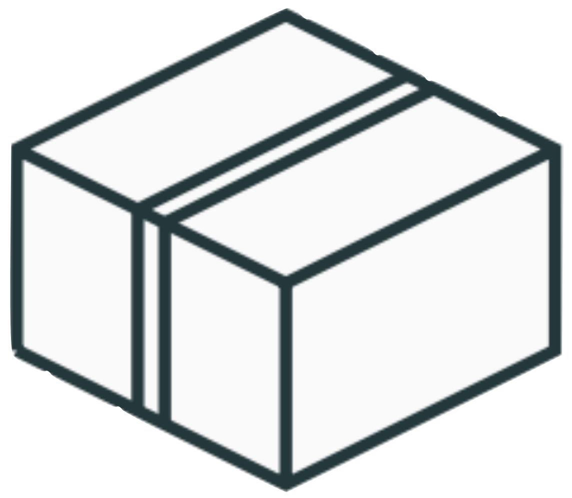

$$w_2 = \begin{cases} 0, \text{ w.p. } 1-\varepsilon \\ \dfrac{1+2\varepsilon}{\varepsilon}, \text{ w.p. } \varepsilon \end{cases}$$

Best ALG for **RoE** : Always choose the 2nd box

RoE $= \dfrac{1}{2}$ **EoR** $= \varepsilon$

2. Find right generalization of PbM in **combinatorial** settings.

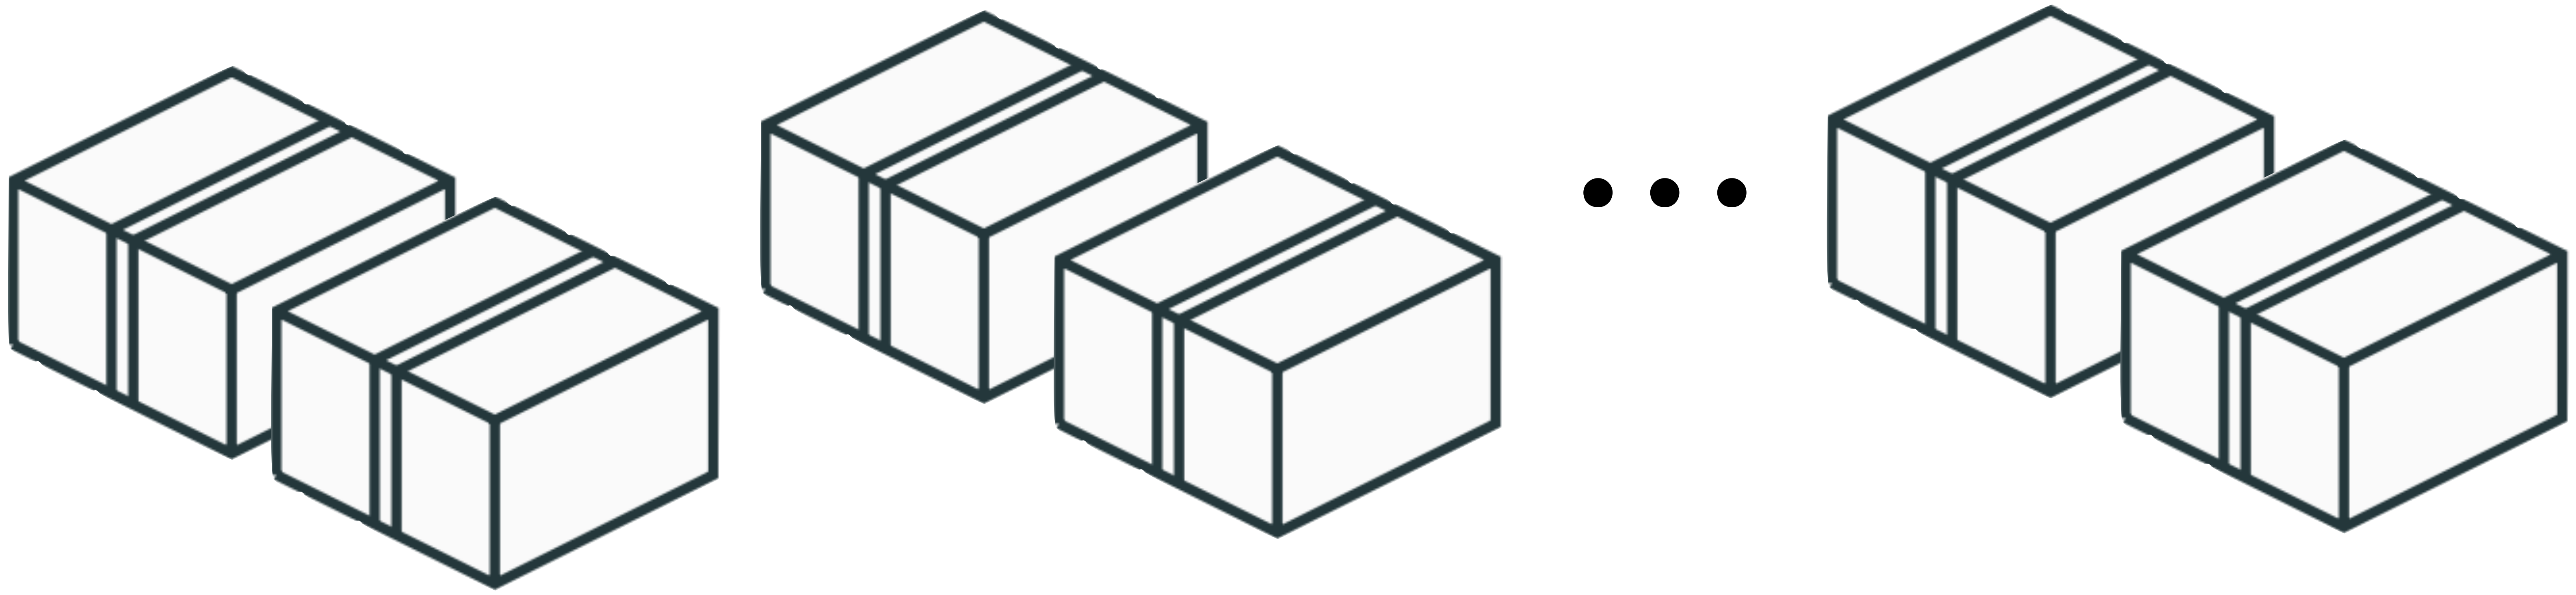

For each pair:

$w_{1,i} = 1$ $w_{2,i} = \begin{cases} 0, \text{ w.p. } 1/2 \\ 2, \text{ w.p. } 1/2 \end{cases}$

<u>Constraint</u>: Select **one** box from **each** pair

Motivation

1. RoE can perform poorly wrt the ex-post outcome, does not capture **risk aversion** (i.e., avoid the possibility of extremely bad outcomes).

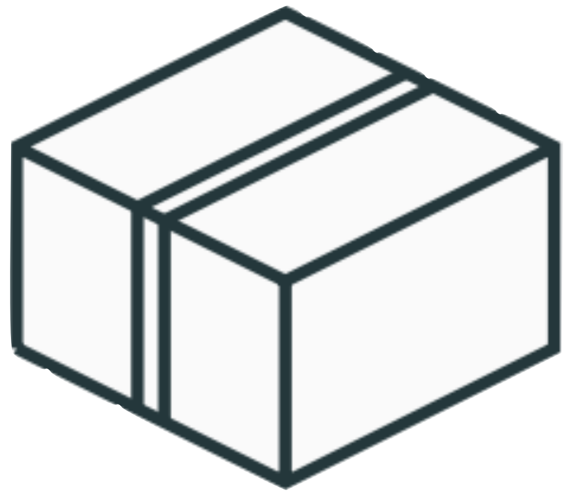

$w_1 = 1$

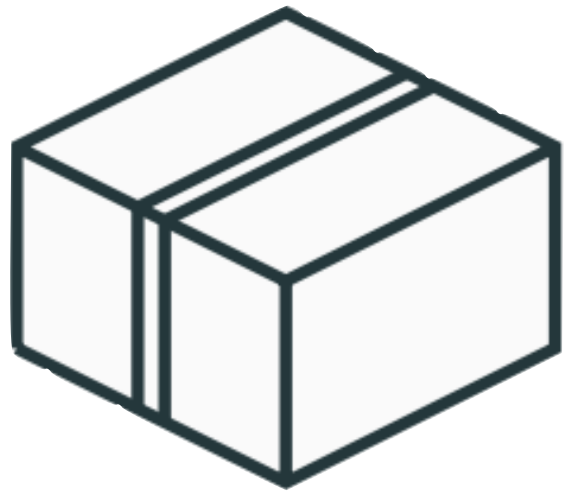

$$w_2 = \begin{cases} 0, \text{ w.p. } 1-\varepsilon \\ \dfrac{1+2\varepsilon}{\varepsilon}, \text{ w.p. } \varepsilon \end{cases}$$

Best ALG for **RoE** : Always choose the 2nd box

$$\textbf{RoE} = \frac{1}{2} \qquad \textbf{EoR} = \varepsilon$$

2. Find right generalization of PbM in **combinatorial** settings.

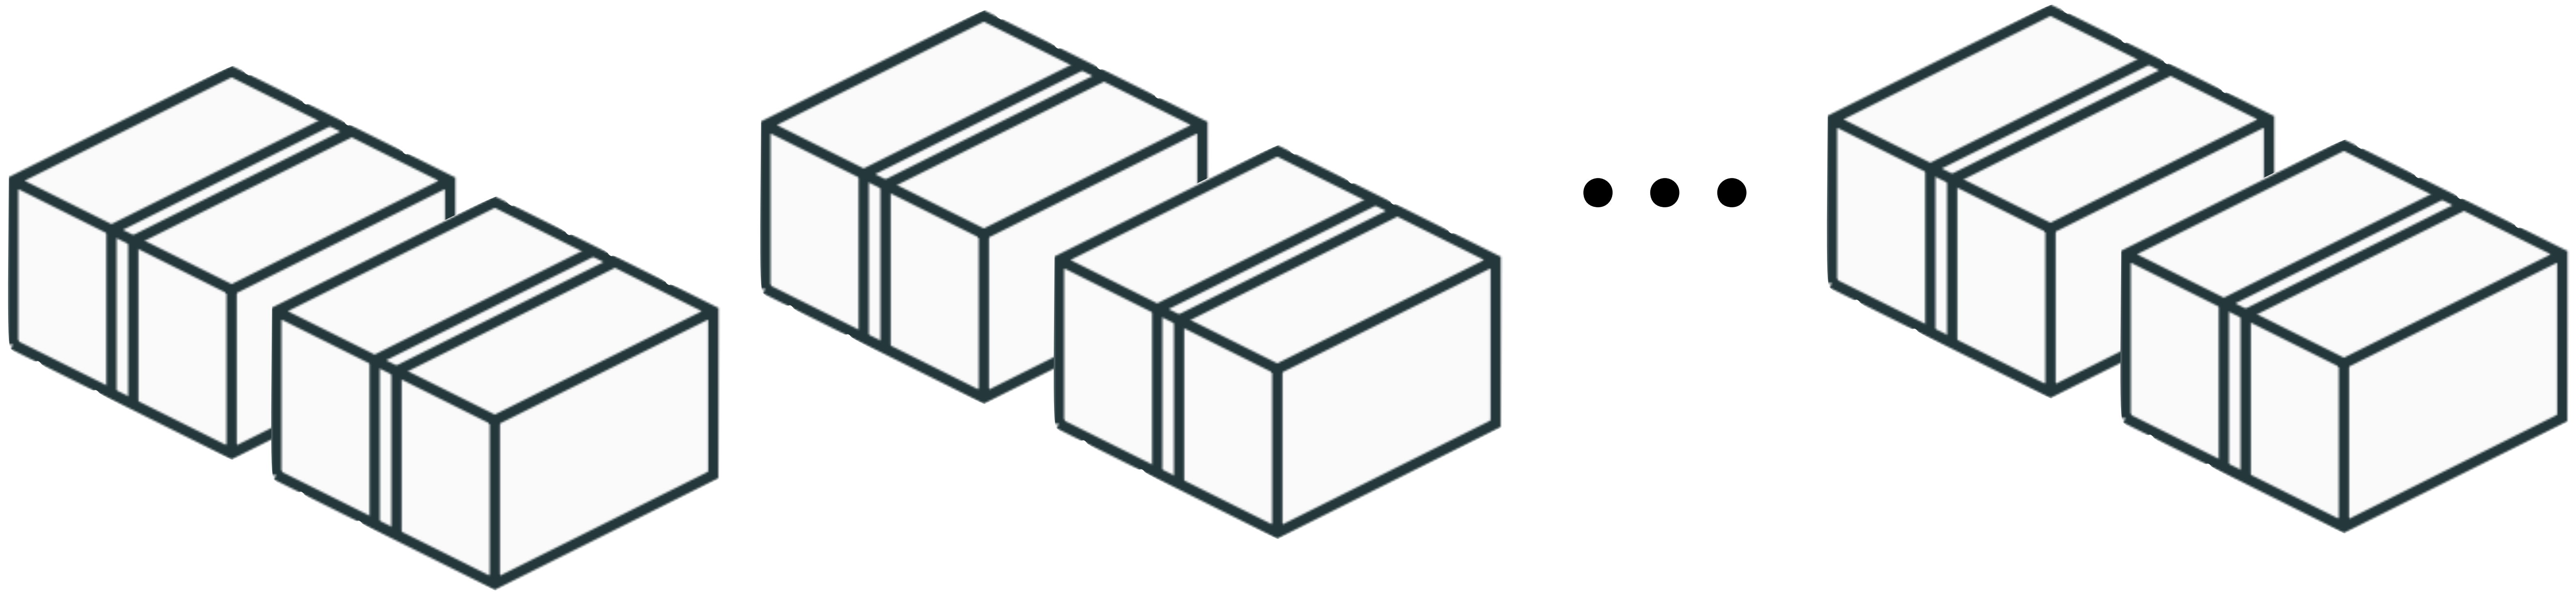

Constraint: Select **one** box from **each** pair

For each pair:

$$w_{1,i} = 1 \qquad w_{2,i} = \begin{cases} 0, \text{ w.p. } 1/2 \\ 2, \text{ w.p. } 1/2 \end{cases}$$

Jensen's ineq

$$\textbf{PbM} \le \frac{1}{2^n} \qquad \textbf{EoR} \ge \frac{2}{3}$$

Naive RoE to EoR (and vice versa)

Naive RoE to EoR (and vice versa)

- Evaluate RoE ALG with EoR

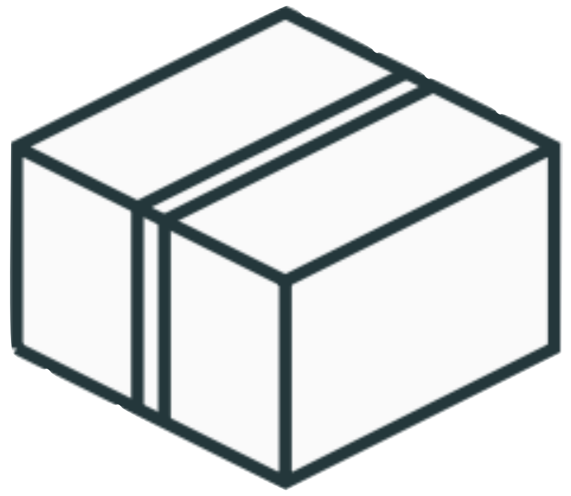

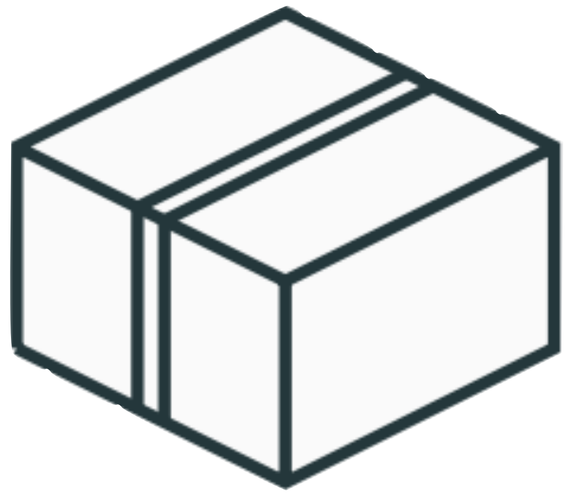

$w_1 = 1$

$$w_2 = \begin{cases} 0, \text{ w.p. } 1-\varepsilon \\ \dfrac{1+2\varepsilon}{\varepsilon}, \text{ w.p. } \varepsilon \end{cases}$$

Naive RoE to EoR (and vice versa)

- Evaluate RoE ALG with EoR

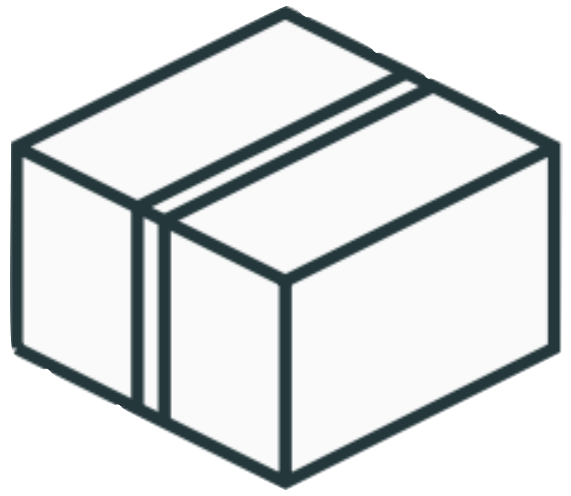

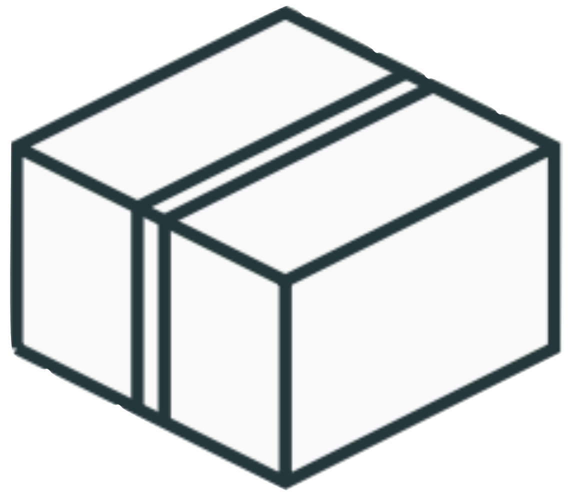

Optimal RoE ALG : Set threshold $\tau = \frac{\mathbb{E}[\text{OPT}]}{2}$

$w_1 = 1$

$$w_2 = \begin{cases} 0, \text{ w.p. } 1 - \varepsilon \\ \frac{1 + 2\varepsilon}{\varepsilon}, \text{ w.p. } \varepsilon \end{cases}$$

Naive RoE to EoR (and vice versa)

- Evaluate RoE ALG with EoR

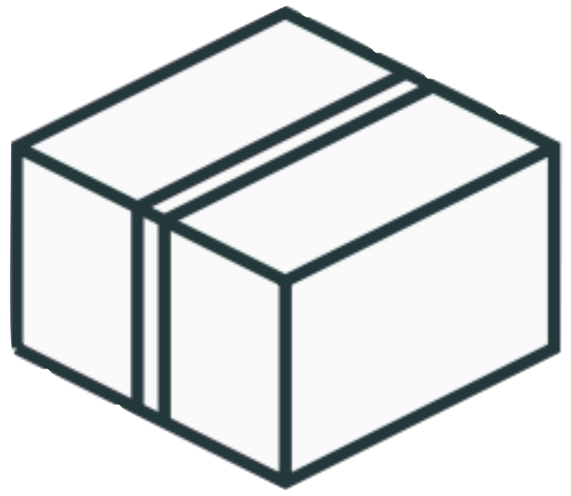

$w_1 = 1$

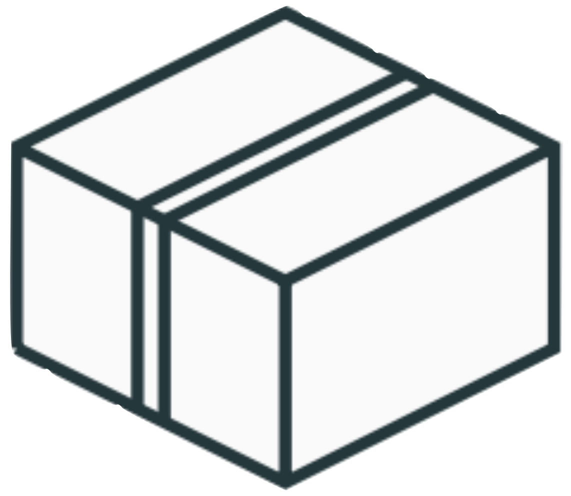

$$w_2 = \begin{cases} 0, \text{ w.p. } 1-\varepsilon \\ \frac{1+2\varepsilon}{\varepsilon}, \text{ w.p. } \varepsilon \end{cases}$$

Optimal RoE ALG : Set threshold $\tau = \frac{\mathbb{E}[\text{OPT}]}{2}$

This gives (tight) **RoE** $= \frac{1}{2}$ but **EoR** $= \varepsilon$!

Naive RoE to EoR (and vice versa)

- Evaluate RoE ALG with EoR

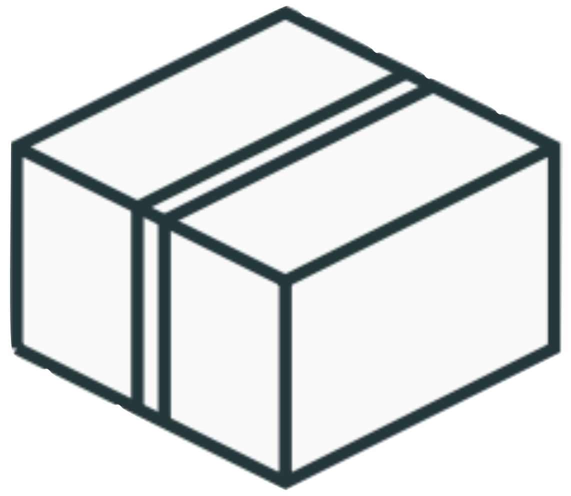

$w_1 = 1$

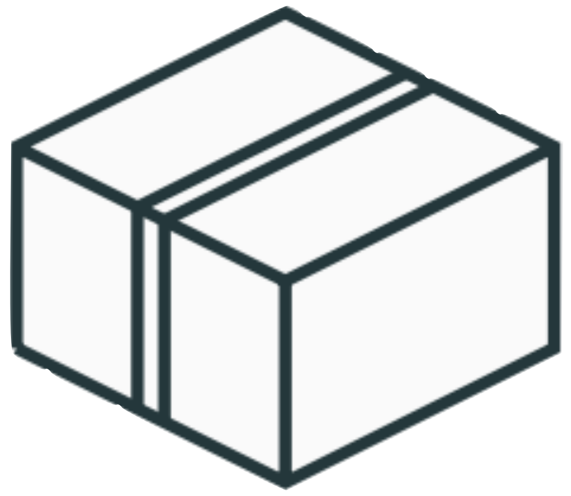

$$w_2 = \begin{cases} 0, \text{ w.p. } 1-\varepsilon \\ \frac{1+2\varepsilon}{\varepsilon}, \text{ w.p. } \varepsilon \end{cases}$$

Optimal RoE ALG : Set threshold $\tau = \frac{\mathbb{E}[\mathrm{OPT}]}{2}$

This gives (tight) **RoE** $= \frac{1}{2}$ but **EoR** $= \varepsilon$!

- Evaluate EoR ALG with RoE

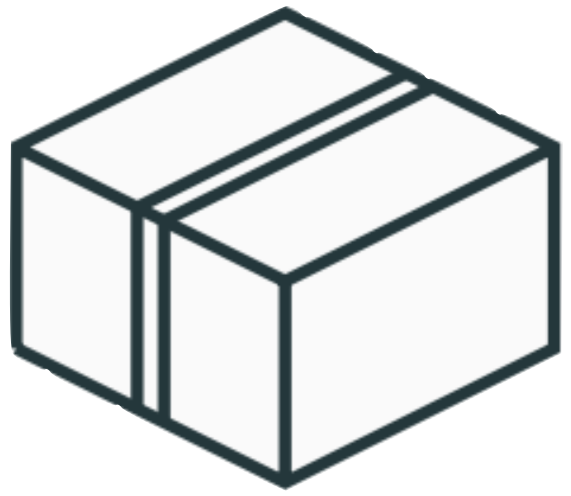

$w_1 = 1$

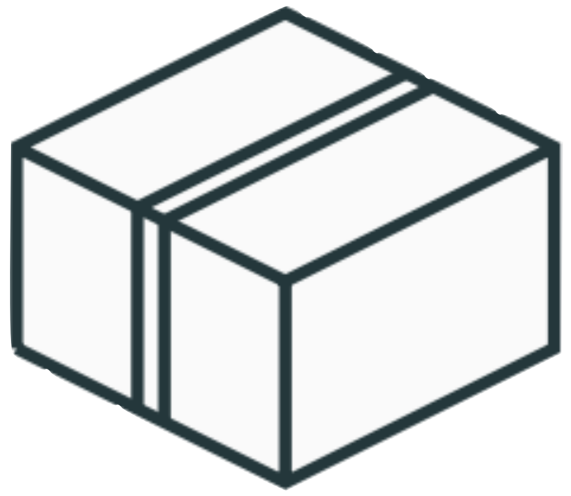

$$w_2 = \begin{cases} \varepsilon^2, \text{ w.p. } 1-\varepsilon \\ \frac{1}{\varepsilon^2}, \text{ w.p. } \varepsilon \end{cases}$$

Naive RoE to EoR (and vice versa)

- Evaluate RoE ALG with EoR

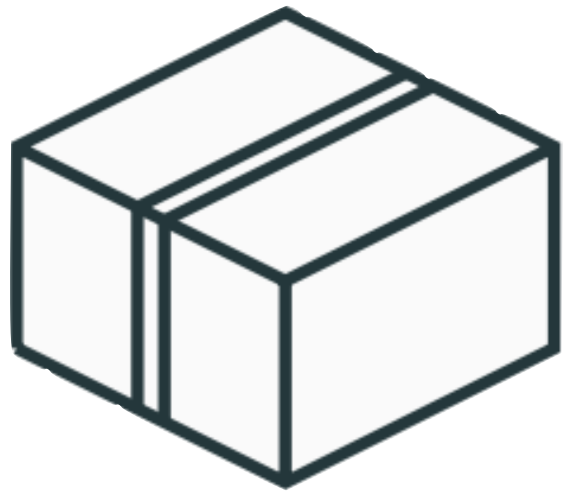

$w_1 = 1$

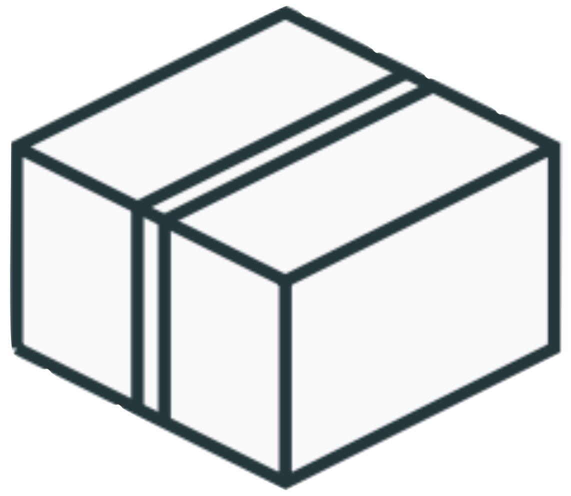

$$w_2 = \begin{cases} 0, \text{ w.p. } 1-\varepsilon \\ \frac{1+2\varepsilon}{\varepsilon}, \text{ w.p. } \varepsilon \end{cases}$$

Optimal RoE ALG : Set threshold $\tau = \frac{\mathbb{E}[\text{OPT}]}{2}$

This gives (tight) **RoE** $= \frac{1}{2}$ but **EoR** $= \varepsilon$!

- Evaluate EoR ALG with RoE

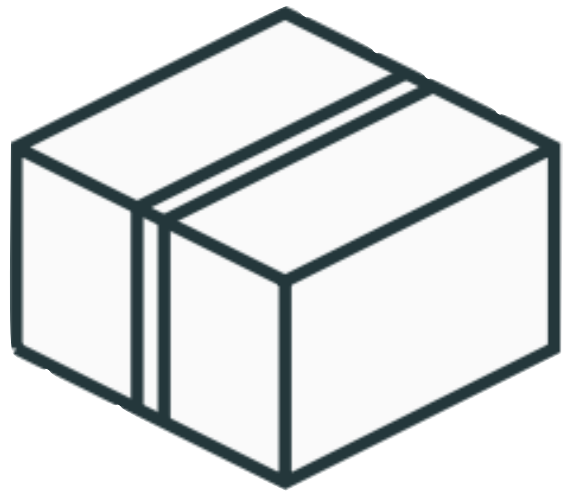

$w_1 = 1$

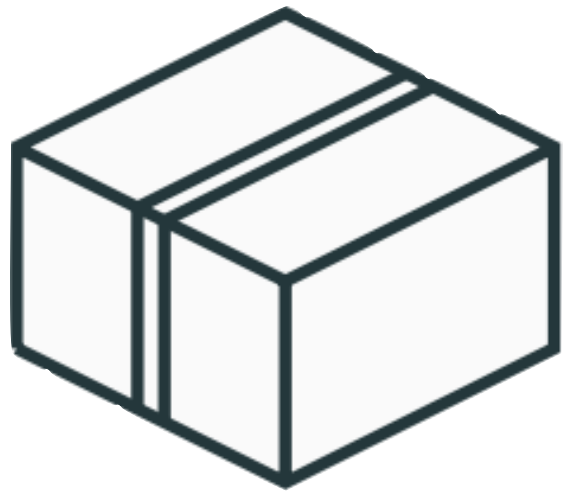

$$w_2 = \begin{cases} \varepsilon^2, \text{ w.p. } 1-\varepsilon \\ \frac{1}{\varepsilon^2}, \text{ w.p. } \varepsilon \end{cases}$$

Optimal EoR ALG : Always pick the first box .

Naive RoE to EoR (and vice versa)

- Evaluate RoE ALG with EoR

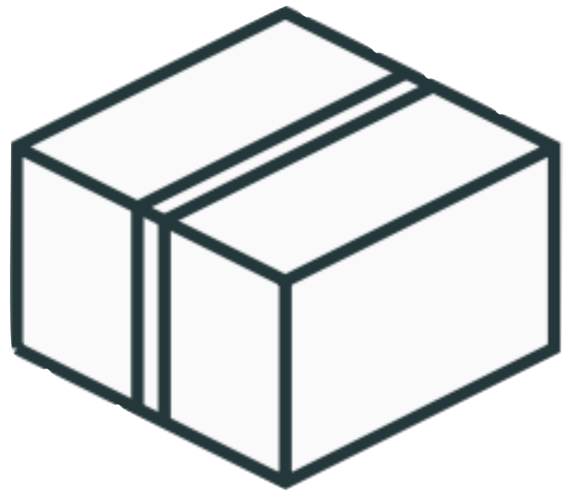

$w_1 = 1$

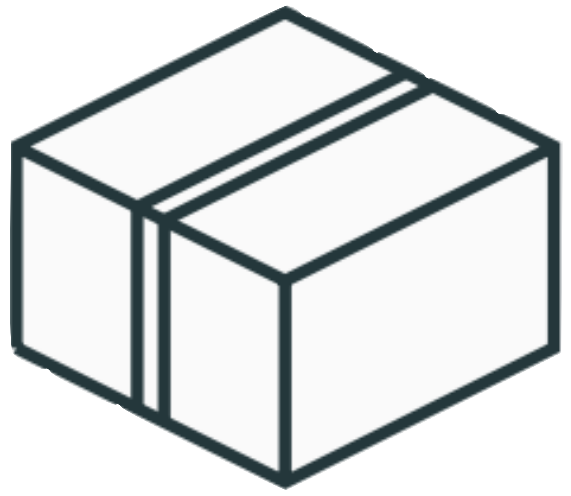

$$w_2 = \begin{cases} 0, \text{ w.p. } 1 - \varepsilon \\ \frac{1 + 2\varepsilon}{\varepsilon}, \text{ w.p. } \varepsilon \end{cases}$$

Optimal RoE ALG : Set threshold $\tau = \frac{\mathbb{E}[\text{OPT}]}{2}$

This gives (tight) **RoE** $= \frac{1}{2}$ but **EoR** $= \varepsilon$!

- Evaluate EoR ALG with RoE

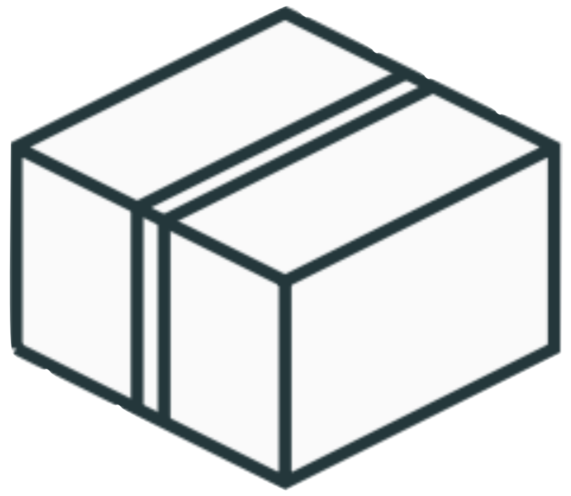

$w_1 = 1$

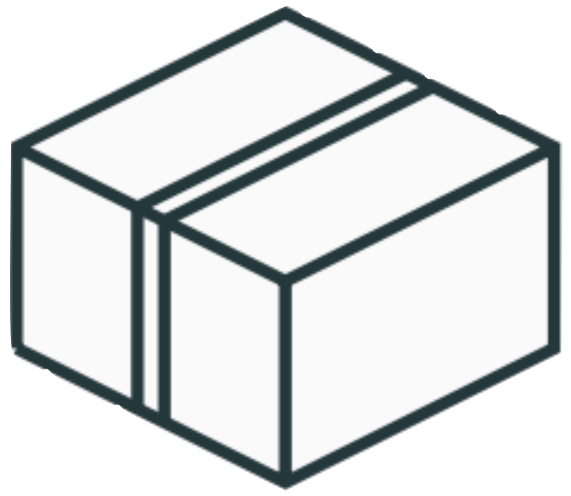

$$w_2 = \begin{cases} \varepsilon^2, \text{ w.p. } 1 - \varepsilon \\ \frac{1}{\varepsilon^2}, \text{ w.p. } \varepsilon \end{cases}$$

Optimal EoR ALG : Always pick the first box .

This gives **EoR** $> 1 - \varepsilon$ but **RoE** $< \varepsilon$!

What is the relation between RoE and EoR
in settings with general combinatorial constraints ?

Our results

Our results

Main Result (informal)

Two-way blackbox reduction: For every downward-closed constraint, RoE and EoR are at most a multiplicative **constant factor** apart.

Our results

Main Result (informal)

Two-way blackbox reduction: For every downward-closed constraint, RoE and EoR are at most a multiplicative **constant factor** apart.

Result 1 (warmup)

For single-choice settings **EoR = PbM** .

Our results

Main Result (informal)

Two-way blackbox reduction: For every downward-closed constraint, RoE and EoR are at most a multiplicative **constant factor** apart.

Result 1 (warmup)

For single-choice settings **EoR = PbM** .

Result 2 (RoE → EoR)

EoR$(\mathrm{F}) \geq$ **RoE**$(\mathrm{F})/12$

Our results

Main Result (informal)

Two-way blackbox reduction: For every downward-closed constraint, RoE and EoR are at most a multiplicative **constant factor** apart.

Result 1 (warmup)

For single-choice settings **EoR** $=$ **PbM** .

Result 2 (RoE $\rightarrow$ EoR)

EoR$(F) \geq$ **RoE**$(F)/12$

Result 3 (EoR $\rightarrow$ RoE)

RoE$(F) \geq$ **EoR**$(F)/18$

Preliminaries

Preliminaries

For $w_e \sim D_e, D = \times_{e \in E} D_e, w \in \mathbb{R}_{\geq 0}^{|E|}$, we define $\mathsf{OPT}(w) = \arg\max_{S \in \mathrm{F}} \sum_{e \in S} w_e$.

Preliminaries

For $w_e \sim D_e$, $D = \times_{e \in E} D_e$, $w \in \mathbb{R}_{\geq 0}^{|E|}$, we define $\textbf{OPT}(w) = \arg\max_{S \in F} \sum_{e \in S} w_e$.

With abuse of notation, our additive (for this talk) weight function is $w(S) := \sum_{e \in S} w_e$.

Preliminaries

For $w_e \sim D_e$, $D = \times_{e \in E} D_e$, $w \in \mathbb{R}_{\geq 0}^{|E|}$, we define $\mathsf{OPT}(w) = \arg\max_{S \in \mathrm{F}} \sum_{e \in S} w_e$.

With abuse of notation, our additive (for this talk) weight function is $w(S) := \sum_{e \in S} w_e$.

We denote by

$$f_{\mathrm{F}}(w) := w(\mathsf{OPT}(w))$$
$$a_{ALG}(w) := w(\mathsf{ALG}(w))$$

.

Preliminaries

For $w_e \sim D_e, D = \times_{e \in E} D_e, w \in \mathbb{R}_{\geq 0}^{|E|}$, we define $\mathsf{OPT}(w) = \arg\max_{S \in \mathrm{F}} \sum_{e \in S} w_e$.

With abuse of notation, our additive (for this talk) weight function is $w(S) := \sum_{e \in S} w_e$.

We denote by
$$f_{\mathrm{F}}(w) := w(\mathsf{OPT}(w))$$
$$a_{ALG}(w) := w(\mathsf{ALG}(w))$$
.

Now the objective is $\mathrm{EoR}(\mathbf{F}, \boldsymbol{D}, \mathrm{ALG}) := \mathbb{E}\left[\frac{a(w)}{f(w)}\right] \Rightarrow \mathrm{EoR}(\mathbf{F}) := \inf_{D} \sup_{ALG} \mathbb{E}\left[\frac{a(w)}{f(w)}\right]$.

Preliminaries

For $w_e \sim D_e, D = \times_{e \in E} D_e, w \in \mathbb{R}_{\geq 0}^{|E|}$, we define $\textbf{OPT}(w) = \arg\max_{S \in \mathrm{F}} \sum_{e \in S} w_e$.

With abuse of notation, our additive (for this talk) weight function is $w(S) := \sum_{e \in S} w_e$.

We denote by $\begin{matrix} f_{\mathrm{F}}(w) := w(\textbf{OPT}(w)) \\ a_{ALG}(w) := w(\textbf{ALG}(w)) \end{matrix}$.

Now the objective is $\mathrm{EoR}(\mathbf{F}, \boldsymbol{D}, \mathrm{ALG}) := \mathbb{E}\left[\frac{a(w)}{f(w)}\right] \Rightarrow \mathrm{EoR}(\mathbf{F}) := \inf_{D} \sup_{ALG} \mathbb{E}\left[\frac{a(w)}{f(w)}\right]$.

Analogously: $\mathrm{RoE}(\mathbf{F}) := \inf_{D} \sup_{ALG} \frac{\mathbb{E}\left[a(w)\right]}{\mathbb{E}\left[f(w)\right]}$ and $\mathrm{PbM}(\mathbf{F}) := \inf_{D} \sup_{ALG} \Pr\left[a(w) = f(w)\right]$.

Single-choice PI: EoR = PbM

Single-choice PI: EoR = PbM

In fact, the statement says sth stronger:

Thm : For each product distr. D, we can construct a new product distr. D' for which EoR is abritrarily close to the PbM of the original distribution.

Single-choice PI: EoR = PbM

In fact, the statement says sth stronger:

Thm : For each product distr. D, we can construct a new product distr. D' for which EoR is abritrarily close to the PbM of the original distribution.

Corollary: The gap between RoE and EoR is at least $2/e$, since

$$\text{RoE}(\mathbf{F}) = \frac{1}{2} > \frac{1}{e} = \text{EoR}(\mathbf{F}) \text{ for fixed order.}$$

Reduction RoE $\rightarrow$ EoR : some tools

Reduction RoE → EoR : some tools

We define the **threshold** $\Pr\left[\tau \geq \max_{e \in E} w_e\right] = \gamma \in (0,1)$.

Reduction RoE $\rightarrow$ EoR : some tools

We define the **threshold** $\Pr\left[\tau \geq \max_{e\in E} w_e\right] = \gamma \in (0,1)$.

For events $\mathscr{E}_0 := \left\{\forall e \in E : w_e \leq \tau\right\}$, $\mathscr{E}_1 := \left\{\exists! e \in E : w_e > \tau\right\}$ we have $\Pr\left[\mathscr{E}_0\right] = \gamma$, $\Pr\left[\mathscr{E}_1\right] \geq \gamma \log\left(\frac{1}{\gamma}\right)$.

Reduction RoE → EoR : some tools

We define the **threshold** $\Pr\left[\tau \geq \max_{e \in E} w_e\right] = \gamma \in (0,1)$.

For events $\mathscr{E}_0 := \{\forall e \in E : w_e \leq \tau\}$, $\mathscr{E}_1 := \{\exists ! e \in E : w_e > \tau\}$ we have $\Pr\left[\mathscr{E}_0\right] = \gamma$, $\Pr\left[\mathscr{E}_1\right] \geq \gamma \log\left(\frac{1}{\gamma}\right)$.

Truncated distr. $\overline{D}_e := D_{e | w_e \leq \tau}$.

Observation: $\overline{D} \Leftrightarrow D \mid \mathscr{E}_0$.

Reduction RoE → EoR : some tools

We define the **threshold** $\Pr\left[\tau \geq \max_{e \in E} w_e\right] = \gamma \in (0,1)$.

For events $\mathscr{E}_0 := \{\forall e \in E : w_e \leq \tau\}$, $\mathscr{E}_1 := \{\exists! e \in E : w_e > \tau\}$ we have $\Pr[\mathscr{E}_0] = \gamma$, $\Pr[\mathscr{E}_1] \geq \gamma \log\left(\frac{1}{\gamma}\right)$.

Truncated distr. $\overline{D}_e := D_{e|w_e \leq \tau}$.

Observation: $\overline{D} \Leftrightarrow D \mid \mathscr{E}_0$.

Useful **Lemma** : $f(w) \leq f(\overline{w}) + \sum_{e \in E} w_e \cdot 1\left[w_e > \tau\right]$.

Reduction RoE → EoR : our algorithm

Reduction RoE → EoR : our algorithm

Idea: Case distinction when prophet's value comes from a small or a large number of boxes.

Reduction RoE → EoR : our algorithm

Idea: Case distinction when prophet's value comes from a small or a large number of boxes.

Assumption: We have ALG_{RoE} for which $\mathbb{E}\left[a(w) \mid \mathcal{E}_0\right] \geq \alpha \cdot \mathbb{E}\left[f(w) \mid \mathcal{E}_0\right]$.

Reduction RoE → EoR : our algorithm

Idea: Case distinction when prophet's value comes from a small or a large number of boxes.

Assumption: We have ALG_{RoE} for which $\mathbb{E}\left[a(w) \mid \mathscr{E}_0\right] \geq \alpha \cdot \mathbb{E}\left[f(w) \mid \mathscr{E}_0\right]$.

Parameters: $\gamma \in (0,1)$, $c > 0$.

Output: Feasible set $\text{ALG}(w)$.

Reduction RoE → EoR : our algorithm

Idea: Case distinction when prophet's value comes from a small or a large number of boxes.

Assumption: We have ALG_{RoE} for which $\mathbb{E}\left[a(w) \mid \mathscr{E}_0\right] \geq \alpha \cdot \mathbb{E}\left[f(w) \mid \mathscr{E}_0\right]$.

Parameters: $\gamma \in (0,1)$, $c > 0$.

Output: Feasible set $\text{ALG}(w)$.

If $\mathbb{E}[f(\overline{w})] \leq c \cdot \tau$ then:

"Catch the superstar"

else:

"Run the Combinatorial Algorithm"

Analysis of RoE-to-EoR algo: "Catch the superstar"

Analysis of RoE-to-EoR algo: "Catch the superstar"

If "normalized" OPT $\mathbb{E}\left[f(\overline{w})\right]$ not too big, then imagine that few boxes contribute the most.

Analysis of RoE-to-EoR algo: "Catch the superstar"

If "normalized" OPT $\mathbb{E}\left[f(\overline{w})\right]$ not too big, then imagine that few boxes contribute the most.

Count **contribution only** from the cases for which exactly one $w_e \geq \tau$ (i.e., superstar element).
Ignore other cases.

Analysis of RoE-to-EoR algo: "Catch the superstar"

If "normalized" OPT $\mathbb{E}\left[f(\overline{w})\right]$ not too big, then imagine that few boxes contribute the most.

Count **contribution only** from the cases for which exactly one $w_e \geq \tau$ (i.e., superstar element).
Ignore other cases.

$$\Pr\left[\mathscr{E}_1\right] \geq \gamma \log\left(\frac{1}{\gamma}\right)$$

Analysis of RoE-to-EoR algo: "Catch the superstar"

If "normalized" OPT $\mathbb{E}\left[f(\overline{w})\right]$ not too big, then imagine that few boxes contribute the most.

Count **contribution only** from the cases for which exactly one $w_e \geq \tau$ (i.e., superstar element).
Ignore other cases.

$$\Pr\left[\mathscr{E}_1\right] \geq \gamma \log\left(\frac{1}{\gamma}\right)$$

If we catch the (unique) superstar and stop, we recover **constant fraction** $\Lambda(c)$ of $f(w)$.

Analysis of RoE-to-EoR algo: "Catch the superstar"

If "normalized" OPT $\mathbb{E}\left[f(\overline{w})\right]$ not too big, then imagine that few boxes contribute the most.

Count **contribution only** from the cases for which exactly one $w_e \geq \tau$ (i.e., superstar element). Ignore other cases.

$$\Pr\left[\mathscr{E}_1\right] \geq \gamma \log\left(\frac{1}{\gamma}\right)$$

If we catch the (unique) superstar and stop, we recover **constant fraction** $\Lambda(c)$ of $f(w)$.

Then, for Case 1: $\mathbb{E}\left[\frac{a(w)}{f(w)}\right] \geq \Pr[\mathscr{E}_1] \cdot \Lambda(c) = O(1)$.

Analysis of RoE-to-EoR: "Run the Combinatorial algorithm"

Analysis of RoE-to-EoR: "Run the Combinatorial algorithm"

If "normalized" OPT $\mathbb{E}\left[f(\overline{w})\right]$ is big, then imagine that many boxes contribute to it.

Analysis of RoE-to-EoR: "Run the Combinatorial algorithm"

If "normalized" OPT $\mathbb{E}\left[f(\overline{w})\right]$ is big, then imagine that many boxes contribute to it.

Count contribution only when no w_e exceeds τ. Ignore other cases.

Analysis of RoE-to-EoR: "Run the Combinatorial algorithm"

If "normalized" OPT $\mathbb{E}\left[f(\overline{w})\right]$ is big, then imagine that many boxes contribute to it.

Count contribution only when no w_e exceeds τ. Ignore other cases.

$$\mathbf{Pr}\left[\mathscr{E}_0\right] = \gamma$$

Analysis of RoE-to-EoR: "Run the Combinatorial algorithm"

If "normalized" OPT $\mathbb{E}\left[f(\overline{w})\right]$ is big, then imagine that many boxes contribute to it.

Count contribution only when no w_e exceeds τ. Ignore other cases.

$$\mathbf{Pr}\left[\mathscr{E}_0\right] = \gamma$$

Technical ingredient: For every $\tau > 0$, the function f/τ restricted to $[0,\tau]^{|E|}$ is ***self-bounding***

Analysis of RoE-to-EoR: "Run the Combinatorial algorithm"

If "normalized" OPT $\mathbb{E}\left[f(\overline{w})\right]$ is big, then imagine that many boxes contribute to it.

Count contribution only when no w_e exceeds τ. Ignore other cases.

$$\Pr\left[\mathcal{E}_0\right] = \gamma$$

Technical ingredient: For every $\tau > 0$, the function f/τ restricted to $[0,\tau]^{|E|}$ is ***self-bounding***

↓

Sharp concentration around its expectation!
[Boucheron, Lugosi, Massart '00]

Analysis of RoE-to-EoR: "Run the Combinatorial algorithm"

If "normalized" OPT $\mathbb{E}\left[f(\overline{w})\right]$ is big, then imagine that many boxes contribute to it.

Count contribution only when no w_e exceeds τ. Ignore other cases.

$$\Pr\left[\mathcal{E}_0\right] = \gamma$$

Technical ingredient: For every $\tau > 0$, the function f/τ restricted to $[0,\tau]^{|E|}$ is ***self-bounding***

↓

Sharp concentration around its expectation!
[Boucheron, Lugosi, Massart '00]

Then proof proceeds in two high-level steps:

1. UB value of ex-post OPT with *high enough* (constant) probability.
2. LB expected value of ALG_{RoE} | ex-post OPT not too large.

Analysis of RoE-to-EoR: "Run the Combinatorial algorithm"

If "normalized" OPT $\mathbb{E}\left[f(\overline{w})\right]$ is big, then imagine that many boxes contribute to it.

Count contribution only when no w_e exceeds τ. Ignore other cases.

$$\Pr\left[\mathscr{E}_0\right] = \gamma$$

Technical ingredient: For every $\tau > 0$, the function f/τ restricted to $[0,\tau]^{|E|}$ is ***self-bounding***

↓

Sharp concentration around its expectation!
[Boucheron, Lugosi, Massart '00]

Then proof proceeds in two high-level steps:

1. UB value of ex-post OPT with *high enough* (constant) probability.
2. LB expected value of ALG_{RoE} | ex-post OPT not too large.

All in all, for Case 2: $\mathbb{E}\left[\frac{a(w)}{f(w)}\right] \geq \Pr[\mathscr{E}_0] \cdot \alpha \cdot \Gamma(k,\delta) = O(1)$.

What's the best we can hope for the ex-post opt ?

What's the best we can hope for the ex-post opt ?

Observation: Our reduction achieves an $O(\alpha)$ – approximation of the ex-post OPT with constant probability (independent of α).

What's the best we can hope for the ex-post opt ?

Observation: Our reduction achieves an $O(\alpha)$ – approximation of the ex-post OPT with constant probability (independent of α).

Maximizing the EoR implies the above and is the best we can achieve (up to constant terms).

What's the best we can hope for the ex-post opt ?

Observation: Our reduction achieves an $O(\alpha) -$ approximation of the ex-post OPT with constant probability (independent of α).

Maximizing the EoR implies the above and is the best we can achieve (up to constant terms).

No $O(1) -$ approximation with $(1 - \varepsilon)$ prob.

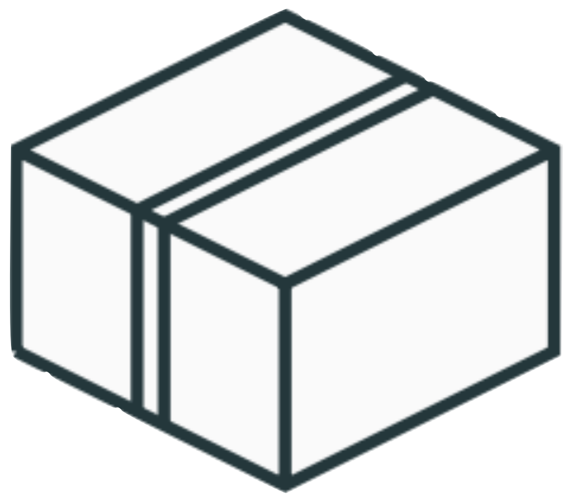

$w_1 = 1$

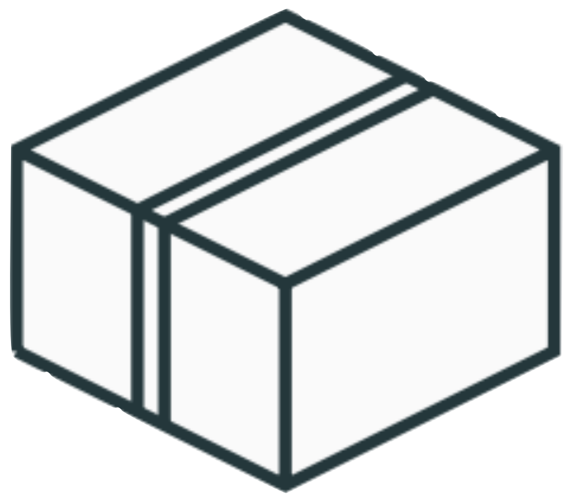

$$w_2 = \begin{cases} \varepsilon, \text{ w.p. } 1/2 \\ \dfrac{1}{\varepsilon}, \text{ w.p. } 1/2 \end{cases}$$

What's the best we can hope for the ex-post opt ?

Observation: Our reduction achieves an $O(\alpha)$ – approximation of the ex-post OPT with constant probability (independent of α).

Maximizing the EoR implies the above and is the best we can achieve (up to constant terms).

No $O(1)$ – approximation with $(1-\varepsilon)$ prob.

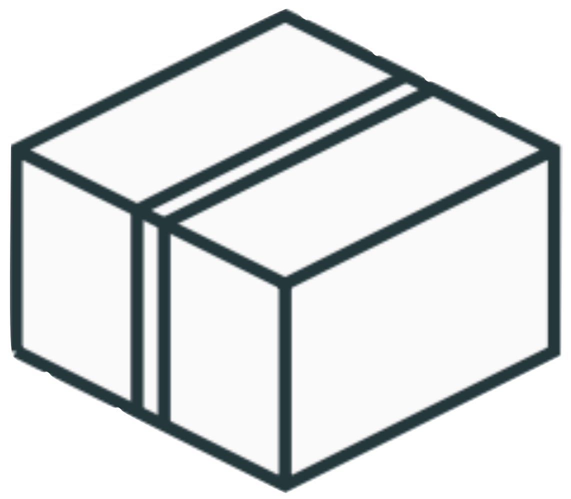

$w_1 = 1$

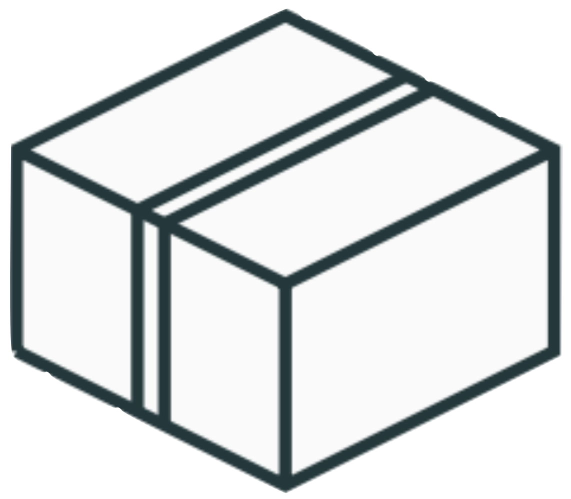

$$w_2 = \begin{cases} \varepsilon, \text{ w.p. } 1/2 \\ \frac{1}{\varepsilon}, \text{ w.p. } 1/2 \end{cases}$$

No algo can do $> \varepsilon$ – approx.
with prob. $> 1/2$.

What's the best we can hope for the ex-post opt ?

Observation: Our reduction achieves an $O(\alpha)$ – approximation of the ex-post OPT with constant probability (independent of α).

Maximizing the EoR implies the above and is the best we can achieve (up to constant terms).

No $O(1)$ – approximation with $(1 - \varepsilon)$ prob.

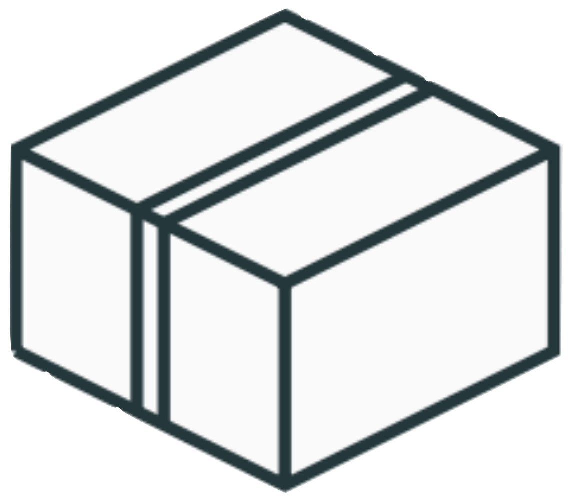

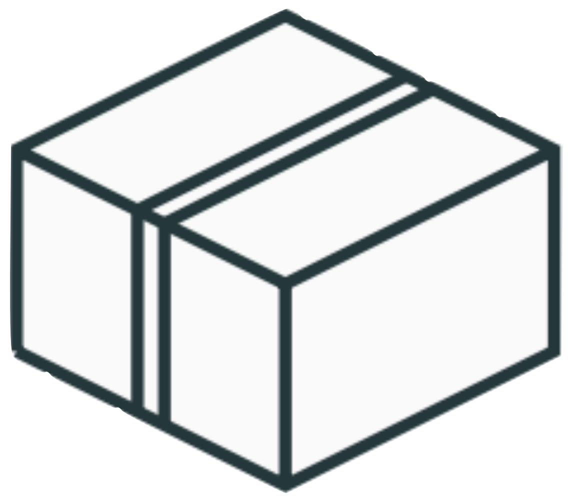

$w_1 = 1$

$$w_2 = \begin{cases} \varepsilon, \text{ w.p. } 1/2 \\ \frac{1}{\varepsilon}, \text{ w.p. } 1/2 \end{cases}$$

No algo can do $> \varepsilon$ – approx. with prob. $> 1/2$.

No $(1 - \varepsilon)$ – approximation with $O(1)$ prob.

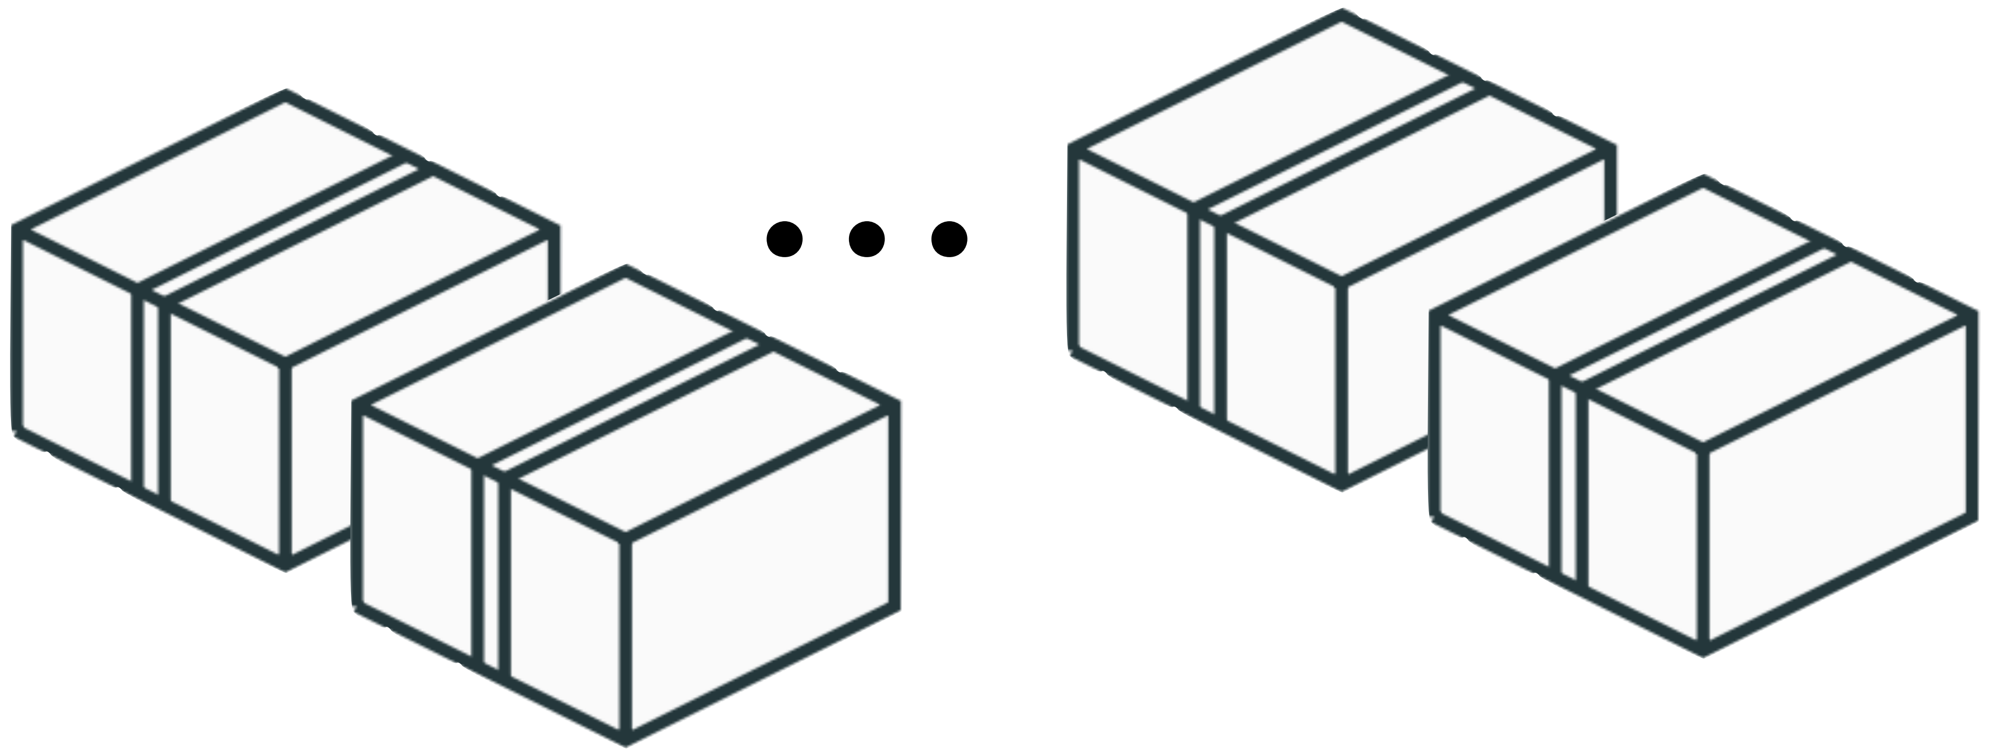

For each pair: $w_{1,i} = 1$ $\quad w_{2,i} = \begin{cases} 0, \text{ w.p. } 1/2 \\ 2, \text{ w.p. } 1/2 \end{cases}$

What's the best we can hope for the ex-post opt ?

Observation: Our reduction achieves an $O(\alpha)$ – approximation of the ex-post OPT with constant probability (independent of α).

Maximizing the EoR implies the above and is the best we can achieve (up to constant terms).

No $O(1)$ – approximation with $(1-\varepsilon)$ prob.

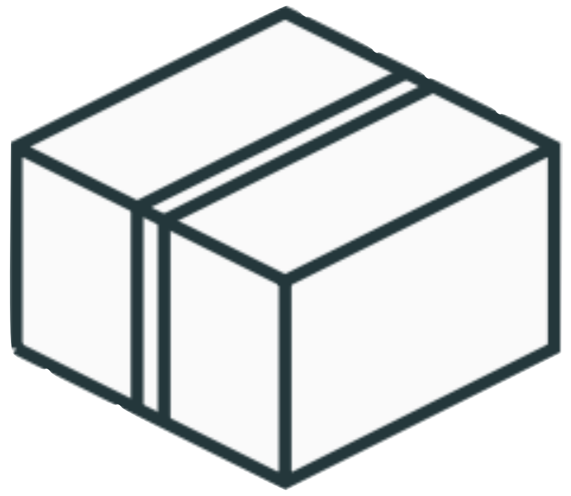

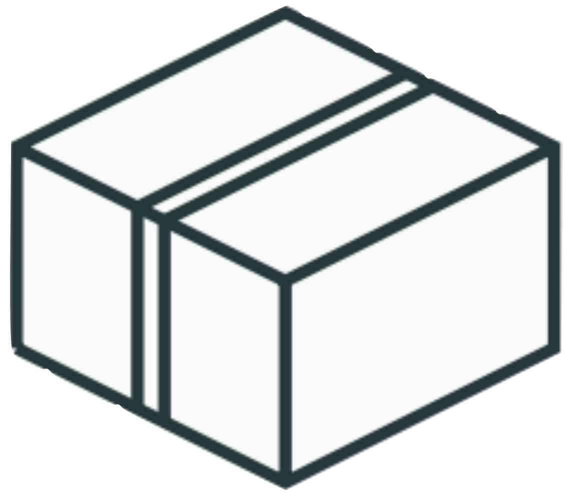

$w_1 = 1$

$$w_2 = \begin{cases} \varepsilon, \text{ w.p. } 1/2 \\ \frac{1}{\varepsilon}, \text{ w.p. } 1/2 \end{cases}$$

No algo can do $> \varepsilon$ – approx. with prob. $> 1/2$.

No $(1-\varepsilon)$ – approximation with $O(1)$ prob.

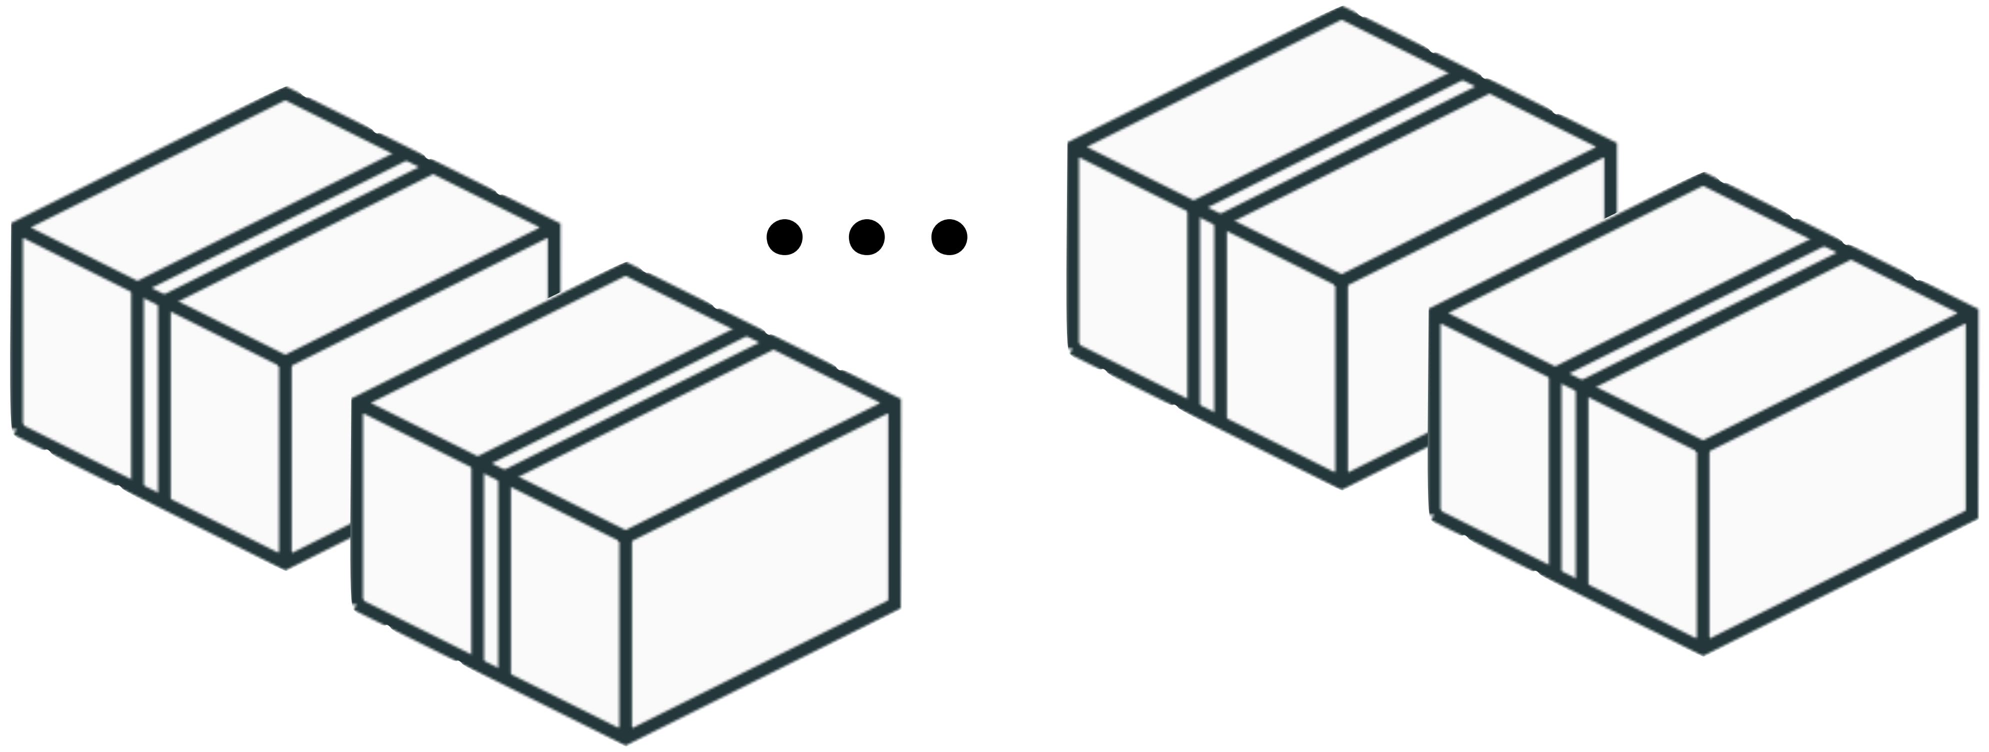

For each pair: $w_{1,i} = 1$ $\quad w_{2,i} = \begin{cases} 0, \text{ w.p. } 1/2 \\ 2, \text{ w.p. } 1/2 \end{cases}$

For $(1-\varepsilon)$ – approx. need to guess max in $> 2/3$ pairs $\rightarrow$ arbitrarily small prob.

Extensions and observation

Extensions and observation

- The reductions hold for any arrival order (random [EHLM '17], free [Yan '11], etc.)

Extensions and observation

- The reductions hold for any arrival order (random [EHLM '17], free [Yan '11], etc.)

- They can be adjusted (with worse constants) to scenarios where we have **a single sample** from each distribution.

Extensions and observation

- The reductions hold for any arrival order (random [EHLM '17], free [Yan '11], etc.)

- They can be adjusted (with worse constants) to scenarios where we have **a single sample** from each distribution.

- We can extend the same techniques up to **XOS weight functions** (again, losing an extra constant factor).

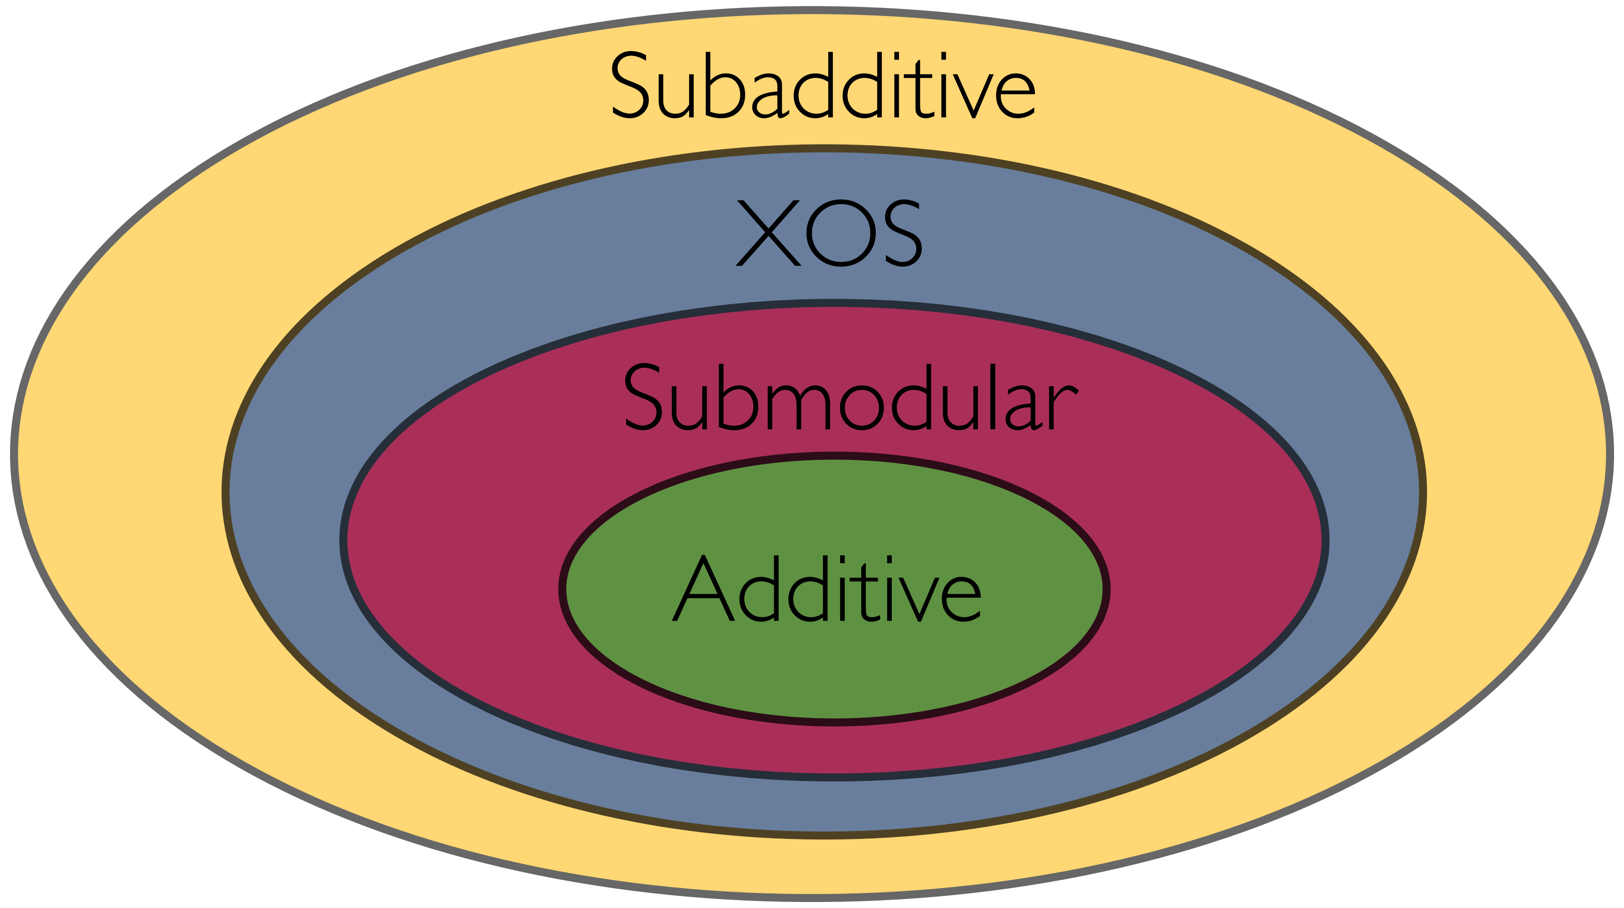

Conclusion & outlook

Conclusion & outlook

- We propose the EoR as a measure of performance for PIs, motivated by *risk-averse* decision-makers.

Conclusion & outlook

- We propose the EoR as a measure of performance for PIs, motivated by *risk-averse* decision-makers.

- For every downward-closed feasibility constraint, arrival order, and XOS weight functions, we establish a two-way blackbox reduction: RoE and EoR are a constant factor apart.

Conclusion & outlook

- We propose the EoR as a measure of performance for PIs, motivated by *risk-averse* decision-makers.

- For every downward-closed feasibility constraint, arrival order, and XOS weight functions, we establish a two-way blackbox reduction: RoE and EoR are a constant factor apart.

Open Questions:

- RoE and EoR have at least a $(2/e)-$ gap. What's the tight factor ?

Conclusion & outlook

- We propose the EoR as a measure of performance for PIs, motivated by *risk-averse* decision-makers.
- For every downward-closed feasibility constraint, arrival order, and XOS weight functions, we establish a two-way blackbox reduction: RoE and EoR are a constant factor apart.

Open Questions:

- RoE and EoR have at least a $(2/e)$ – gap. What's the tight factor ?
- Is it maybe always RoE$(\mathbf{F}) \geq$ EoR$(\mathbf{F})$?

Conclusion & outlook

- We propose the EoR as a measure of performance for PIs, motivated by *risk-averse* decision-makers.

- For every downward-closed feasibility constraint, arrival order, and XOS weight functions, we establish a two-way blackbox reduction: RoE and EoR are a constant factor apart.

Open Questions:

- RoE and EoR have at least a $(2/e)$ – gap. What's the tight factor ?
- Is it maybe always RoE$(\mathbf{F}) \geq$ EoR$(\mathbf{F})$?
- Can we apply similar ideas to online *minimization* problems ?

  Note: [Garg, Gupta, Leonardi, Sankowski '08] briefly discuss EoR for the online Steiner tree.

Conclusion & outlook

- We propose the EoR as a measure of performance for PIs, motivated by *risk-averse* decision-makers.
- For every downward-closed feasibility constraint, arrival order, and XOS weight functions, we establish a two-way blackbox reduction: RoE and EoR are a constant factor apart.

Open Questions:

- RoE and EoR have at least a $(2/e)-$ gap. What's the tight factor ?
- Is it maybe always RoE$(\mathbf{F}) \geq$ EoR$(\mathbf{F})$?
- Can we apply similar ideas to online *minimization* problems ?

  Note: [Garg, Gupta, Leonardi, Sankowski '08] briefly discuss EoR for the online Steiner tree.
- What can we say when we have more samples from each distribution ?

Thank you for your attention!