IPML

IMPRECISE PROBABILISTIC MACHINE LEARNING

Lecture 6: Decision Making under Imprecision

Krikamol Muandet

28 November 2025

Overview

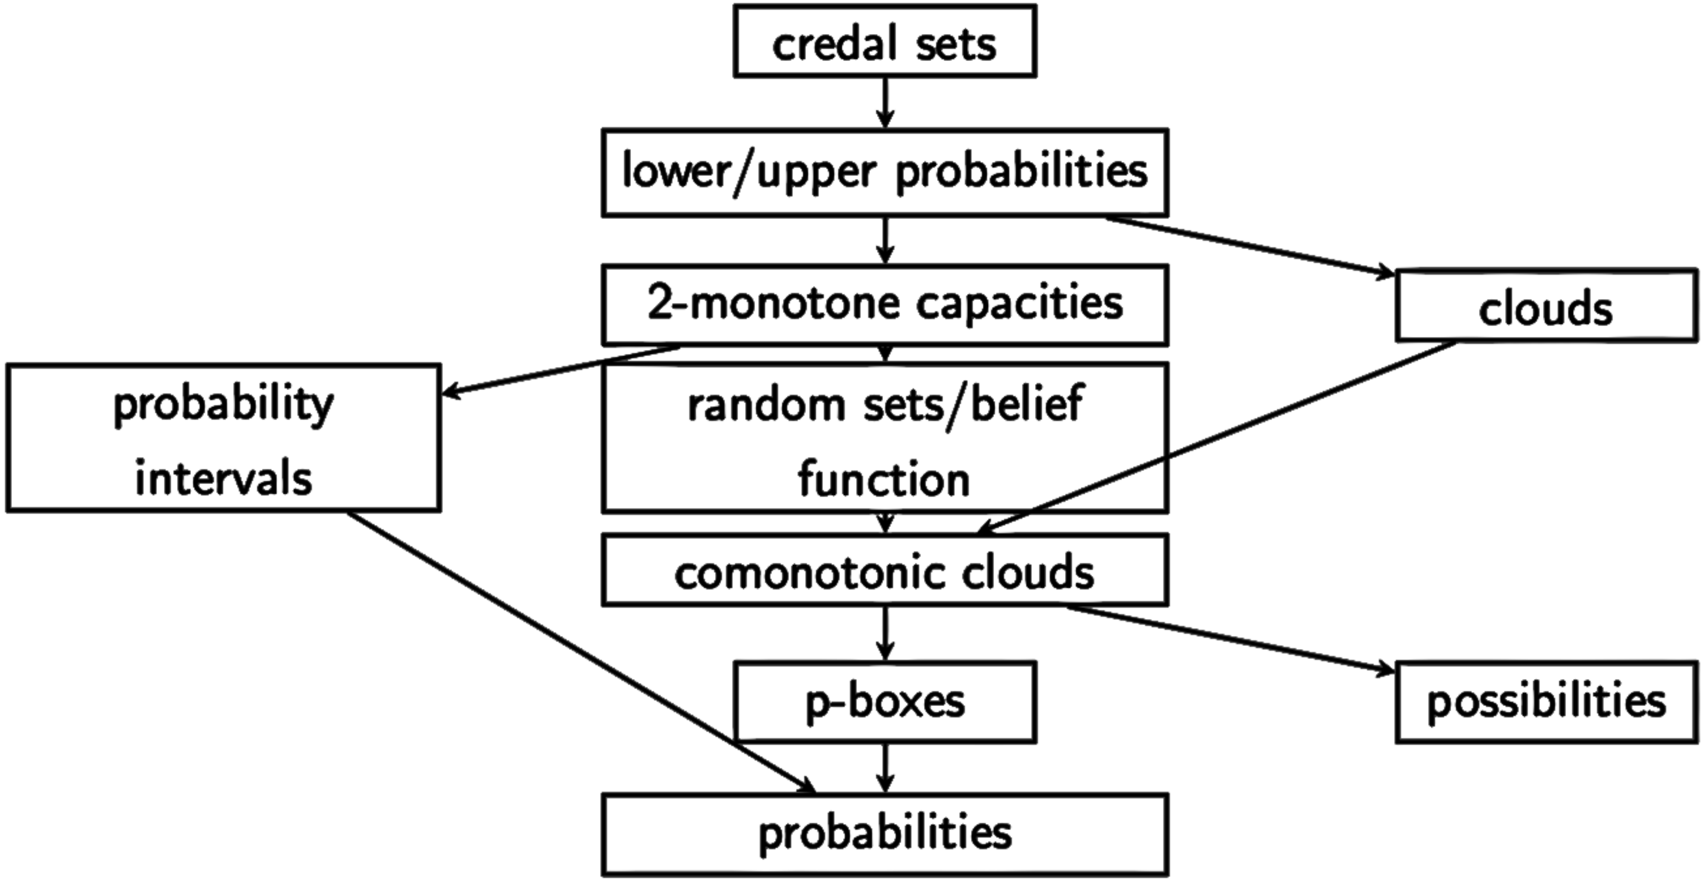
credal sets
lower/upper probabilities
2-monotone capacities
clouds
probability intervals
random sets/belief function
comonotonic clouds
p-boxes
possibilities
probabilities

Outline

1. Decision Making under Uncertainty
2. Imprecise Decision Rules
3. Containment of Rules

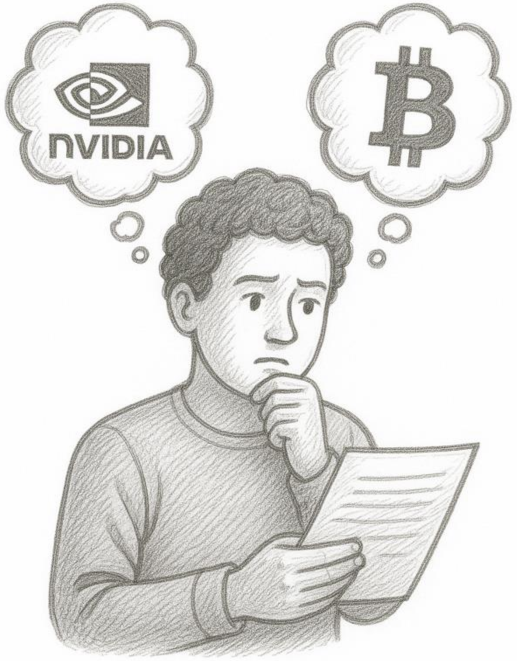

Decision Making under Uncertainty

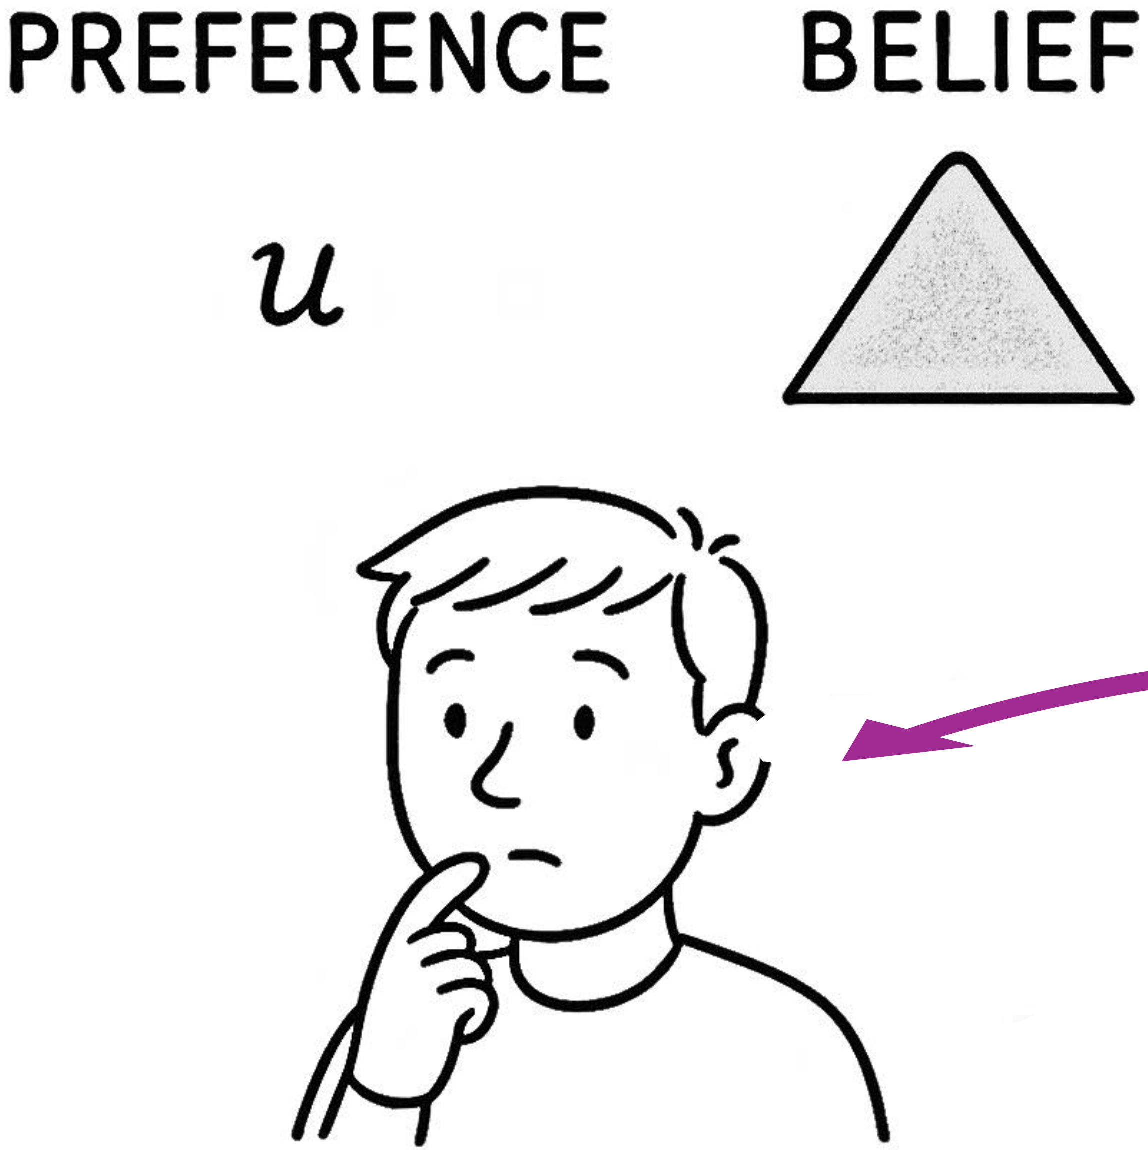

Quasi-Bayesianism

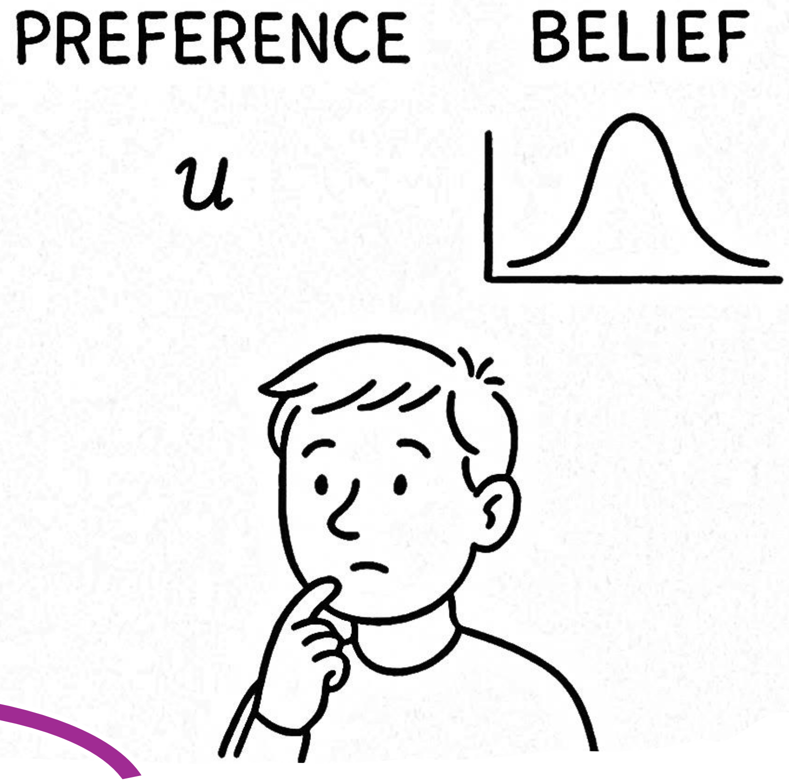

Bayesianism

F. J. Giron and S. Rios, “**Quasi-Bayesian behaviour: A more realistic approach to decision making?**”, In: Bayesian Statistics, University Press, 1980, pp. 17–38.

Quasi-Bayesian Framework

- Let $\prec$ be a partial order of preferences on $\mathcal{A}$ which fulfills the **quasi-Bayesian rationality axioms**.
- Then, there exists a unique nonempty convex set $\mathcal{K}$ of finitely additive probability measures such that

$$a_1 \prec a_2 \Leftrightarrow \mathbb{E}_p[u(\omega, a_1)] < \mathbb{E}_p[u(\omega, a_2)], \forall p \in \mathcal{K}$$

- The set $\mathcal{K}$ is the **credal set** representing $\prec$.
- Some acts are better than others, and some acts cannot be compared.

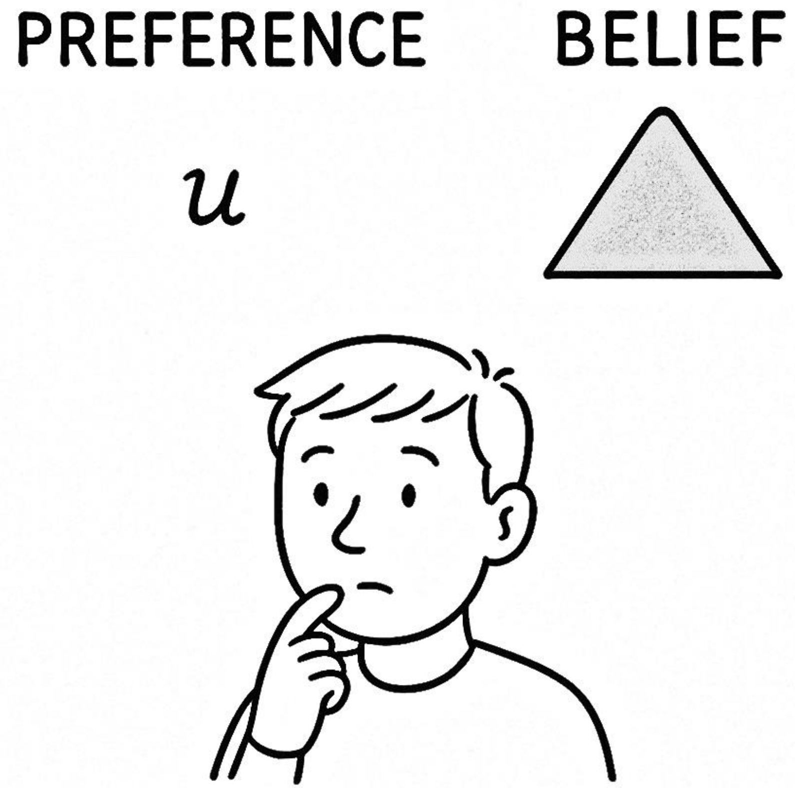

F. J. Giron and S. Rios, "**Quasi-Bayesian behaviour: A more realistic approach to decision making?**", In: Bayesian Statistics, University Press, 1980, pp. 17–38.

Gamble

- A **gamble** is a bounded real-valued function on the state space Ω:

$$f: \Omega \rightarrow R$$

- It represents an **uncertain reward** over the state of affair.

State $\boldsymbol{\omega}$	Sunny	Rainy	Cloudy
$f(\boldsymbol{\omega})$	10	-5	2

Indicator Gamble

- An **indicator gamble** is an indicator function on the state space Ω:

$$f_{\omega}\colon \Omega \rightarrow R, \qquad \omega \in \Omega$$

- The gamble $f_{\omega}\,(\tilde{\omega}) = 1$ when $\omega = \tilde{\omega}$ and 0 otherwise.

State $\boldsymbol{\omega}$	Sunny	Rainy	Cloudy
$f_{\mathrm{sunny}}(\boldsymbol{\omega})$	1	0	0
$f_{\mathrm{rainy}}(\boldsymbol{\omega})$	0	1	0
$f_{\mathrm{cloudy}}(\boldsymbol{\omega})$	0	0	1

De Finetti's Price Functional

- How much would you pay for a gamble f? What is a fair price?

$$P(f_{\text{sunny}}) = \mathbb{E}_{\omega\in\Omega}[f_{\text{sunny}}] = \sum_{\omega\in\Omega} f_{\text{sunny}}(\omega)\cdot \boxed{P(\omega)}$$

$$= P(\text{sunny})$$

Classical Event Probability

State ω	Sunny	Rainy	Cloudy
$f_{\text{sunny}}(\omega)$	1	0	0
$f_{\text{rainy}}(\omega)$	0	1	0
$f_{\text{cloudy}}(\omega)$	0	0	1

Expected Utility

- A **utility function** u_a associated with an action $a \in \mathcal{A}$ is a **gamble** on Ω:

$$u_a\colon \Omega \to R$$

- An expected utility is the *fair price* of this gamble:

$$P(u_a) = \mathbb{E}_{\omega \in \Omega}[u_a(\omega)] = \mathrm{EU}(a)$$

Decision Making

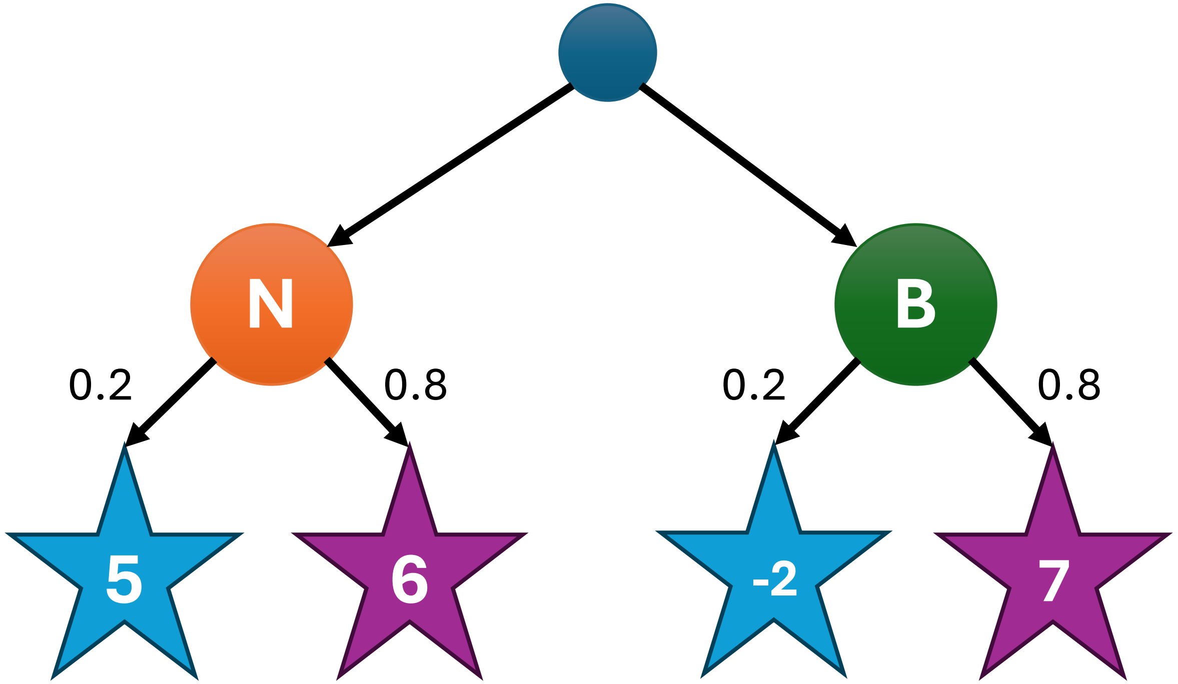

EU(N) = 5.8

EU(B) = 5.4

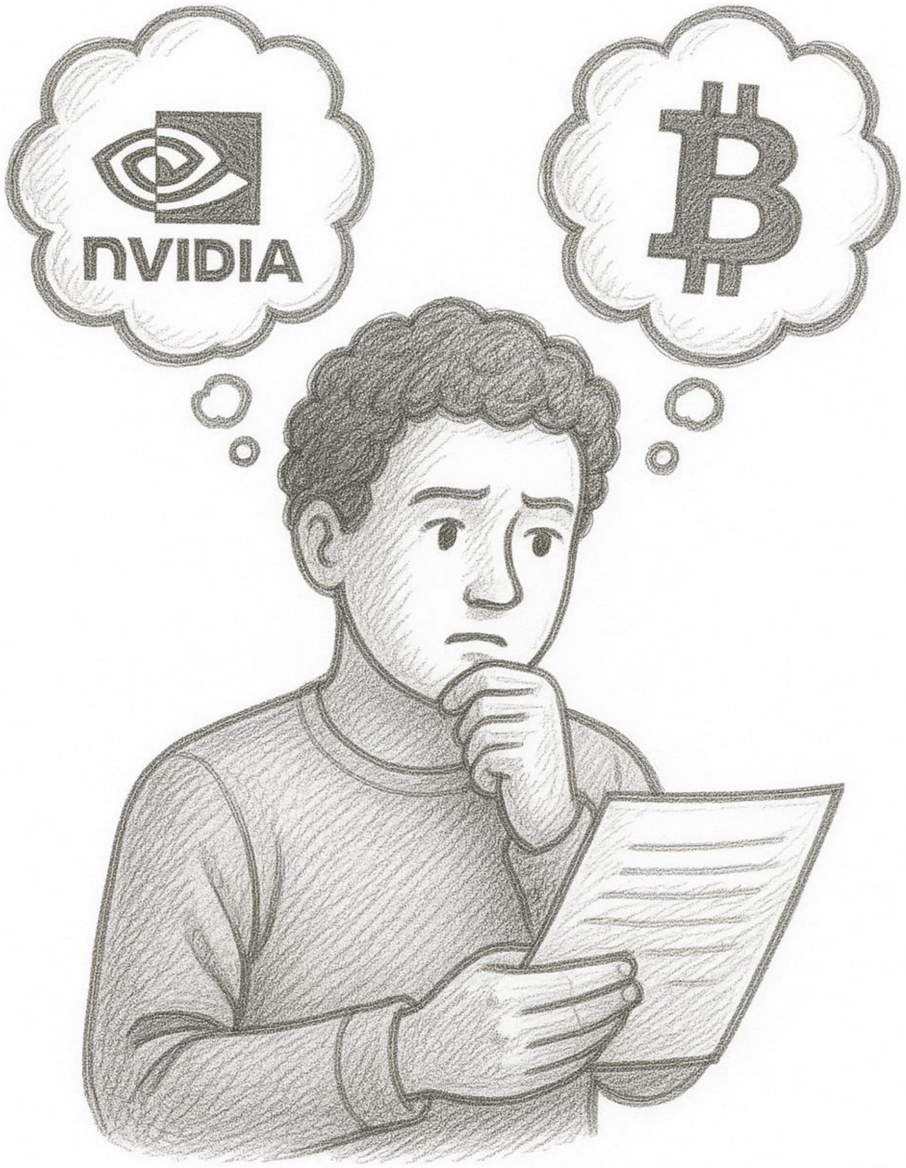

Lower and Upper Expected Utility

- A **utility function** u_a associated with an action $a \in \mathcal{A}$ is a **gamble** on Ω:

$$u_a : \Omega \to R$$

- The **lower and upper expected utilities** for a credal set $\mathcal{K} \subseteq \mathcal{P}(\Omega)$:

$$\underline{\mathrm{EU}}(a) = \min_{P \in \mathcal{K}} \mathbb{E}_{\omega \sim P}[u_a(\omega)]$$

$$\overline{\mathrm{EU}}(a) = \max_{P \in \mathcal{K}} \mathbb{E}_{\omega \sim P}[u_a(\omega)]$$

- How should the agent make decision from $\underline{\mathrm{EU}}(a)$?

Imprecise Decision Rules

Decision Rules for Imprecise Probability

Decision Rule	**Description**
Statewise Admissibility	Remove actions that don’t make sense **regardless of the probability**
Interval Dominance	Remove actions that have **worst best** case than some other’s **worst** case
Γ-maximin	Prepare for the **worst** case
Γ-maximax	Aim for the **best** case
Maximality	Remove actions that are **strictly dominated** in expected utility over the entire credal set.
E-Admissibility	Is it the **best** action **at least in one** case?

Statewise Admissibility

- Remove actions that don't make sense **regardless of the probability**

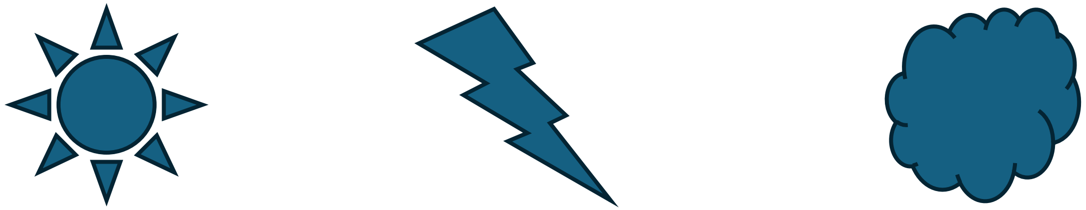

	Sunny (ω_1)	Rainy (ω_2)	Cloudy (ω_3)
a_1	2	2	2
a_2	3	0	1
a_3	1	2	1

Statewise Admissibility

- Remove actions that don't make sense **regardless of the probability**

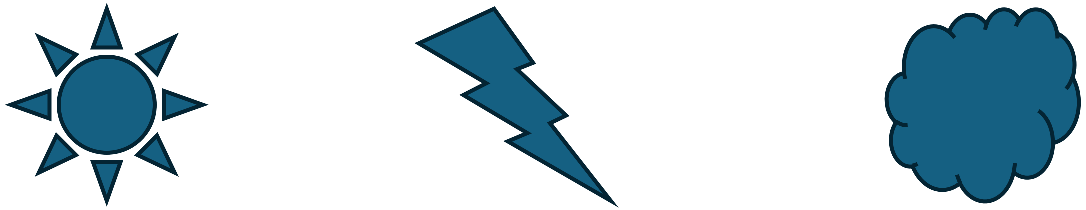

	Sunny (ω_1)	Rainy (ω_2)	Cloudy (ω_3)
a_1	2	2	2
a_2	3	0	1
~~a_3~~	~~1~~	~~2~~	~~1~~

Interval Dominance

- Remove actions that have **worst best** case than some other's **worst** case

	EU_1	EU_2	EU_3	$\underline{\mathrm{EU}}$	$\overline{\mathrm{EU}}$
a_1	3	5	2	$\boxed{2}$	5
~~a_2~~	~~0~~	~~1~~	~~1~~	~~0~~	~~$\boxed{1}$~~
a_3	-1	3	0	-1	3

- Action a_2 has lower $\overline{\mathrm{EU}}$ (best case) than $\underline{\mathrm{EU}}$ (worst case) of action a_1.

Interval Dominance

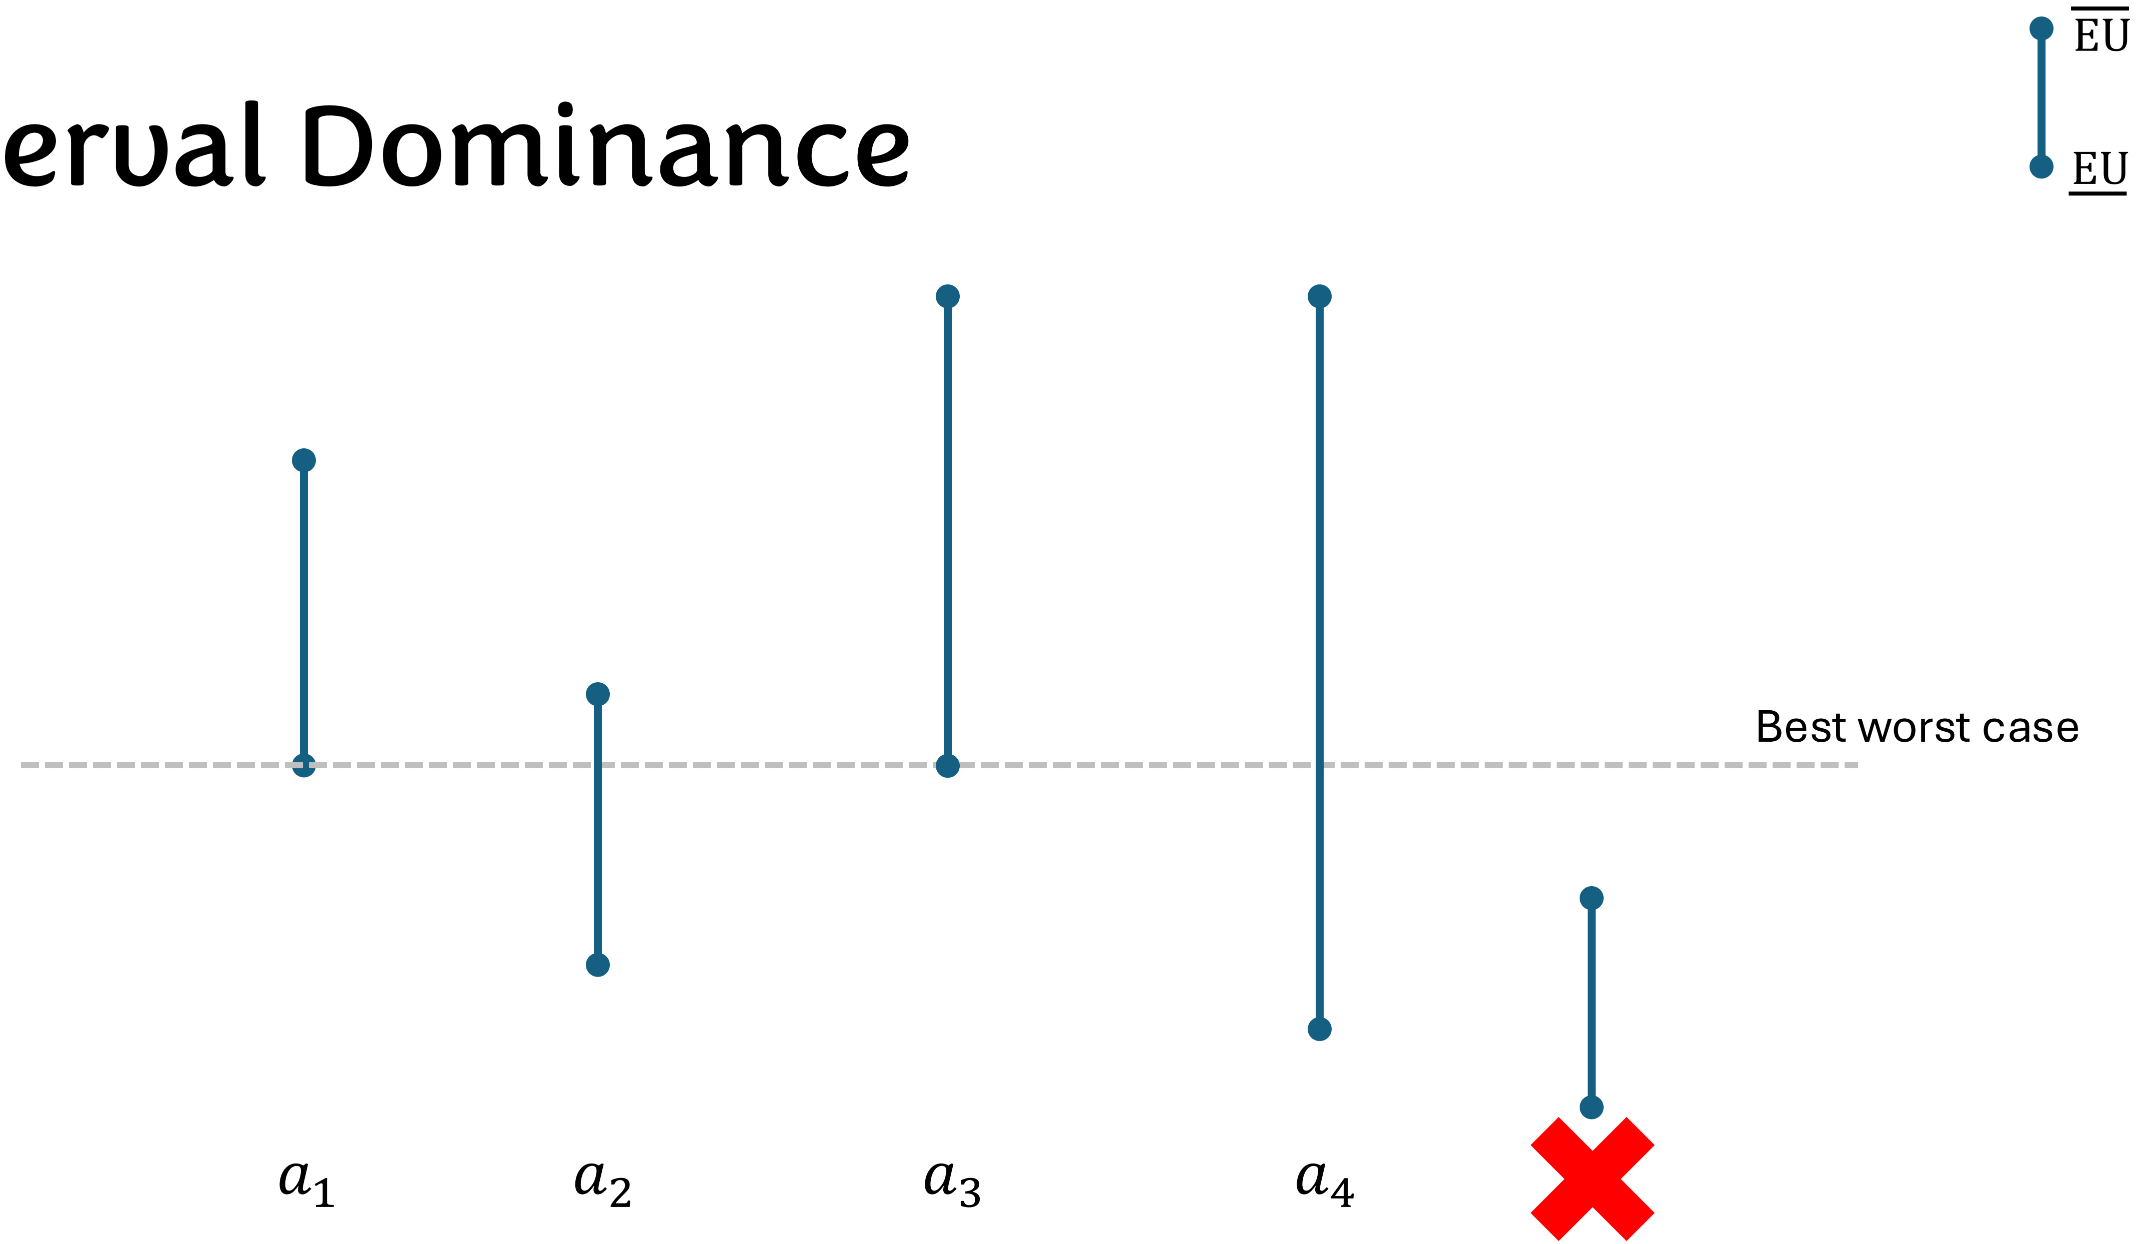

Γ-Maximin

- Prepare for the **worst** case

	EU_1	EU_2	EU_3	$\underline{\text{EU}}$	$\overline{\text{EU}}$
a_1	3	5	2	$\boxed{2}$	5
~~a_2~~	~~0~~	~~1~~	~~2~~	~~0~~	~~2~~
~~a_3~~	~~1~~	~~3~~	~~0~~	~~1~~	~~3~~

- Action a_2 and a_3 have lower $\underline{\text{EU}}$ (worst case) than a_1.

Γ-Maximin

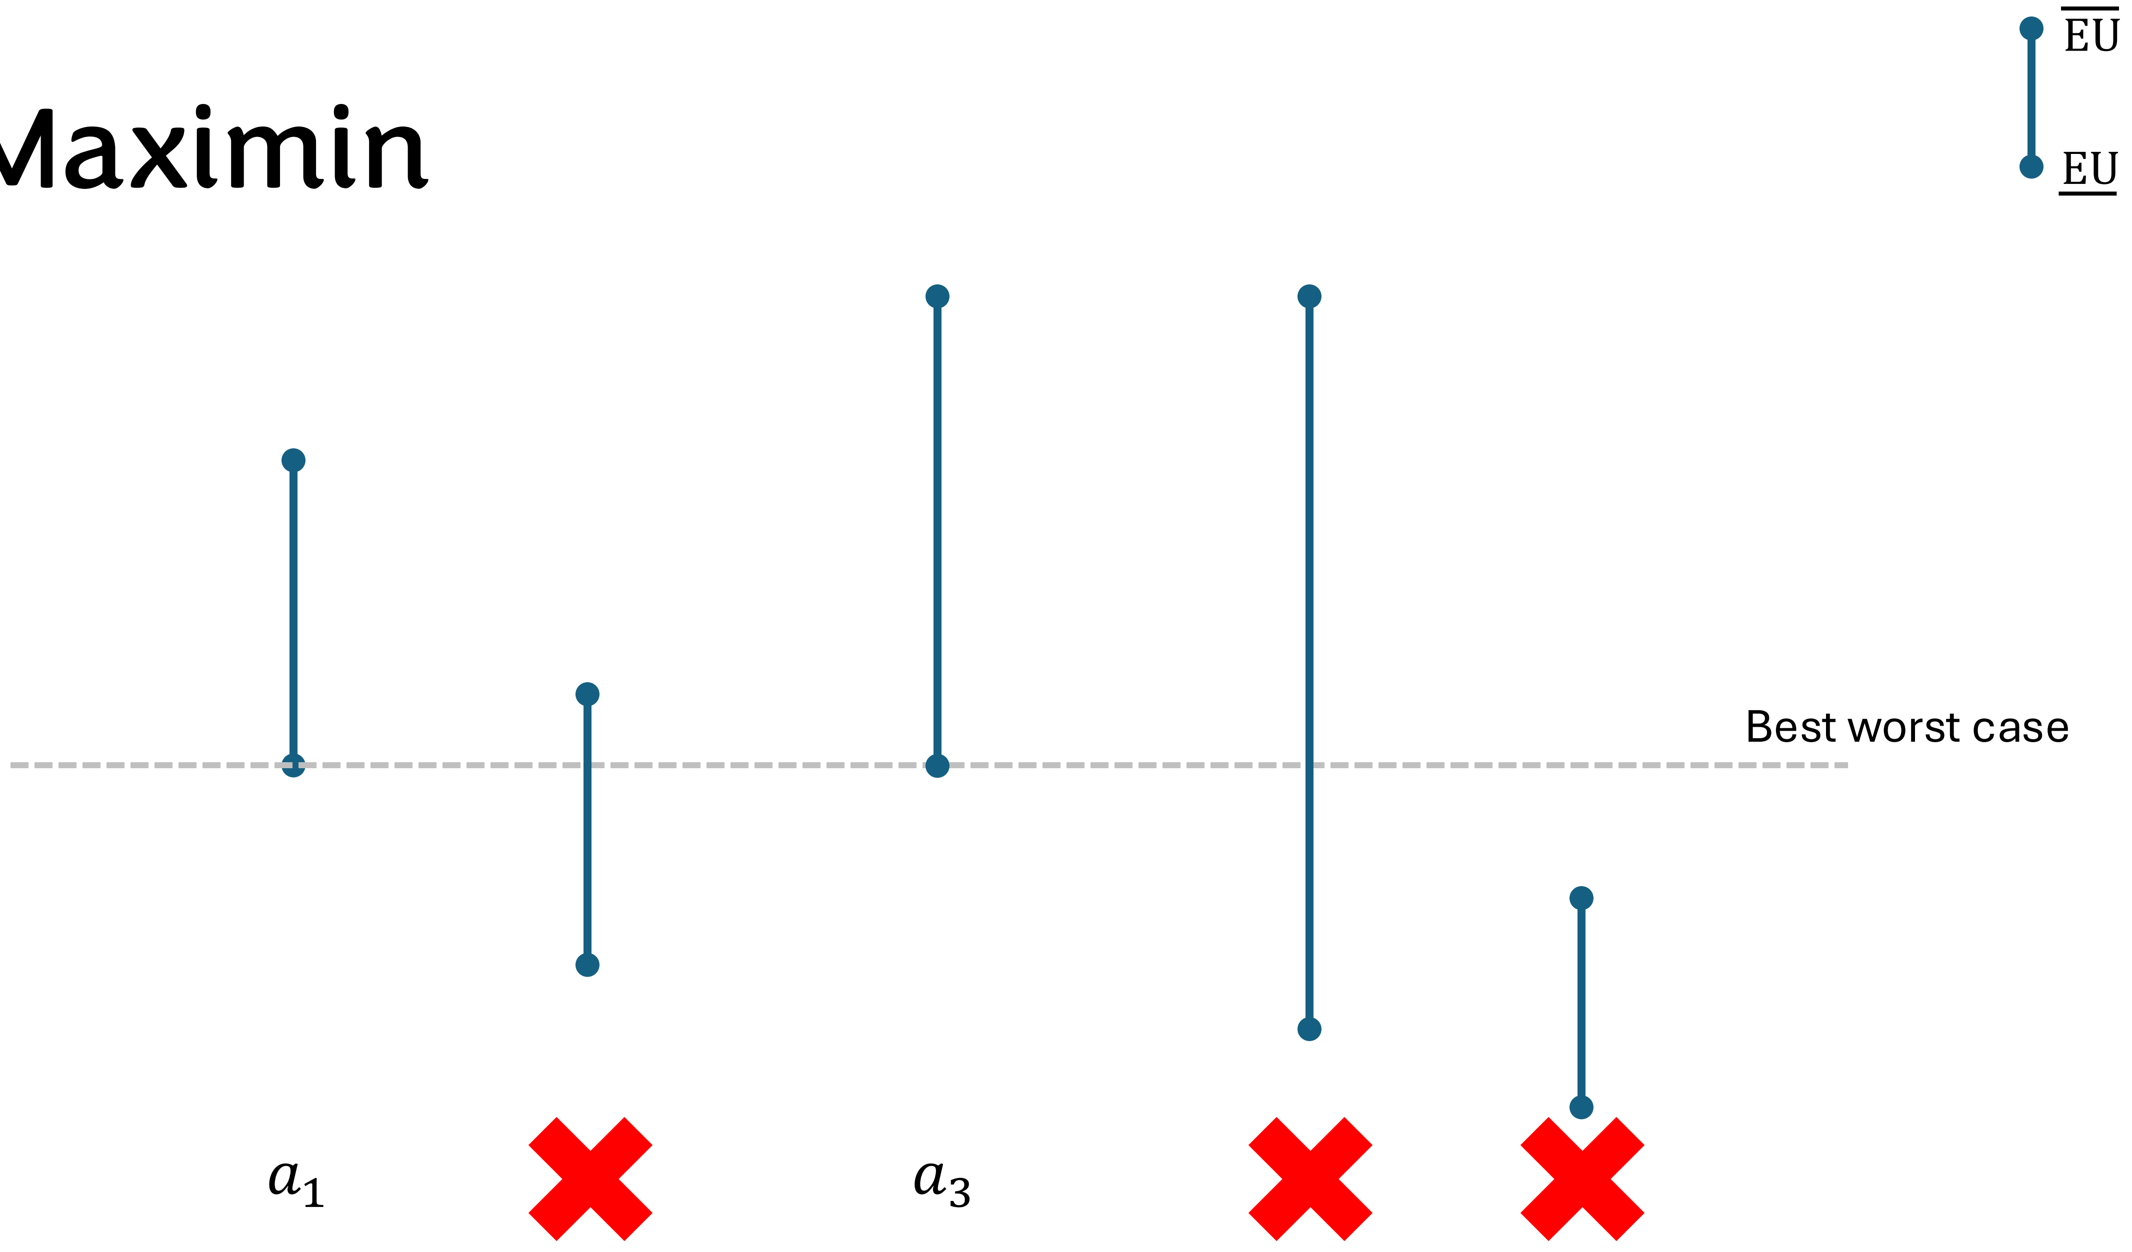

$\overline{\mathrm{EU}}$
$\underline{\mathrm{EU}}$
Best worst case
a_1
a_3

Γ-Maximax

- Aim for the **best** case

	EU_1	EU_2	EU_3	$\underline{\mathrm{EU}}$	$\overline{\mathrm{EU}}$
a_1	3	5	2	2	$\boxed{5}$
~~a_2~~	~~0~~	~~1~~	~~2~~	~~0~~	~~2~~
~~a_3~~	~~1~~	~~3~~	~~0~~	~~1~~	~~3~~

- Action a_2 and a_3 have lower $\overline{\mathrm{EU}}$ (best case) than a_1.

Γ-Maximin

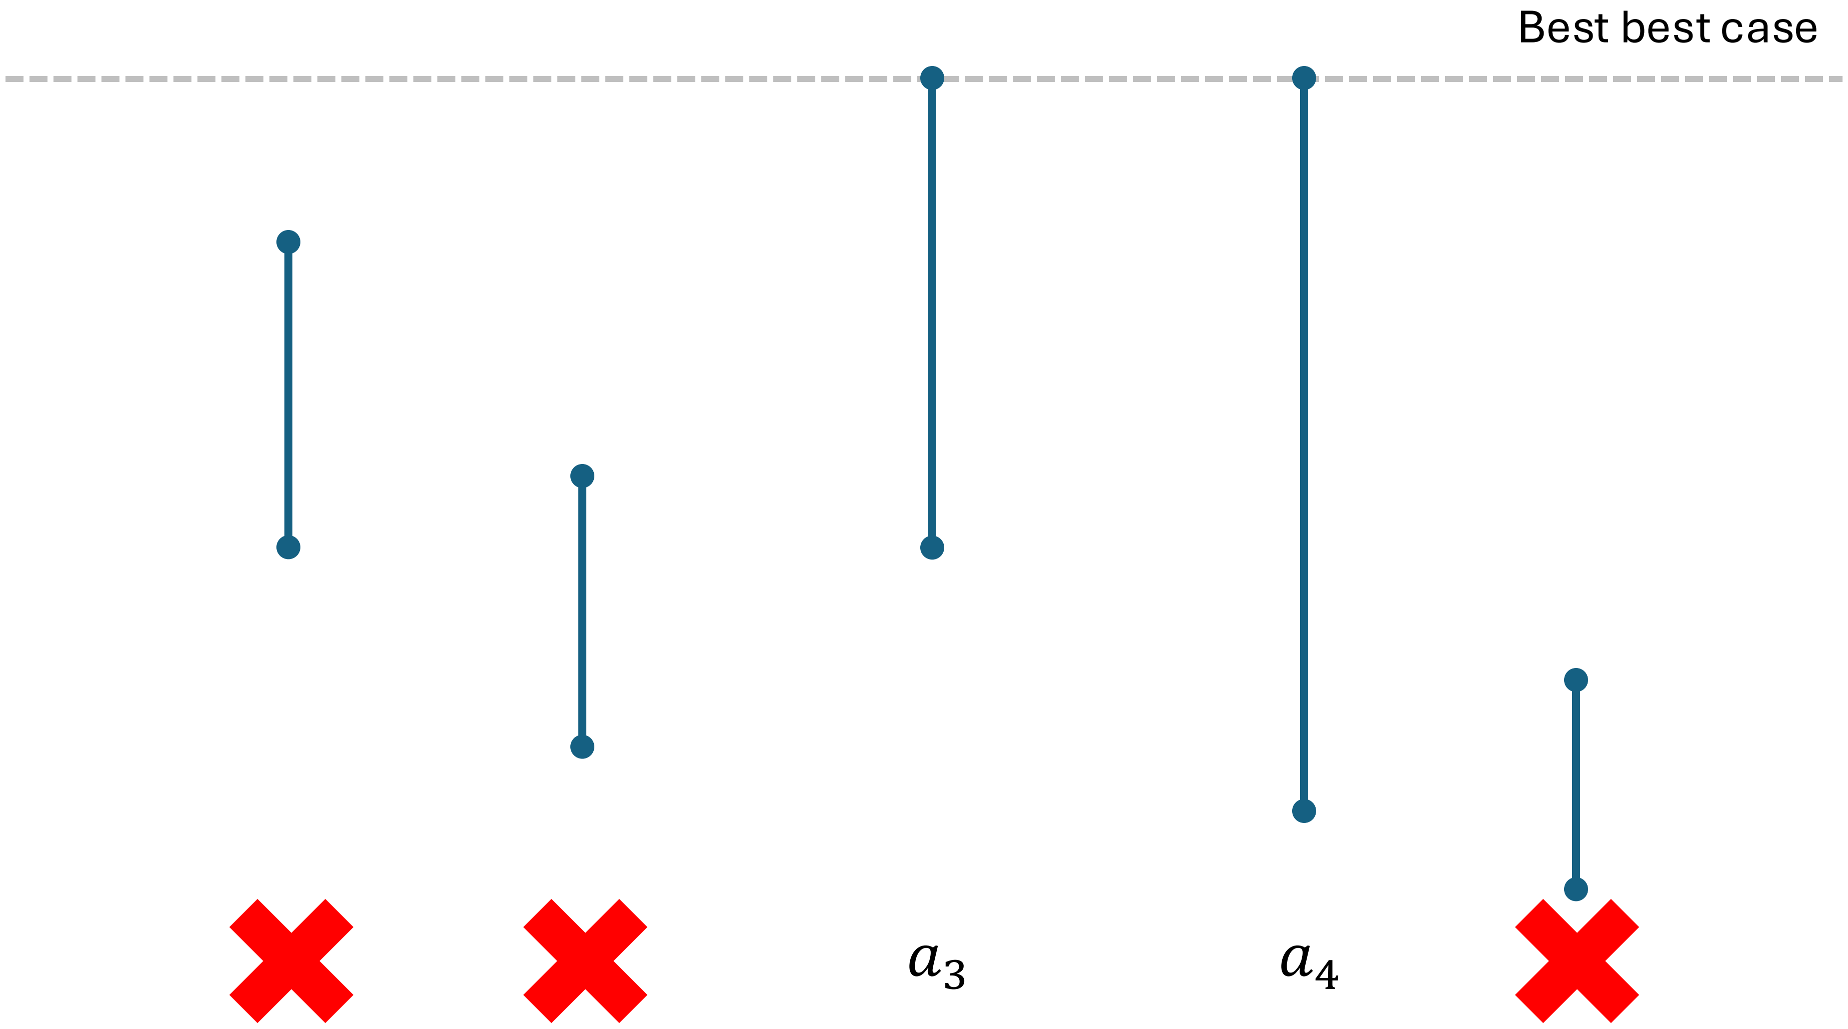

Maximality

- Remove actions that are **strictly dominated** in expected utility over the entire credal set.

	EU_1	EU_2	EU_3	EU_4	EU_5
a_1	3	5	2	2	5
a_2	-1	3	0	5	4
a_3	0	4	1	0	2

Maximality

- Remove actions that are **strictly dominated** in expected utility over the entire credal set.

	EU_1	EU_2	EU_3	EU_4	EU_5
a_1	3	5	2	2	5
a_2	-1	3	0	5	4
a_3	0	4	1	0	2

- Nothing to eliminate!

Maximality

- Remove actions that are **strictly dominated** in expected utility over the entire credal set.

	EU_1	EU_2	EU_3	EU_4	EU_5
a_1	3	5	2	2	5
a_2	-1	3	0	5	4
a_3	0	4	1	0	2

- Nothing to eliminate!

Maximality

- Remove actions that are **strictly dominated** in expected utility over the entire credal set.

	EU_1	EU_2	EU_3	EU_4	EU_5
a_1	3	5	2	2	5
a_2	-1	3	0	5	4
~~a_3~~	~~0~~	~~4~~	~~1~~	~~0~~	~~2~~

- Action a_1 is always better. Remove a_3.

E-Admissibility

- Is it the **best** action **at least in one** case?

	EU_1	EU_2	EU_3	EU_4	EU_5
a_1	3	-1	4	5	-4
a_2	2	2	3	2	2
a_3	1	3	-3	0	3

E-Admissibility

- Is it the **best** action **at least in one** case?

	EU_1	EU_2	EU_3	EU_4	EU_5
a_1	3	-1	4	5	-4
~~a_2~~	~~2~~	~~2~~	~~3~~	~~2~~	~~2~~
a_3	1	3	-3	0	3

- Action a_2 is never the best.

Computational Complexity

$$\min\{r^{\mathrm{T}}p : p \in \mathcal{K}\} \qquad \max\{r^{\mathrm{T}}p : p \in \mathcal{K}\}$$

Decision Rule	LPs required	Description
Statewise Admissibility	0	Simple vector domination test
Interval Dominance	0	Array comparison only
Γ-maximin / Γ-maximax	0	Pick extrema from the two arrays above
Maximality	$\leq \lvert\mathrm{D}\rvert(\lvert\mathrm{D}\rvert - 1)$	One $\min(c_{\mathrm{d}} - c_{\mathrm{e}})^{\mathrm{T}}p$ per ordered pair
E-Admissibility	$\leq \lvert\mathrm{D}\rvert$	One feasibility LP per surviving act
Lower / upper EU bounds	$2\lvert\mathrm{D}\rvert$	One min and one max per action

Containment of Rules

Existence Guarantees

- We are always **filtering** the decisions that are **worse**.
- For a finite set of actions $\mathcal{A}$, we can always find **at least one** **winner**.

Containment of Rules

1. Remove **statewise inadmissible** actions
2. Calculate $\underline{\mathrm{EU}}$ and $\overline{\mathrm{EU}}$ for all actions
   $2|\mathrm{D}|$ LP shared by later rules
3. (Optional) Read Γ**-maximax** and Γ**-maximin**
4. Read **interval dominance** by comparing $\underline{\mathrm{EU}}$ and $\overline{\mathrm{EU}}$s
5. Solve reward difference LP to obtain **maximality**
   $\leq |\mathrm{D}| \cdot (|\mathrm{D}| - 1)$ LPs
6. Solve a feasibility LP to obtain the E-admissible set
   $\leq |\mathrm{D}|$ feasibility LPs

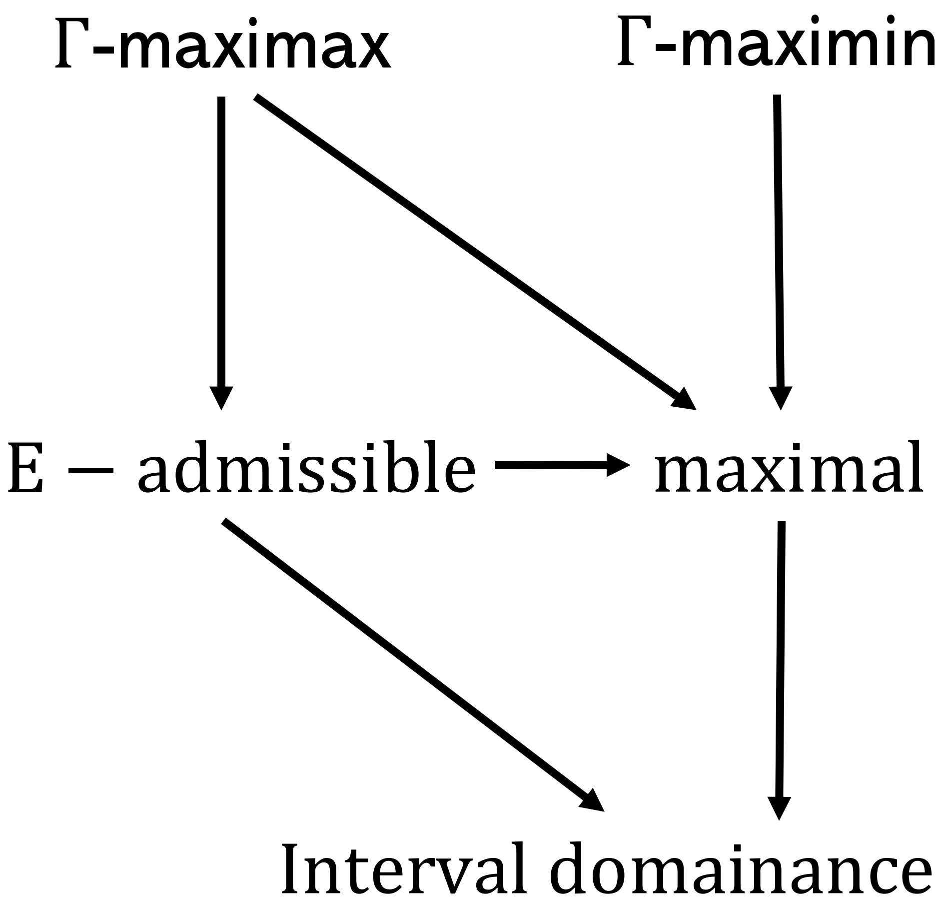

Recommended Reading

- Decision Making under Uncertainty using Imprecise Probabilities by Matthias C. M. Troffaes
- SIPTA School 2024: Decisions by Matthias C.M. Troffaes

This lecture is based in part on the presentation of **Bartlomiej Pogodzinski** submitted for the Imprecise Probabilistic Machine Learning (IPML) seminar at Saarland University.