Towards Long-term Preference Elicitation

Paolo Viappiani, CNRS and University Pais Dauphine PSL

DA2PL 2026 - Brussels (Belgium)

Motivation

- Preference elicitation (PE) schemes typically assume the decision problem to be fully defined
- In dynamic or open-ended decision environments, however, these assumptions are not valid:
  - the sets of options may evolve over time, or even be unknown at the outset
  - different similar decisions need to be made in different points in time,
  - the context may change
- The need for reasoning in terms of PE with longer horizon of application

Key challenge: elicit preferences without knowing the precise future decision problems.

Classical approach to elicitation vs Adaptive Elicitation

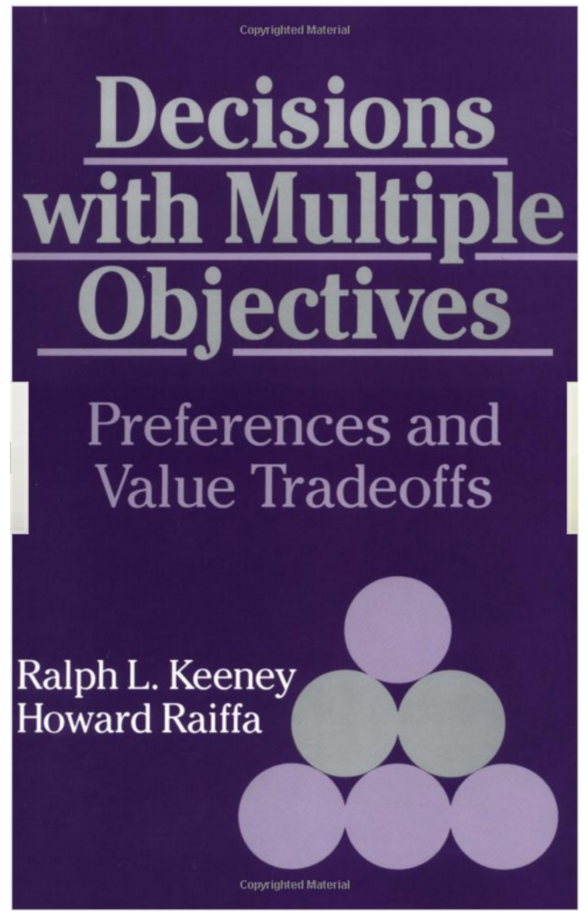

- Classic approach to elicitation
  - Learn the parameters of the utility function precisely
  - Typically long list of (often difficult) questions
  - Focus on high stakes decisions

- Interactive preference elicitation
  - Focus on what is feasible: do not need to know the preference model precisely to make a recommendation
  - Ask « simple » questions (comparison queries, ranking queries)
  - Ask preference questions in an *adaptive* way

Preference Elicitation Loop

- An **incremental elicitation** procedure asks questions to the user at each step.
  - Recommendation problem formulated as decision-making under *preference uncertainty.*
  - Responses gradually reduce uncertainty about the utility function, enabling the selection of a near-optimal choice at the end of the procedure.

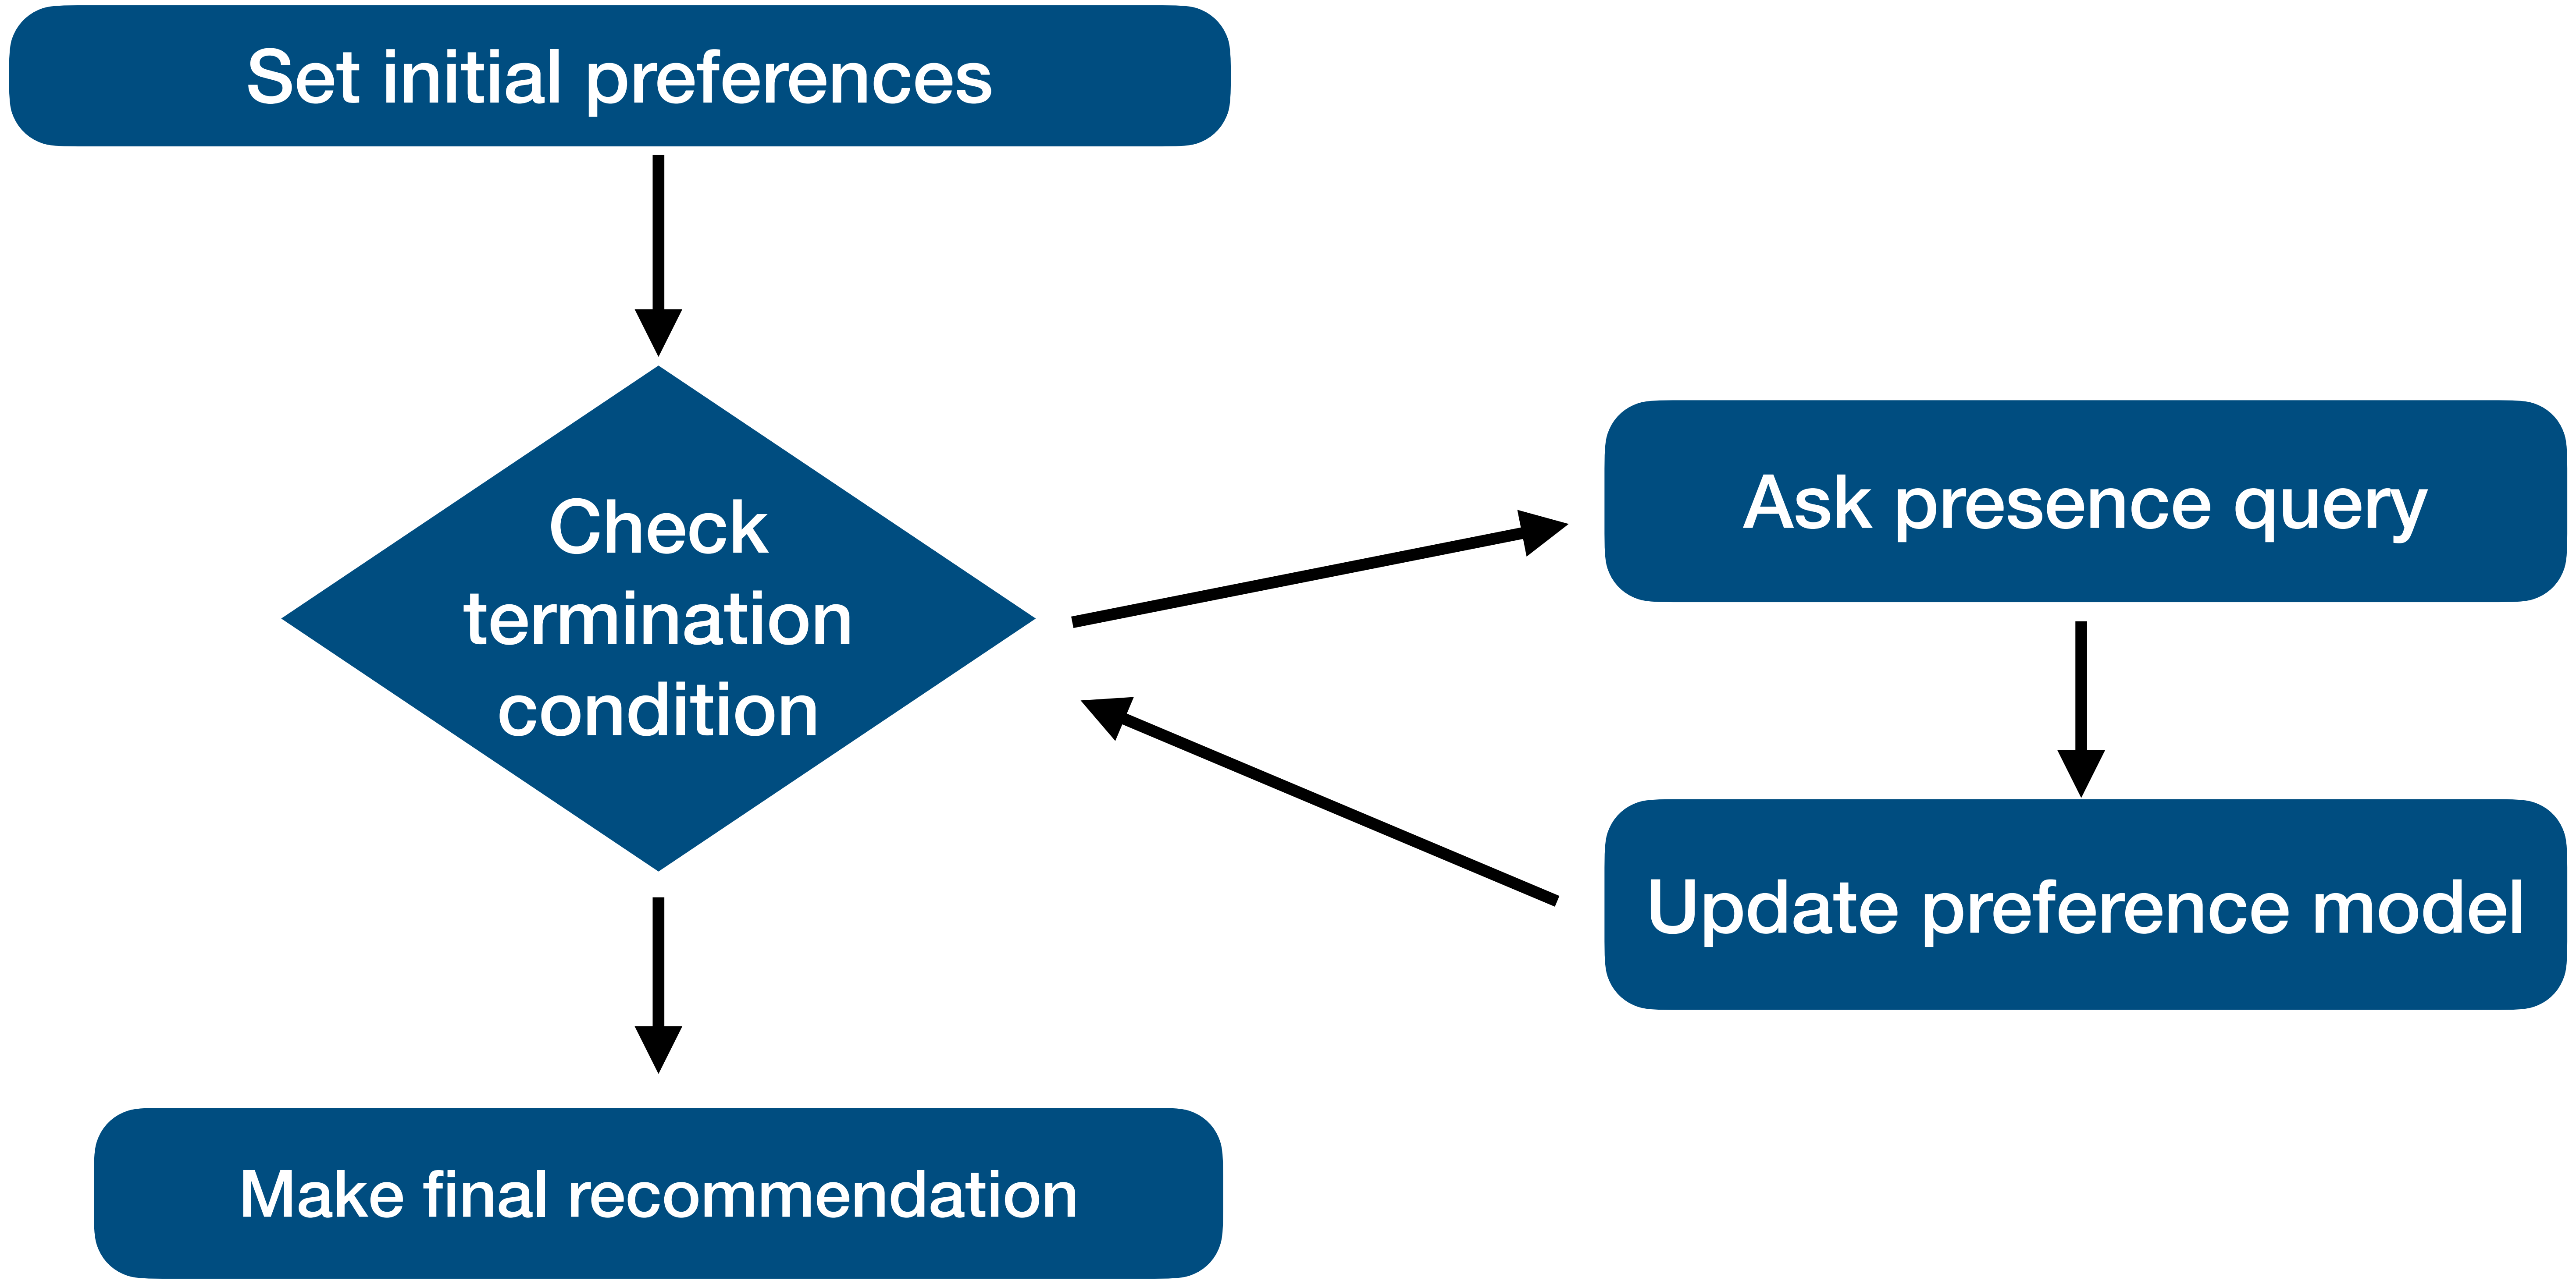

Models for preference uncertainty and elicitation

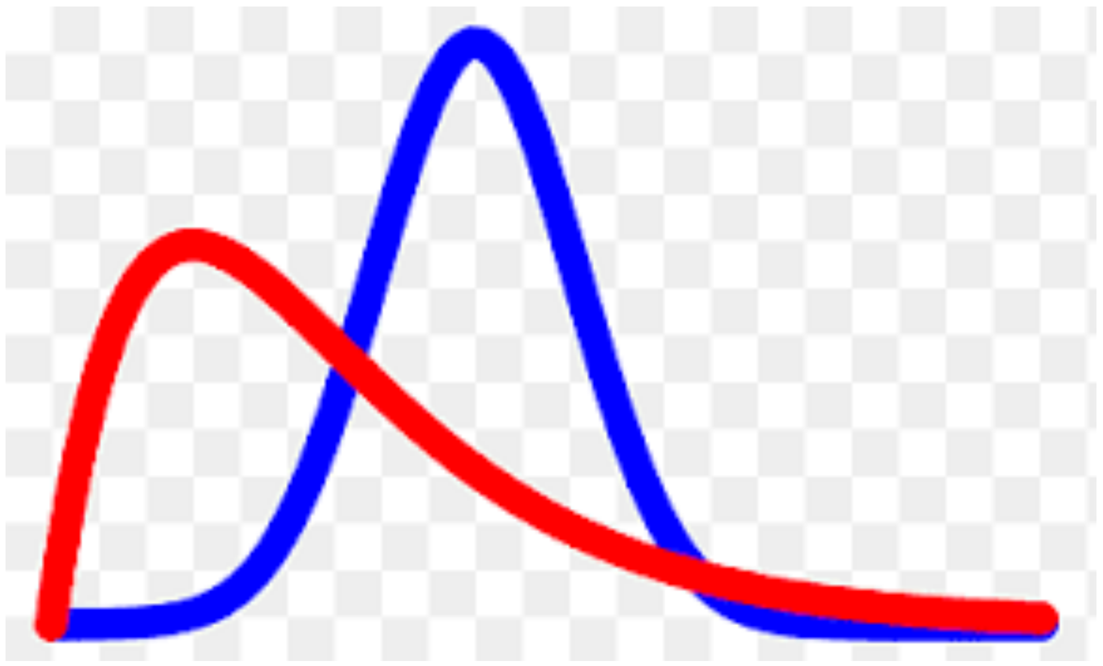

- **Bayesian** approaches: a distribution represents the subjective belief about the user's utility function

  + can model user mistake (answer « noise »)

  + can exploit prior information

  + query strategies to ask informative queries

  - computationally expensive (for Bayesian updates, for making recommendation, choosing the next query)

  - priors might not available, arbitrary choices of probabilistic model

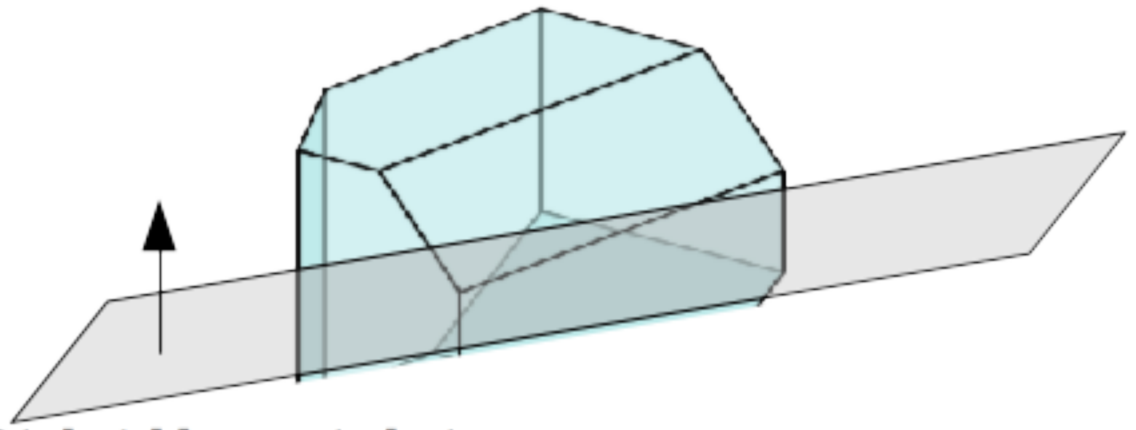

- **Regret-based** approach based on **strict uncertainty:** only a feasible region of utility function parameters is available

  + give robust recommendations

  + usually faster than Bayesian elicitation

  + query strategies to ask informative queries

  - cannot handle user mistakes

  - computation may still expensive when dataset is large

Short term versus long term elicitation

Long-term elicitation

- Preference information is acquired at the onset but decisions occur later
  - PE could be used by chatbots and personal assistants
  - Middle ground between classic elicitation and adaptive one
- Key research questions
  - How to adapt existing elicitation approaches to long-term scenarios?
  - What are the most efficient elicitation strategies for long-term use?
  - Can we derive a policy for recommendation?

Motivating example

- *E-commerce website:* ask a few preference questions when the user registers the first time.

- In *Spring*: choose a book on Preference Learning
  - Ask a few additional preference questions before recommendation

- In *Summer*: book to read at the beach
  - Ask a few additional preference questions before recommendation

- At *Christmas*: book as a presented for a loved one
  - Ask a few additional preference questions before recommendation

Recommendation and Elicitation under Uncertain Availability

- Assume a set of alternatives A
- Consider different stages *1,…,n*
  - At stage i, a subset S of A is revealed
  - The user has to choose an alternative in S
- $P_i(S)$ is the probability that set S is available at stage i
  - Dependence on i can model *temporal* aspects
  - Simplified model: each alternative x has a fixed probability p_x of being available
  - Even simpler: each alternative has the same fixed probability p
- Motivated by the unavailable candidate model in social choice by *Lu and Boutilier*

The Bayesian View

- Probabilities represent the **subjective belief** about the user's utility function
  - **P(u):** distribution over possible utility function
  - **P(r|u):** response model dictating likelihood of responses to queries
  - We use **expected utility** to sort the alternatives: x* is the one with highest EU
  - **Expected loss** can be used as *stopping criterion* $\int_{x \in A} [\max u(x) - u(x^*)] P(u) du$
- Different response models consistent with ***Luce axiom of choice:***

$$P(S \to x) = P(S \to R) P(R \to x)$$

  - In particular, in the **Logistic model** choice probability follows Boltzmann distribution

$$P(S \to x) = \frac{e^{\gamma u(x)}}{\sum_{y \in S} e^{\gamma u(y)}}$$

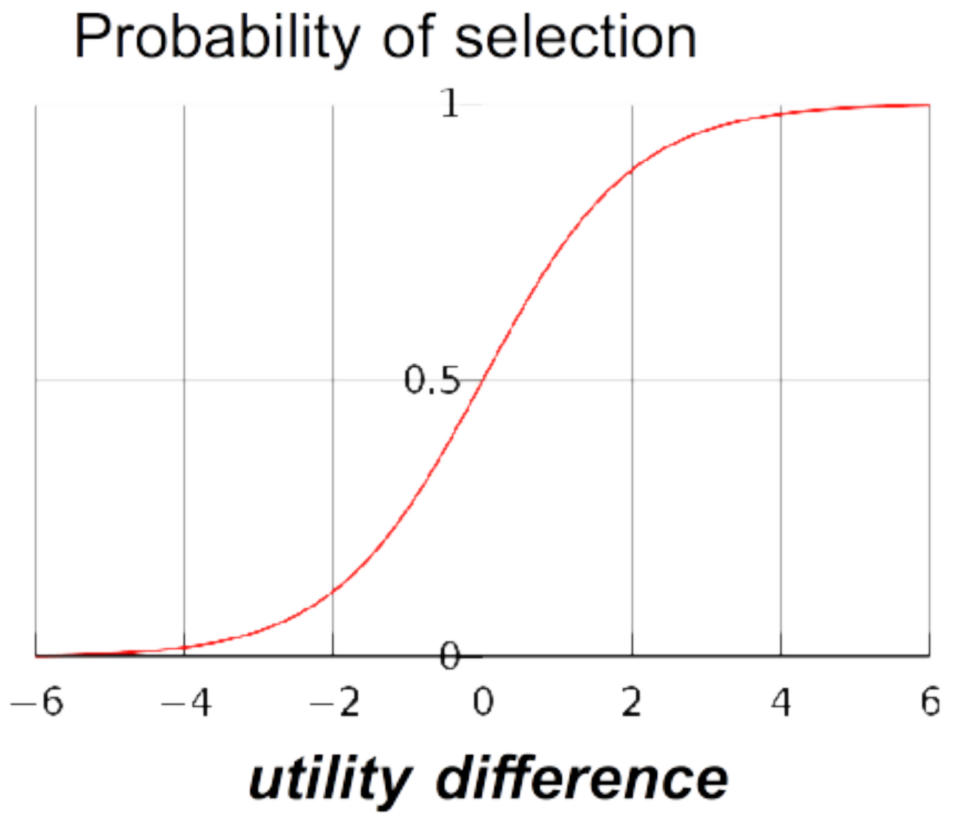

What Query to Ask Next ?

- Ask query with highest Expected Value of Information (EVOI):

$$EVOI(q) = \sum_{r \in R_q} P(r) EU^*_{u \sim P(u|r)} - EU^*_{u \sim P(u)}$$

where R_q is the set of possible answers

- Link between **optimal recommendation sets** and **optimal query sets**
  - *We provided* efficient computation methods for computing myopically choice set
  - Greedy methods with guarantees using **submodularity**

« Do you prefer apple or banana? »

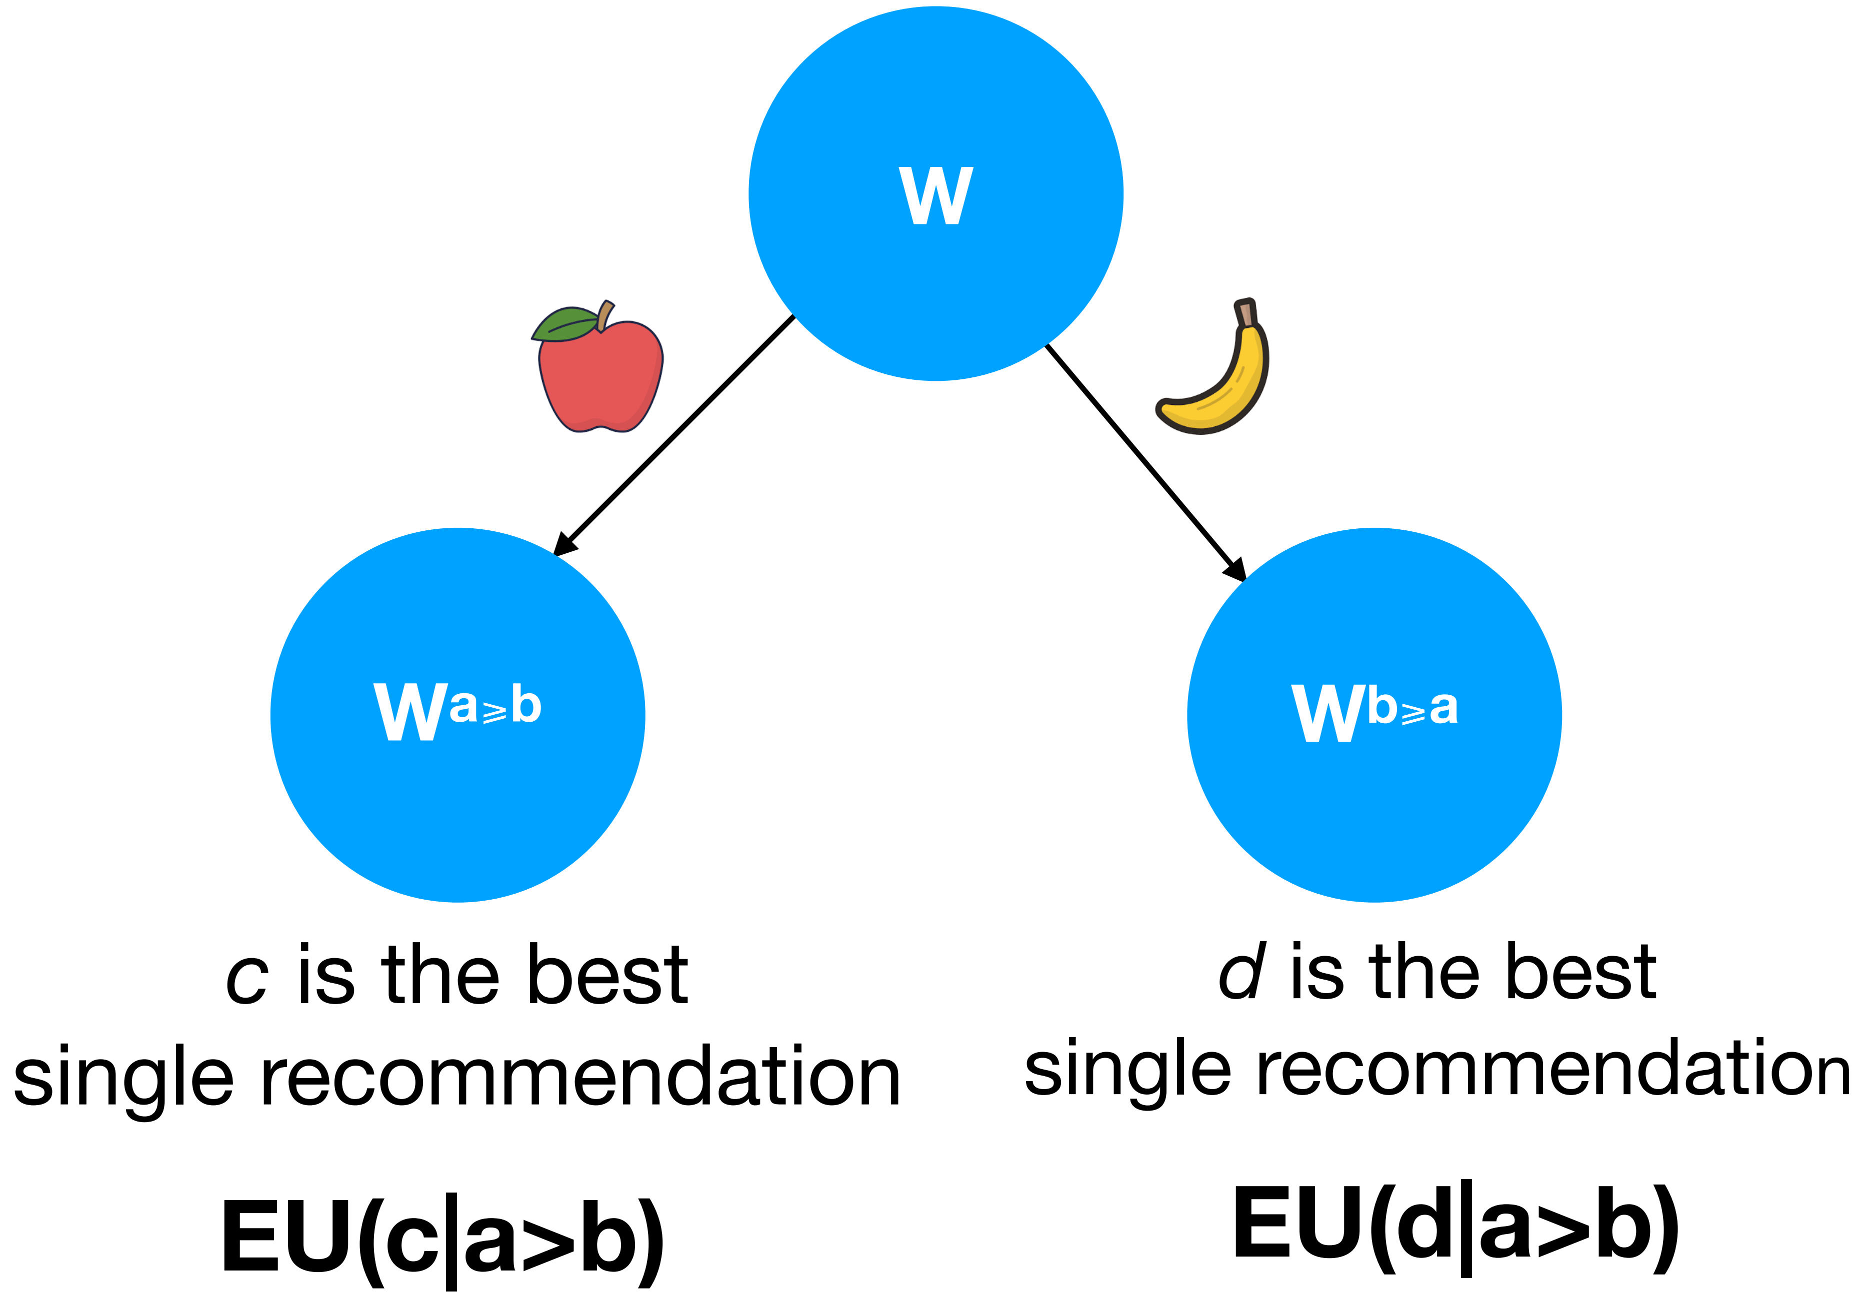

Paolo Viappiani, Craig Boutilier. On the equivalence of optimal recommendation sets and myopically optimal query sets. Artificial Intelligence 2020.

Bayesian approach under uncertain availability

- Formulas in expectation over *both* the uncertain utility and availability
- **Expected Expected Utility (EEU):** *to evaluate how « good » we are doing*

$$EEU = \sum_{S \subseteq A} P(S) \max_{x \in S} EU(x)$$

- **Expected Loss:** *to decide when to stop*

$$\ell := \sum_{S \subseteq A} P(S) \int_U [\max_{y \in S} u(y) - u(x)] p(u) du$$

- **Expected VoI:** *to evaluate the goodness of asking « do you prefer a or b ? »*

$$p_{a \succ b} EEU_{u|a \succ b} + p_{b \succ a} EEU_{u|b \succ a} - EEU^*$$

EEU Computation

- Computation: challenging, currently use sampling in a « smart » way

- Particular case where computation is efficient: each alternative has probability **p** iid of being available
  - Compute expected utility EU for each alternative
  - Sort $x_1,\ldots,x_n$ with respect to EU: $EU(x_1) \geq \ldots \geq EU(x_n)$

$$EEU = \sum_{i=1}^{n} p(1-p)^{i-1} EU(x_i)$$

Ranking Policy

- Sort alternatives according to expected utility: the resulting ranking is recommendation **policy** under uncertain availability
  - Given **a > b > c**
  - If a is available, recommend a
  - otherwise, if b is available, recommend b
  - Otherwise, recommend c
- **Contingency plan** for recommendation under uncertain availability
  - the alternative with the highest EU (the top in the ranking) among those in S is returned.
- In the Bayesian model, we do not loose anything by restricting to a ranking policy

Reasoning with Strict Uncertainty

- The weight space *W* is the set of weights that satisfy all preference statement
- Preference inequalities are often obtained from answers to *preference questions*
  - For example: *(8, 0, 3)* > *(0, 7, 0)*
  - From *a* at least as good as *b* we infer $(a - b) \cdot w \geq 0$

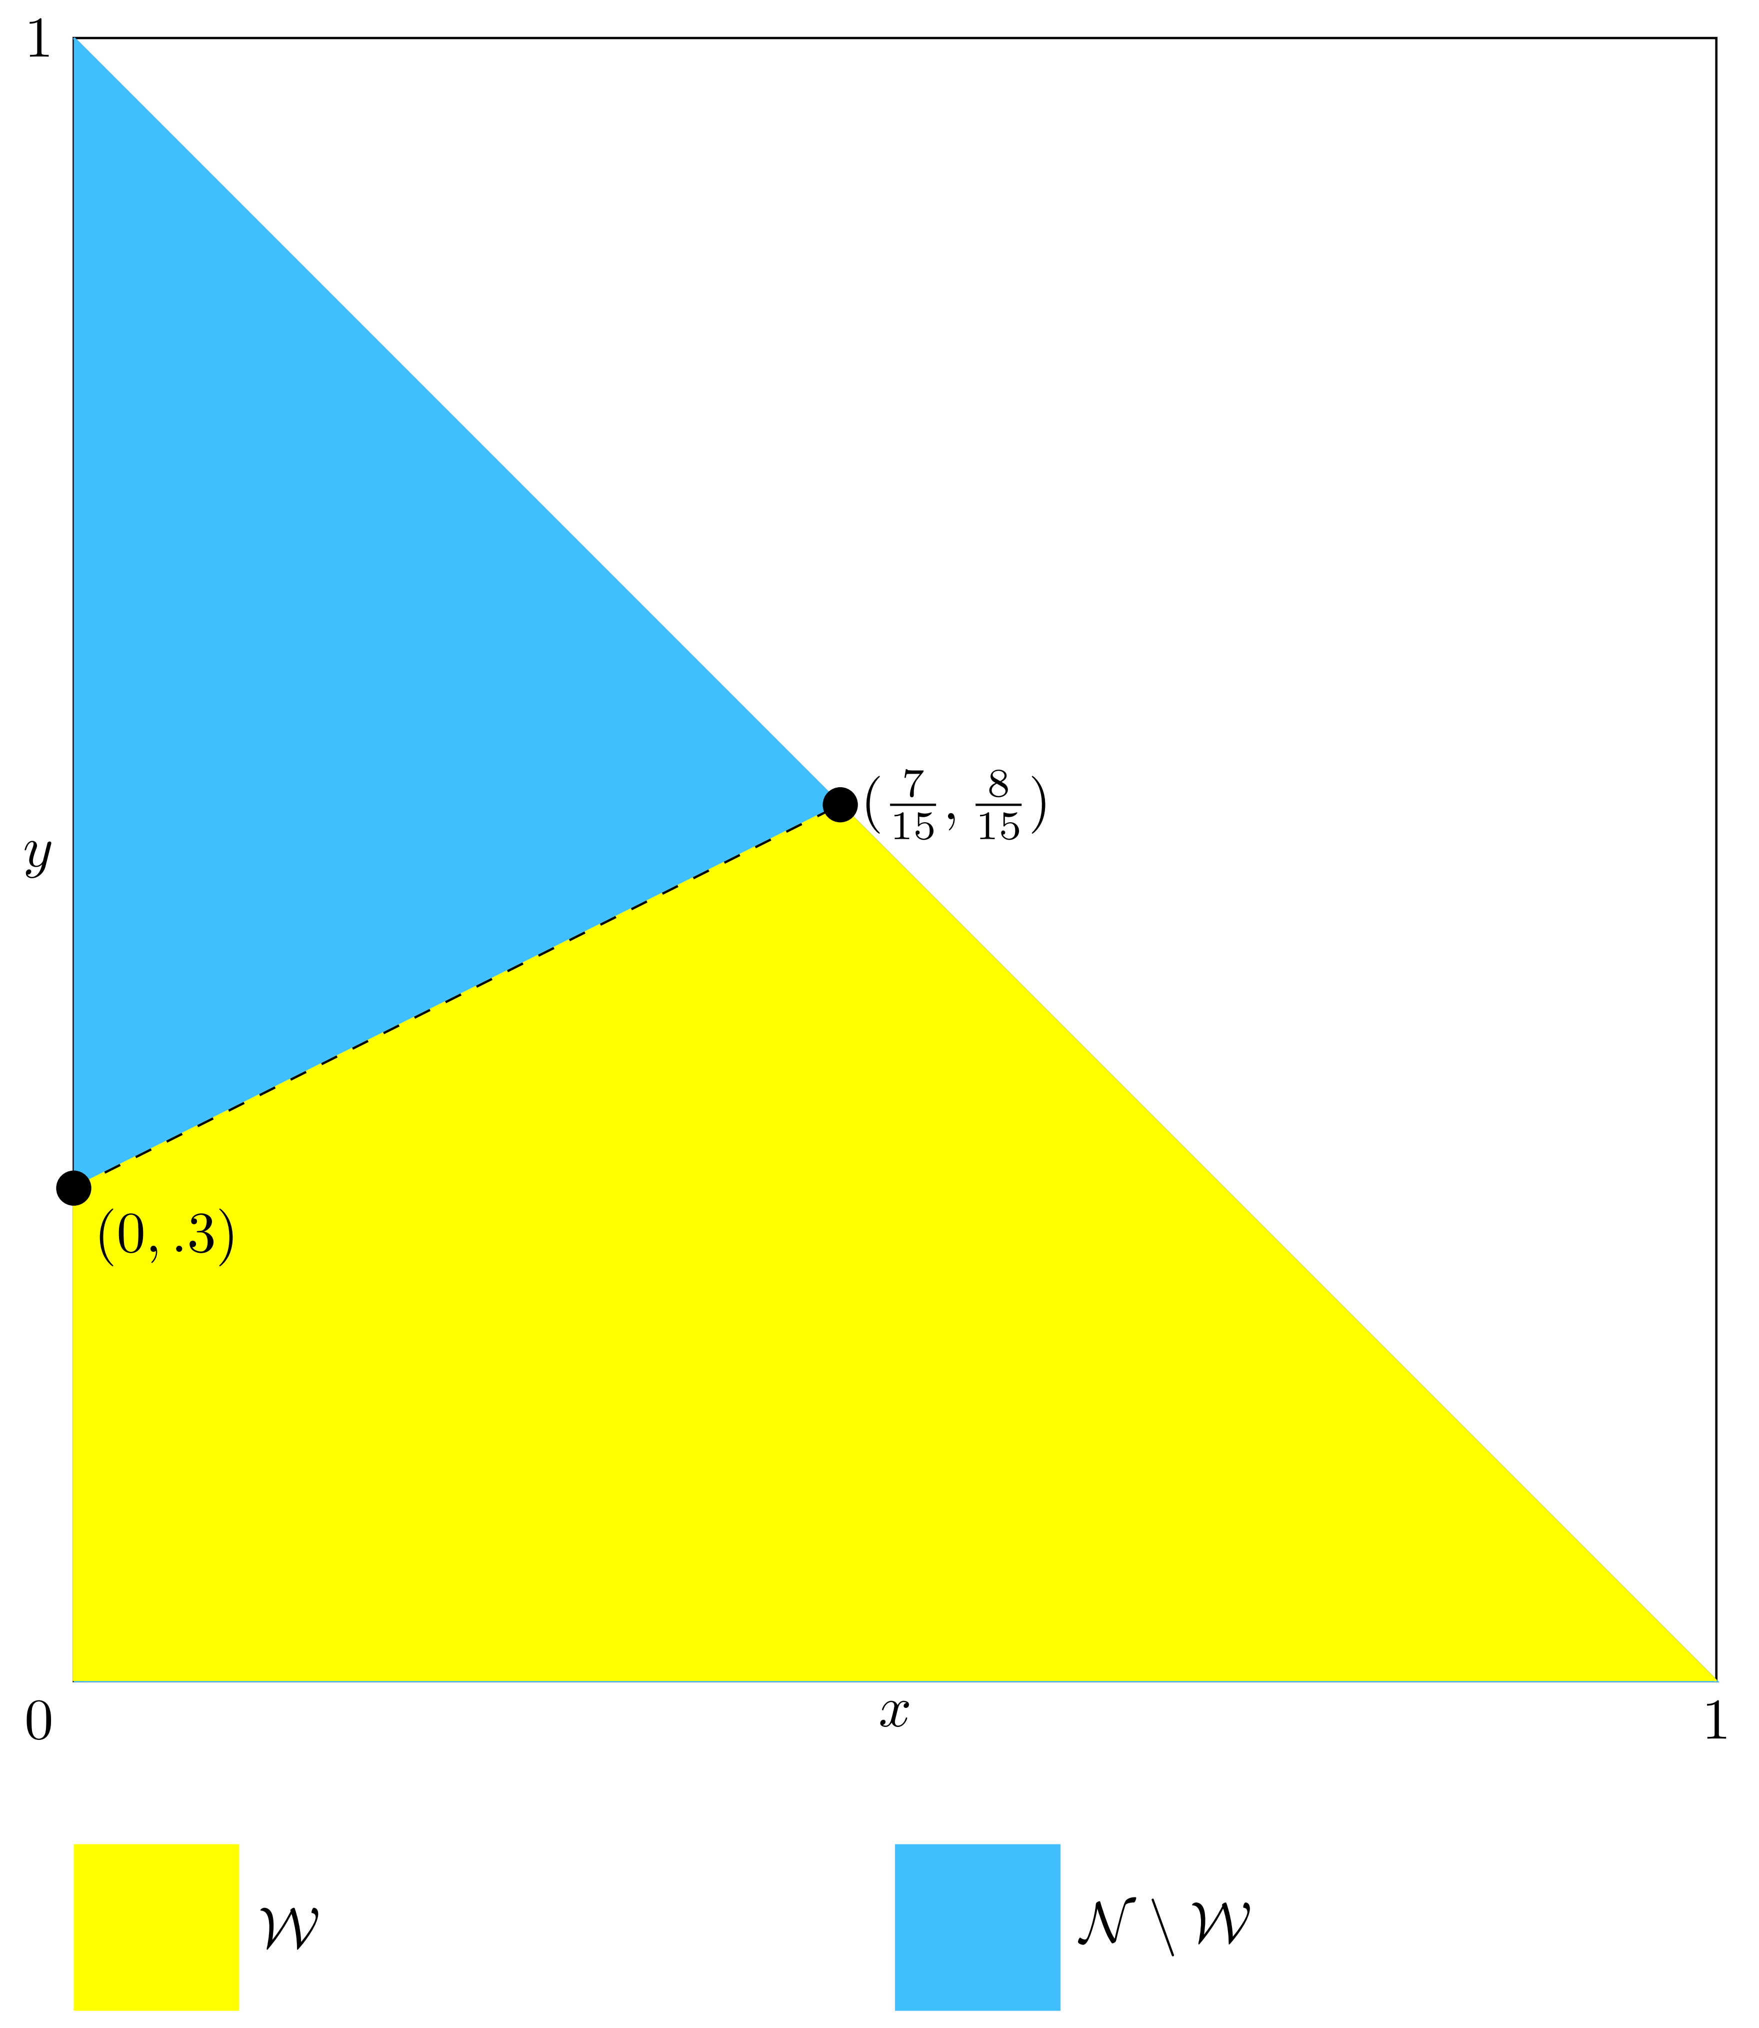

Minimax Regret

- *Assumptions:* a set of feasible utility parameters *W* is given; a set of items (or choices) A
- Max regret is the maximum loss associated with x *when the adversary chooses the utility function*

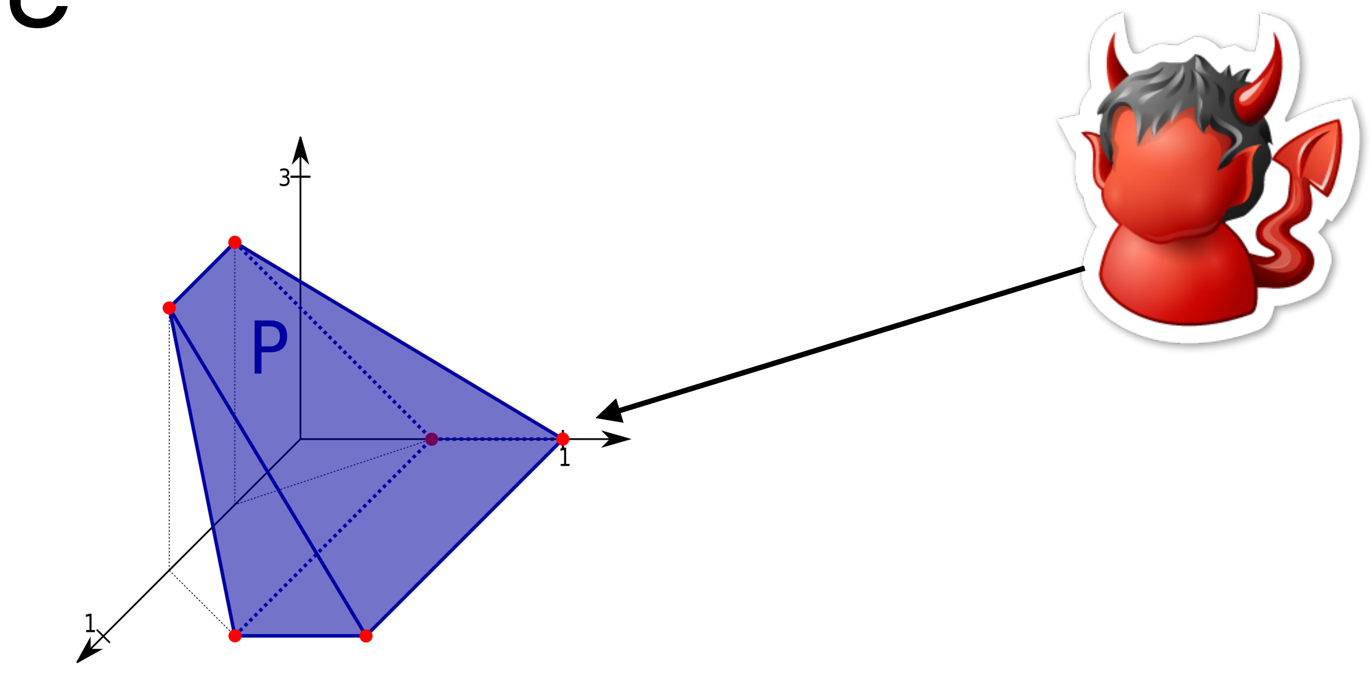

$$MR(x, A; W) = \max_{y \in A} \max_{w \in W} u_w(y) - u_w(x) =$$

$$\max_{w \in W} [\max_{y \in A} u_w(y)] - u_w(x).$$

- The regret-optimal recommendation is the one associated with minimax regret

$$MMR(W) = \min_{x \in A} MR(x, A; W)$$

Expected Minimax Regret

- In the case of uncertain availability, we consider

$$EMMR = \sum_{S \subseteq A} P(S) \min_{x \in S} MR(x, S; U) .$$

- Depending on the hypothesis on P , the computation of the expected minimax regret may be simplified.
- Queries can be generated
  - according to an analogue of value of information
  - Heuristics accounting for probability of being available and max regret
- There is no compact « ranking policy » (as shown in next slide)

IIA Violation by Minimax Regret

- Assume 3 objects and 2 possible utility functions

	u_1	u_2
a	4	2
b	2	3
c	0	5

PMR	**a**	**b**	**c**	***MR***
a	-	1	3	***3***
b	2	-	2	***2***
c	4	3	-	***4***

- Now assume c is not available: a is now the minimax optimal alternative

	u_1	u_2
a	4	2
b	2	3
[redacted]		

PMR	**a**	**b**	[redacted]	**MR**
a	-	1		**1**
b	2	-		**2**
[redacted]				

Conclusions

- PE is a well established subfield, with different approaches (Bayesian, regret,…)
  - Most works focus on eliciting preferences for a single given choice problem
- Main idea: consider the case where elicitation takes place in advance, before actually facing a decision problem
  - Shift from a short-term focus to a multi-decision perspective
  - Dynamic environments
- Ongoing works
  - Full development of the idea: make computation more efficient, study different assumptions for P
  - Consider not just varying the set of alternatives but also the context
  - Experimental validation

adt2026.sciencesconf.org

Login

MAIN MENU

- Home
- Call for papers
- Committees

HELP

- @ Contact

ADT 2026 - 9TH INTERNATIONAL CONFERENCE ON ALGORITHMIC DECISION THEORY

The 9th International Conference on Algorithmic Decision Theory (ADT 2026) will be held at **University Paris Dauphine-PSL** on **November 16-18, 2026**

Established in 2009, the ADT conferences usually take place every two years with the aim of gathering researchers interested in the algorithmic aspects of

IMPORTANT DATES

Abstract deadline: May 4, 2026

Paper deadline: May 11, 2026

Notification: July 13, 2026

Final version: August 11, 2026

ADT conference: November 16-18, 2026