A Single Dice Roll to Satisfy All Goals: Randomisation Requirements for Strategies in Multi-Objective Markov Decision Processes

James C. A. Main Mickaël Randour

F.R.S.-FNRS and UMONS – Université de Mons, Belgium

Highlights of Logic, Games and Automata 2024

Random strategies and multiple objectives

- We study **Markov decision processes** with **multiple payoffs**.
- In general, the satisfaction of multi-objective queries requires **randomised strategies**.

Main questions

- What is the relationship between expected payoffs of **pure strategies** and expected payoffs of **general strategies**?
- What **type of randomisation** do we need for multi-objective queries?

$\rightarrow$ **Goal**: results for the **broadest possible class of payoffs**.

Markov decision processes

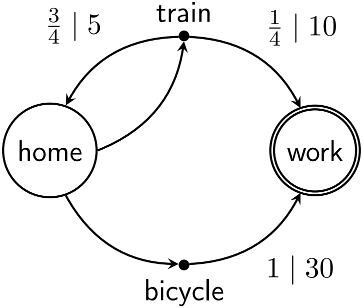

Markov decision process

- **Finite** state space S
- **Finite** action space A
- **Randomised** transitions

Plays are sequences in $(SA)^\omega$ coherent with transitions.

- A **strategy** is a function $\sigma\colon (SA)^*S \to \mathcal{D}(A)$.
- A **payoff** is a measurable function $f\colon \text{Plays} \to \bar{\mathbb{R}}$.

Multi-objective Markov decision processes

We consider **two payoffs**:

- reaching work under 40 minutes with **high probability**;
- minimising the **expectancy** of the time to reach work.

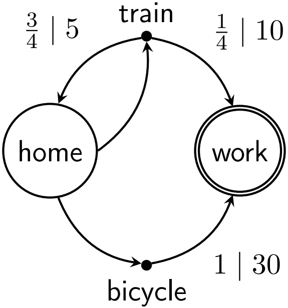

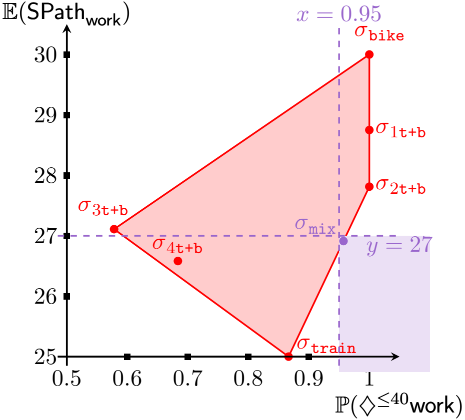

Focus on a specific result

Theorem

Let $\bar{f} = (f_1, \ldots, f_d)$. *Assume that for all strategies* σ, *all states* s *and all* $1 \leq j \leq d$, $\mathbb{E}_s^\sigma(|f_j|) \in \mathbb{R}$. *Then, for all states* s,

$$\mathsf{Pay}_s(\bar{f}) = \mathrm{conv}(\mathsf{Pay}_s^{\mathsf{pure}}(\bar{f})).$$

In particular, to match the expected payoff of any strategy, it suffices to:

- ***mix*** $d+1$ ***pure strategies****;*
- *consider strategies use* ***randomisation at most*** d *along any play.*

We also provide a **variant** of this result for payoff functions violating the assumption of this theorem.

Thank you for your attention.