Different Strokes in Randomised Strategies: Revisiting Kuhn's Theorem Under Finite-memory Assumptions

James C. A. Main Mickael Randour

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

Highlights'22 – June 30 2022

Playing randomly

- In general, one can define randomised strategies in different ways.

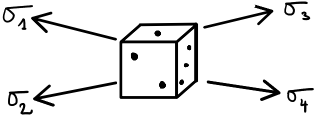

Mixed strategies

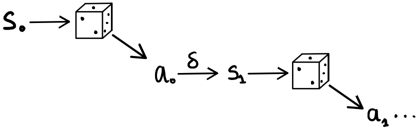

Behavioural strategies

- In general, these two classes of strategies are not comparable.

[1] Aumann, "28. Mixed and Behavior Strategies in Infinite Extensive Games".

Playing randomly

- In general, one can define randomised strategies in different ways.

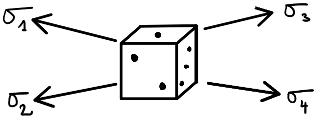

Mixed strategies

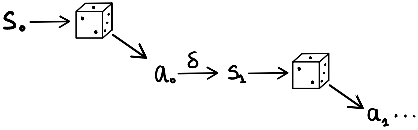

Behavioural strategies

- In general, these two classes of strategies are not comparable.
- Kuhn's theorem [Aum64][1]: in games of perfect recall any mixed strategy has an equivalent behavioural strategy and vice-versa.

Focus of this talk

A Kuhn-inspired classification of finite-memory strategies.

[1] Aumann, "28. Mixed and Behavior Strategies in Infinite Extensive Games".

Setting: finite stochastic games

We consider two-player stochastic games.

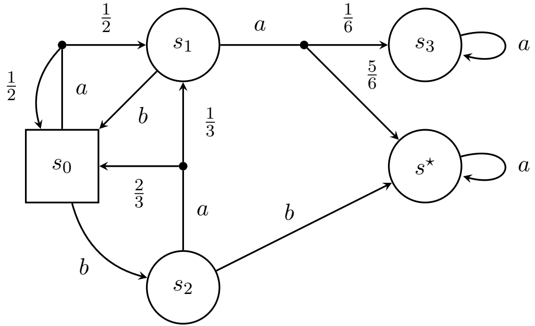

Essential characteristics

- Finite state space $S = S_1 \uplus S_2$ and action space A.
- Players can observe their own actions.

Comparing strategies

Definition

A strategy of $\mathcal{P}_i$ is a function $\sigma_i\colon (SA)^*S_i \to \mathcal{D}(A)$.

Comparing strategies

Definition

A strategy of $\mathcal{P}_i$ is a function $\sigma_i \colon (SA)^* S_i \to \mathcal{D}(A)$.

- We compare strategies independently of any objective or payoff.
- Equality is too restrictive: two different strategies may induce the same behaviour in practice.

Comparing strategies

Definition

A strategy of $\mathcal{P}_i$ is a function $\sigma_i \colon (SA)^*S_i \to \mathcal{D}(A)$.

- We compare strategies independently of any objective or payoff.
- Equality is too restrictive: two different strategies may induce the same behaviour in practice.

Outcome-equivalence

Two strategies σ_1 and τ_1 of $\mathcal{P}_1$ are outcome-equivalent if for all strategies σ_2 of $\mathcal{P}_2$ and all initial states $s_{\mathsf{init}} \in S$, the Markov chain induced from s_{init} by σ_1 and σ_2 is the same than the Markov chain induced from s_{init} by τ_1 and σ_2.

Randomised finite-memory strategies

Definition

A strategy σ_i of $\mathcal{P}_i$ is finite-memory if it is induced by a stochastic Mealy machine $\mathcal{M} = (M, \mu_{\mathsf{init}}, \alpha_{\mathsf{next}}, \alpha_{\mathsf{up}})$ where

- M is a finite set of memory states;
- $\mu_{\mathsf{init}} \in \mathcal{D}(M)$ is an initial distribution;
- $\alpha_{\mathsf{next}} \colon M \times S_i \to \mathcal{D}(A)$ is a stochastic next-move function;
- $\alpha_{\mathsf{up}} \colon M \times S \times A \to \mathcal{D}(M)$ is a stochastic memory update function.

Randomised finite-memory strategies

Definition

A strategy σ_i of $\mathcal{P}_i$ is finite-memory if it is induced by a stochastic Mealy machine $\mathcal{M} = (M, \mu_{\mathsf{init}}, \alpha_{\mathsf{next}}, \alpha_{\mathsf{up}})$ where

- M is a finite set of memory states;
- $\mu_{\mathsf{init}} \in \mathcal{D}(M)$ is an initial distribution;
- $\alpha_{\mathsf{next}} \colon M \times S_i \to \mathcal{D}(A)$ is a stochastic next-move function;
- $\alpha_{\mathsf{up}} \colon M \times S \times A \to \mathcal{D}(M)$ is a stochastic memory update function.

- In the literature, variations of this model where some components are not randomised are sometimes used.
- We can classify Mealy machines following whether their initialisation, updates and outputs are randomised or deterministic.

All classes of Mealy machines are not equally powerful

- Some classes of Mealy machines allow richer behaviours than others.

All classes of Mealy machines are not equally powerful

- Some classes of Mealy machines allow richer behaviours than others.
- For instance, the strategy illustrated on the right cannot be emulated with randomisation only in the outputs.

Game

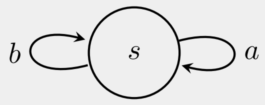

Mealy machine example

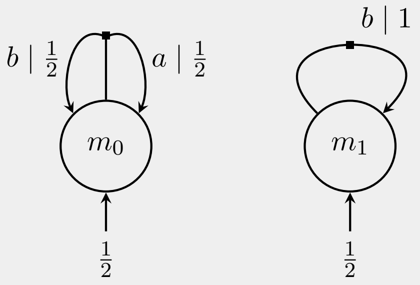

Our results

We use acronyms to define classes of Mealy machines: we use XYZ where X, Y, Z $\in$ {D, R} where D stands for deterministic and R for random, and

- X characterises initialisation,
- Y characterises outputs (next-move function),
- Z characterises updates.

Our results

We use acronyms to define classes of Mealy machines: we use XYZ where X, Y, Z $\in$ {D, R} where D stands for deterministic and R for random, and

- X characterises initialisation,
- Y characterises outputs (next-move function),
- Z characterises updates.

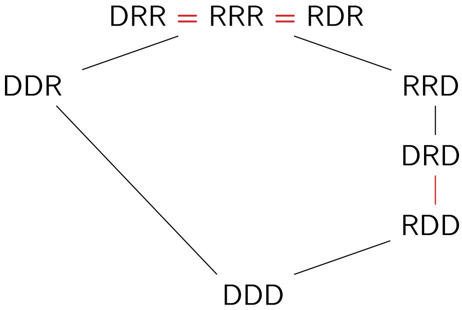

References I

Aumann, Robert J . "28. Mixed and Behavior Strategies in Infinite Extensive Games". In: *Advances in Game Theory. (AM-52), Volume 52*. Princeton University Press, 2016, pp. 627–650. DOI: `doi:10.1515/9781400882014-029`. URL: `https://doi.org/10.1515/9781400882014-029`.