Dresden University of Technology

Computational Logic ∴ Group

Hannes Strass

Faculty of Computer Science, Institute of Artificial Intelligence, Computational Logic Group

Games with Missing Information: Modelling

Lecture 7, 1st June 2026 // Algorithmic Game Theory, SS 2026

Previously . . .

- **Monte Carlo Tree Search** uses random playouts to evaluate moves and keeps statistics on which moves led to which payoffs how many times.
- A **selection policy** balances exploitation and exploration.
- **UCT** is an effective selection policy that applies UCB1 to trees.
- A **playout policy** steers playout simulations towards realistic play.
- MCTS and deep reinforcement learning led to expert-level Go programs.

wide, but shallow

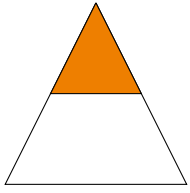

Type A

Alpha-Beta Tree Search

narrow, but deep

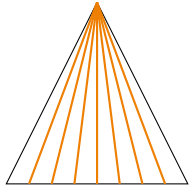

Type B

Monte Carlo Tree Search

Overview

Example: The Monty Hall Problem

Extensive-Form Games

Behaviour Strategies

Kuhn's Theorem

Motivation: Missing Information

- So far, we have considered games with perfect information:
- In every state, all players know the full history of play so far, i.e. they know their (joint) position in the game tree.
- However, e.g. in card games, players typically do not know the cards of opponents.
- This form of incomplete knowledge can be formalised by sets of indistinguishable nodes in the game tree, typically called information sets.
- In this context, we also add another element to games: chance.
- This is modelled via moves by nature and can be used to formalise dealing cards or throwing dice.
- We will see that this also allows us to model games with incomplete information, where e.g. some of the payoffs may be uncertain.
- In principle, however, chance and imperfect information are unrelated and we could model either without the other.

Example: The Monty Hall Problem

Example: The Monty Hall Problem

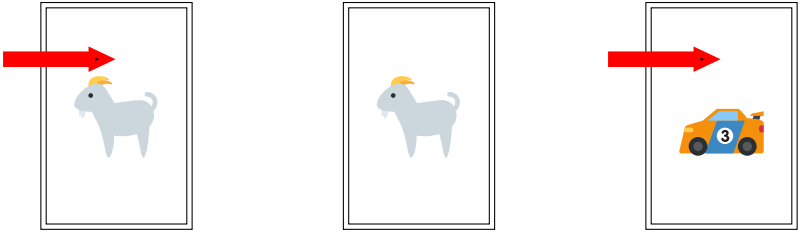

The Monty Hall Problem

A game show participant (`Guest`) is shown three doors behind which there are prizes. Behind one door, there is an expensive car, behind each of the other doors there is a goat. (The participant prefers the car over a goat.) The `Guest` is asked to `Choose` one of the doors. The game show `Host` (the other player) now opens one of the remaining doors that has a goat behind it. The `Guest` then gets their final move: `Stay` with the door they initially picked, or `Switch` to the other door.

What should the participant do?

The Monty Hall Problem: Game Tree Sketch

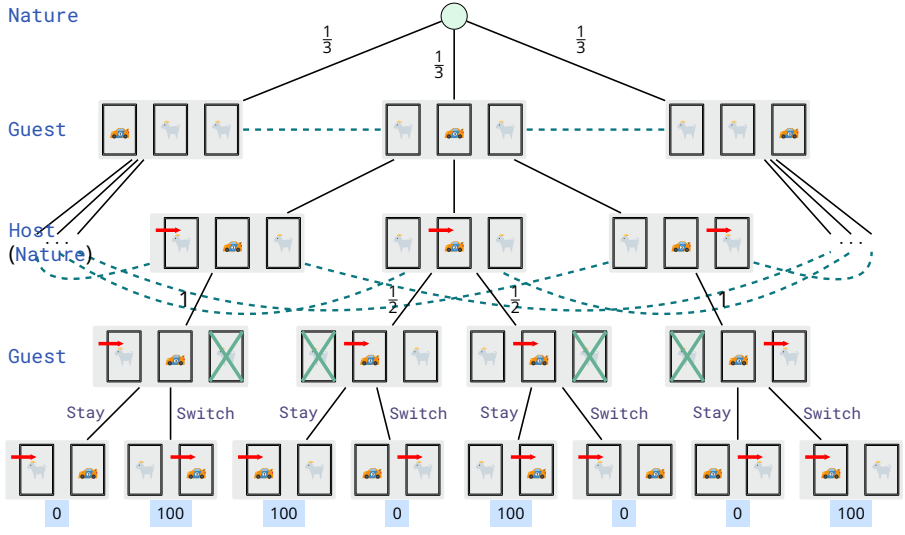

The Monty Hall Problem: Analysis

- Each of the possible states s_1, s_2, s_3 after `Nature`'s move has probability $\frac{1}{3}$.
- For each of these states, the ensuing game is symmetric.
- If `Guest` chooses their door uniformly at random, then:
  - With probability $\frac{1}{3}$, their initial guess is correct; thus `Switch` has a payoff of 0.
  - With probability $\frac{2}{3}$, their initial guess is wrong; thus `Switch` has a payoff of 100.
- Thus in each s_i, `Switch` has an expected payoff of $\frac{1}{3} \cdot 0 + \frac{2}{3} \cdot 100$.
- The overall payoff of `Switch` is thus

$$U_{\texttt{Guest}}(\texttt{Switch}) = 3 \cdot \frac{1}{3} \cdot \left(\frac{1}{3} \cdot 0 + \frac{2}{3} \cdot 100 \right) = 66\frac{2}{3}$$

- Likewise, the overall payoff of `Stay` is obtained as

$$U_{\texttt{Guest}}(\texttt{Stay}) = \frac{1}{3} \cdot 100 = 33\frac{1}{3}$$

- Therefore, a rational player should always choose `Switch` over `Stay`.

Extensive-Form Games

Missing Information: Formalisation

Definition

An **extensive-form game** consists of the following:

1. A set $P = \{1, \dots, n\}$ of players, and possibly `Nature`.
2. An $n+1$-tuple $\mathbf{M} = (M_1, \dots, M_n, M_{\texttt{Nature}})$ of sets M_i of **moves** for all players.
3. A set H of **histories**, sequences of moves $m_j \in M_1 \cup \dots \cup M_n \cup M_{\texttt{Nature}}$.
4. A subset $Z \subseteq H$ of **terminal** histories.
5. A partition $\mathcal{I}_1 \dot\cup \dots \dot\cup \mathcal{I}_k = H \setminus Z$ of non-terminal histories into **information sets** such that for all $1 \leq j \leq k$, all $h_1, h_2 \in \mathcal{I}_j$ have the same legal moves.
6. A **player function** $p\colon \{1, \dots, k\} \to P \cup \{\texttt{Nature}\}$ (stating whose turn it is).
7. An n-tuple $\mathbf{u} = (u_1, \dots, u_n)$ of utility functions $u_i\colon Z \to \mathbb{R}$.

Starting with the **empty history** $[]$, in each history $h = [m_1, \dots, m_\ell] \in H \setminus Z$, player $i = p(h)$ chooses a move $m \in M_i$, leading to the history $[m_1, \dots, m_\ell, m]$. Each $\mathcal{I}_j$ with $p(\mathcal{I}_j) = \texttt{Nature}$ has a probability distribution on possible moves.

Information Sets: Remarks

Intuition of information sets: The player (whose turn it is) does not have the information to distinguish between states in the set (but from other sets).

- $\mathcal{I} = \{\mathcal{I}_1, \ldots, \mathcal{I}_k\}$ being a **partition** means that:
  - for all $1 \leq j \leq k$, we have $\mathcal{I}_j \neq \emptyset$,
  - $\mathcal{I}_1 \cup \ldots \cup \mathcal{I}_k = H \setminus Z$, and
  - for all $1 \leq j, \ell \leq k$, we have $\mathcal{I}_j \cap \mathcal{I}_\ell = \emptyset$.
- Thus every $h \in H \setminus Z$ belongs to exactly one information set $\mathcal{I}_h \in \mathcal{I}$.
- For all $1 \leq j \leq k$ and $h \in \mathcal{I}_j$, we denote $p(h) := p(\mathcal{I}_j)$.
- $\mathcal{I}$ can also be represented by an **equivalence relation** $\sim^G$, where for any $h_1, h_2 \in H$, we have $h_1 \sim^G h_2$ iff there is a $\mathcal{I}_j \in \mathcal{I}$ such that $h_1, h_2 \in \mathcal{I}_j$.
- We graphically represent $\sim^G$ in game trees via dashed edges - - - - -.

Battleship

The initial placement of ships is private to the players and can be modelled via information sets. Some information may later be disclosed through hits.

Information Sets: Example

The Monty Hall Problem

- The (true) initial state is represented by the information set $\mathcal{I}_0 = \{[]\}$.
- The (seemingly) initial state for `Guest` is given by the information set
$$\mathcal{I}_1 = \{[\texttt{Car1}], [\texttt{Car2}], [\texttt{Car3}]\}\,.$$
- For each possible (initial) choice of door for `Guest`, there is one set:
$$\mathcal{I}_{\texttt{Choose1}} = \{[\texttt{Car1}, \texttt{Choose1}], [\texttt{Car2}, \texttt{Choose1}], [\texttt{Car3}, \texttt{Choose1}]\}$$
$$\mathcal{I}_{\texttt{Choose2}} = \{[\texttt{Car1}, \texttt{Choose2}], [\texttt{Car2}, \texttt{Choose2}], [\texttt{Car3}, \texttt{Choose2}]\}$$
$$\mathcal{I}_{\texttt{Choose3}} = \{[\texttt{Car1}, \texttt{Choose3}], [\texttt{Car2}, \texttt{Choose3}], [\texttt{Car3}, \texttt{Choose3}]\}$$
- Some information is disclosed by the host opening a door:
$$\mathcal{I}_{[\texttt{Choose1},\texttt{Open2}]} = \{[\texttt{Car1}, \texttt{Choose1}], [\texttt{Car3}, \texttt{Choose1}]\}$$

Strategic Games and Imperfect Information

- Uncertainty induced by simultaneous moves can be modelled in extensive-form games (that seem to be sequential by definition).
- Main idea: Sequentialise moves, model uncertainty in information sets.

Example: Recall the game penalties. One extensive-form variant is:

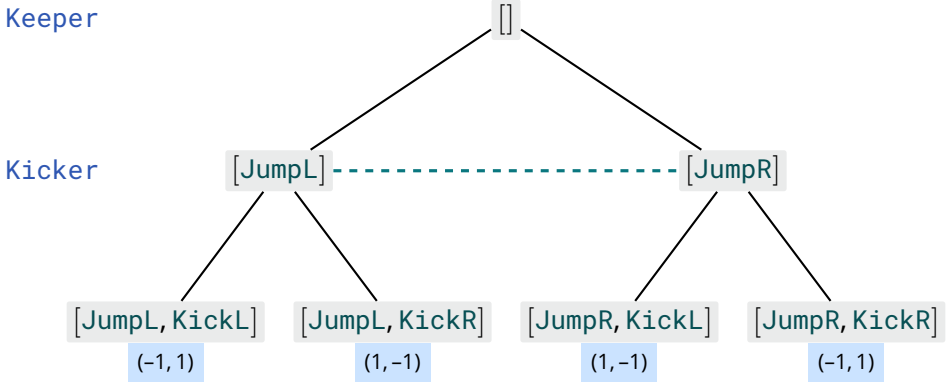

Chance Nodes (Moves by `Nature`)

Intuition of chance nodes: Something happens that is controlled by an entity with no strategic interest in the game's outcome.

Examples

- In card games, `Nature` controls the dealer's shuffling the cards.
- In games involving dice, `Nature` controls the dice throws.

- Probability distributions model uncertainty about effects of such actions.
- We typically use uniform distributions over possible atomic results.
- In the remainder of the course, we use conditional probabilities ...
- ...and how they relate to marginal probabilities via Bayes' Theorem.
- The relevant definitions and results are in the appendix.

Behaviour Strategies

From Extensive Form to Normal Form

Definition

Let $G = (P, \mathbf{M}, H, Z, \{\mathcal{I}_1, \ldots, \mathcal{I}_k\}, p, \mathbf{u})$ be a finite extensive-form game.

1. For $i \in P$, we denote $p^{-1}(i) = \{\mathcal{I}_j \in \mathcal{I} \mid p(\mathcal{I}_j) = i\}$.
2. The **possible moves** of player $i \in P$ at information set $\mathcal{I}_j \in p^{-1}(i)$ are:
$$M_i(\mathcal{I}_j) = \{m \in M_i \mid [h; m] \in H \text{ for some } h \in \mathcal{I}_j\}$$
3. A **pure strategy** for player $i \in P$ is a function
$$s_i \colon p^{-1}(i) \to M_i \quad \text{where} \quad s_i(\mathcal{I}_j) \in M_i(\mathcal{I}_j) \text{ for all } \mathcal{I}_j \in p^{-1}(i)$$

The **associated normal-form game** $G' = (P, \mathbf{S}, \mathbf{u}')$ is as follows:

- $\mathbf{S} = (S_1, \ldots, S_n)$ with each S_i the set of all pure strategies for player i.
- $\mathbf{u}' = (u'_1, \ldots, u'_n)$ with each u'_i mapping profile $\mathbf{s} = (s_1, \ldots, s_n)$ to $u_i(\langle\!\langle \mathbf{s}, [] \rangle\!\rangle)$, where $\langle\!\langle \mathbf{s}, [] \rangle\!\rangle$ is the terminal history obtained from playing as specified by $\mathbf{s}$.

As usual, strategies can be reduced; pure strategies can be mixed.

Extensive to Normal Form: Example

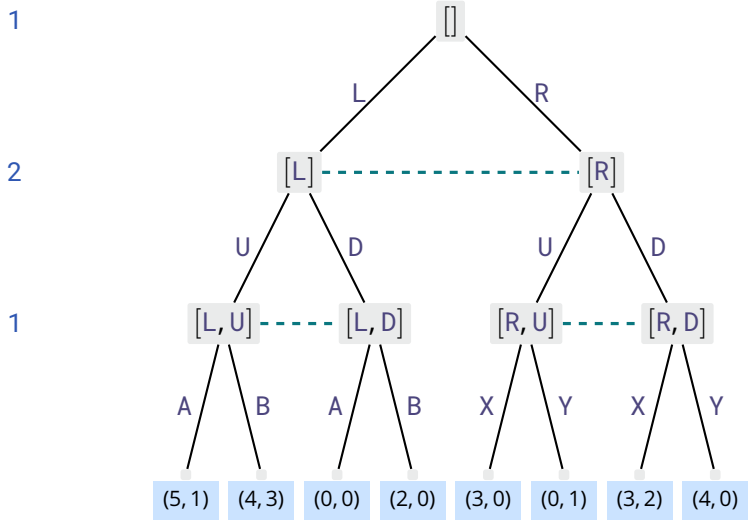

(1, 2)	U	D
LA∗	(5, 1)	(0, 0)
LB∗	(4, 3)	(2, 0)
R∗X	(3, 0)	(3, 2)
R∗Y	(0, 1)	(4, 0)

Strategies and the Normal Form: Remarks

- Playing a mixed strategy $\pi\colon S_i \to [0,1]$ in an extensive-form game can be understood as randomly choosing a pure strategy $s_i \sim \pi$ and then strictly following this pure strategy s_i during gameplay.
- Decisions in different parts of the game tree are determined by the chosen s_i and thus not independent.
- For an extensive-form game $G = (P, \mathbf{M}, H, Z, \mathcal{I}, p, \mathbf{u})$, player $i \in P$ may have

$$|M_i|^{|p^{-1}(i)|}$$

  different pure strategies (in the normal-form game) in the worst case.
- Practically, we are however mainly interested in local decision making.
- In a behaviour strategy σ, a player has a probability distribution $\sigma(\mathcal{I}_j)$ in each information set $\mathcal{I}_j$ and randomly chooses a move $m \sim \sigma(\mathcal{I}_j)$ when in $\mathcal{I}_j$.
- Probability distributions at different information sets are independent.

Behaviour Strategies

Definition

Let G be an extensive-form game with players P and information sets $\mathfrak{I}$. A **behaviour strategy** for player $i \in P$ is a function σ_i that assigns a probability distribution over possible moves to each of player i's information sets.

- $\sigma_i(\mathfrak{I}_j)(m_k)$ denotes the probability that player i will make move m_k at information set $\mathfrak{I}_j$. For readability, we will write this as $\sigma_i(m_k \mid \mathfrak{I}_j)$.
- Every behaviour strategy σ of player $i \in P$ induces a mixed strategy π_σ:

$$\pi_\sigma := \left\{ s_i \mapsto \prod_{\mathfrak{I}_j \in p^{-1}(i)} \sigma(s_i(\mathfrak{I}_j) \mid \mathfrak{I}_j) \;\middle|\; s_i \in S_i \right\}$$

- In particular, a pure strategy s_i can be seen as a behaviour strategy σ_i where any $s_i(\mathfrak{I}_j) = m_k$ induces $\sigma_i(m_k \mid \mathfrak{I}_j) = 1$ and $\sigma_i(m_\ell \mid \mathfrak{I}_j) = 0$ for $m_\ell \in M_i(\mathfrak{I}_j)$, $\ell \neq k$.

Behaviour vs. Mixed Strategies: Examples

Consider the following extensive-form game G_5 and its normal form:

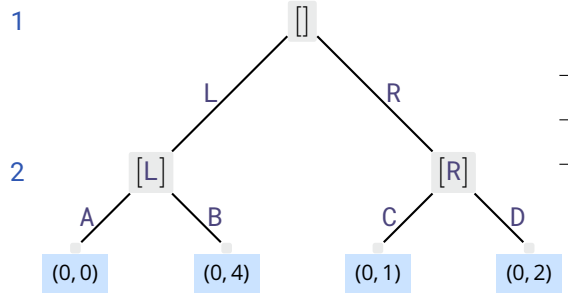

(1, 2)	AC	AD	BC	BD
L	(0, 0)	(0, 0)	(0, 4)	(0, 4)
R	(0, 1)	(0, 2)	(0, 1)	(0, 2)

- A mixed strategy for player 2 is $\pi = \left(\frac{1}{2}, 0, 0, \frac{1}{2}\right) \mathrel{\hat{=}} \left\{\mathsf{AC} \mapsto \frac{1}{2}, \mathsf{BD} \mapsto \frac{1}{2}, \ldots\right\}$.
- The "closest" behaviour strategy σ assigns $\sigma([\mathsf{L}]) = \left\{\mathsf{A} \mapsto \frac{1}{2}, \mathsf{B} \mapsto \frac{1}{2}\right\}$ and $\sigma([\mathsf{R}]) = \left\{\mathsf{C} \mapsto \frac{1}{2}, \mathsf{D} \mapsto \frac{1}{2}\right\}$, inducing mixed strategy $\pi_\sigma = \left(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}\right) \neq \pi$.

Kuhn’s Theorem

Idea and Intuition

Do we always need the “full power” of mixed strategies to analyse extensive-form games?

- Representing a mixed strategy requires up to $|M_i|^{|p^{-1}(i)|}$ numbers.
- Representing a behaviour strategy requires up to $|M_i| \cdot |p^{-1}(i)|$ numbers.

Question

Can mixed strategies always be equivalently written as behaviour strategies?

What does “equivalent” even mean in this context?

$\rightsquigarrow$ Intuitively: When all histories are reached with the same probability.

(Formal definition later.)

Equivalent Behaviour Strategies: Example

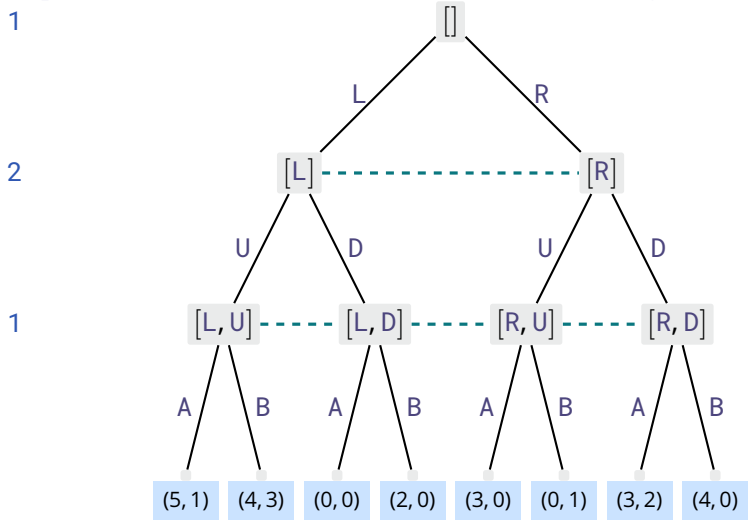

(1, 2)	U	D
LA	(5, 1)	(0, 0)
LB	(4, 3)	(2, 0)
RA	(3, 0)	(3, 2)
RB	(0, 1)	(4, 0)

Player 1's mixed strategy $\left(0, \frac{2}{5}, \frac{3}{5}, 0\right)$ has no equivalent behaviour strategy.

Perfect Recall

Definition

Let $G = (P, \mathbf{M}, H, \mathcal{I}, p, \mathbf{u})$ be an extensive-form game.

- For every player $i \in P$ and history $h \in H$, define the sequence h_i of pairs $(\mathcal{I}_j, m)$ for $\mathcal{I}_j \in \mathcal{I}$ and $m \in M_i$ by induction:

$$[]_i := [] \qquad \text{and} \qquad [h;m]_i := \begin{cases} [h_i; (\mathcal{I}_h, m)] & \text{if } p(h) = i, \\ h_i & \text{otherwise.} \end{cases}$$

  where for $h = [m_1, \ldots, m_\ell]$ we denote $[h;m] = [m_1, \ldots, m_\ell, m]$ as usual.
- Player $i \in P$ has **perfect recall** in G iff
  for all $\mathcal{I}_j \in \mathcal{I}$, for every $h, h' \in \mathcal{I}_j$, it holds that $h_i = h'_i$.
- G has **perfect recall** iff every player $i \in P$ has perfect recall in G.

- h_i extracts all decision points and decisions (of player i) from history h.
- $h_i = h'_i$ means that i made the same moves in the same information sets.
- With perfect recall, players remember their trajectory through the game.

Perfect Recall: Examples (1)

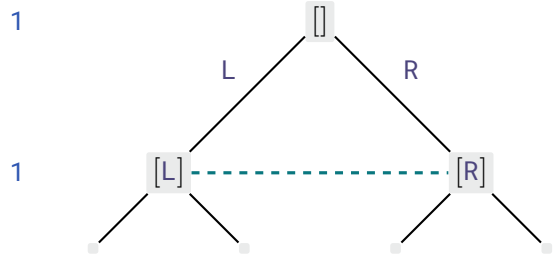

This game does not have perfect recall:

- Denote $\mathcal{I}_0 = \{[]\}$ and $\mathcal{I}_1 = \{[\mathsf{L}], [\mathsf{R}]\}$.
- We have $[\mathsf{L}] \in \mathcal{I}_1$ and $[\mathsf{R}] \in \mathcal{I}_1$, but:

$$[\mathsf{L}]_1 = [(\mathcal{I}_0, \mathsf{L})] \neq [(\mathcal{I}_0, \mathsf{R})] = [\mathsf{R}]_1$$

Perfect Recall: Examples (2)

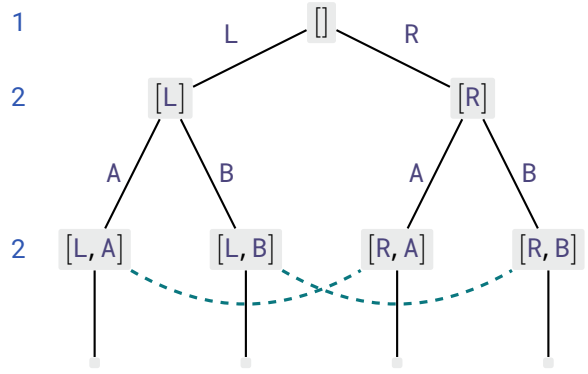

This game does not have perfect recall:

- Denote $\mathcal{I}_0 = \{[]\}$, $\mathcal{I}_1 = \{[\mathsf{L}]\}$, $\mathcal{I}_2 = \{[\mathsf{R}]\}$, $\mathcal{I}_3 = \{[\mathsf{L}, \mathsf{A}], [\mathsf{R}, \mathsf{A}]\}$, $\mathcal{I}_4 = \{[\mathsf{L}, \mathsf{B}], [\mathsf{R}, \mathsf{B}]\}$.
- Then $[\mathsf{L}, \mathsf{A}]_2 = [(\mathcal{I}_1, \mathsf{A})] \neq [(\mathcal{I}_2, \mathsf{A})] = [\mathsf{R}, \mathsf{A}]_2$.

Perfect Recall: Examples (3)

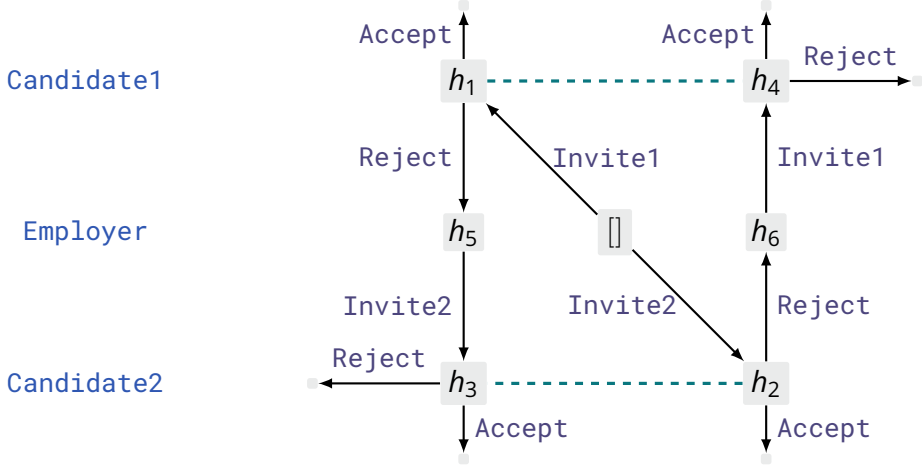

This game has perfect recall: e.g. $h \in \{h_2, h_3\}$ implies $h_{\texttt{Candidate2}} = []$.

Realisation Equivalence

Definition

Let $G = (P, \mathbf{M}, H, Z, \mathcal{I}, p, \mathbf{u})$ be an extensive-form game and $i \in P$.

1. For every history $h \in H$ and mixed strategy $\pi_i \colon S_i \to [0,1]$, the probability of i moving as in h_i when playing according to π is $P(h_i \,|\, \pi)$ with

$$P([(\mathcal{I}_1, m_1), \ldots, (\mathcal{I}_\ell, m_\ell)] \,|\, \pi) := \sum_{\substack{s_i \in S_i, \\ s_i(\mathcal{I}_1) = m_1, \ldots, s_i(\mathcal{I}_\ell) = m_\ell}} \pi(s_i)$$

2. For behaviour strategy σ_i, $P([(\mathcal{I}_1, m_1), \ldots, (\mathcal{I}_\ell, m_\ell)] \,|\, \sigma_i) := \prod_{1 \leq j \leq \ell} \sigma_i(m_j \,|\, \mathcal{I}_j)$.

3. Mixed strategy π and behaviour strategy σ for player $i \in P$ are **realisation-equivalent** iff for all $h \in H$: $\quad P(h_i \,|\, \pi) = P(h_i \,|\, \sigma)$.

In particular $P([] \,|\, \pi) = \sum_{s_i \in S_i} \pi(s_i) = 1 = P([] \,|\, \sigma)$.

Realisation Equivalence: Example & Remark

Consider again extensive-form game G_5:

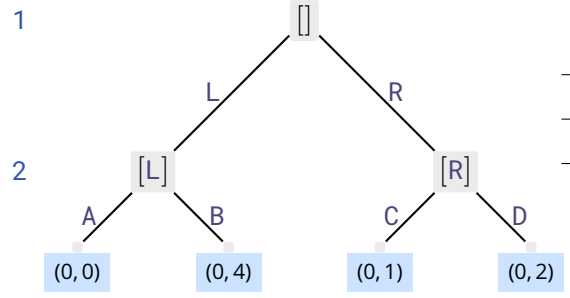

(1, 2)	AC	AD	BC	BD
L	(0, 0)	(0, 0)	(0, 4)	(0, 4)
R	(0, 1)	(0, 2)	(0, 1)	(0, 2)

- Player 2's mixed strategy $\pi = \left(\frac{1}{2}, 0, 0, \frac{1}{2}\right)$ and behaviour strategy σ with $\sigma(\mathsf{A} \,|\, [\mathsf{L}]) = \sigma(\mathsf{B} \,|\, [\mathsf{L}]) = \frac{1}{2}$ and $\sigma(\mathsf{B} \,|\, [\mathsf{R}]) = \sigma(\mathsf{C} \,|\, [\mathsf{R}]) = \frac{1}{2}$ are realisation-equivalent:

$$P([(\{[\mathsf{L}]\}, \mathsf{A})] \,|\, \pi) = \tfrac{1}{2} = P([(\{[\mathsf{L}]\}, \mathsf{A})] \,|\, \sigma) \text{ and similarly for } \{[\mathsf{R}]\}.$$

- Generally, since in particular terminal histories are reached with the same probabilities, realisation-equivalent strategies obtain equal game values.

Kuhn's Theorem

Theorem [Kuhn, 1953]

Let G be an extensive-form game with perfect recall. For every mixed strategy π for player $i \in P$, there exists a realisation-equivalent behaviour strategy σ.

Proof.

- Since G has perfect recall, for any $\mathcal{I}_j \in p^{-1}(i)$ we have that $h', h'' \in \mathcal{I}_j$ satisfy $h'_i = h''_i$ and thus any $h \in \mathcal{I}_j$ uniquely defines $P(\mathcal{I}_j \mid \pi) := P(h_i \mid \pi)$.
- For $\mathcal{I}_h \in p^{-1}(i)$ with $P(\mathcal{I}_h \mid \pi) > 0$ and $m \in M_i(\mathcal{I}_h)$, set $\sigma(m \mid \mathcal{I}_h) := \frac{P([h_i;(\mathcal{I}_h,m)] \mid \pi)}{P(h_i \mid \pi)}$.
- To show realisation equivalence, consider any $h_i = [(\mathcal{I}_1, m_1), \ldots, (\mathcal{I}_\ell, m_\ell)]$:

$$
\begin{aligned}
P(h_i \mid \sigma) &= \sigma(m_1 \mid \mathcal{I}_1) \cdot \sigma(m_2 \mid \mathcal{I}_2) \cdot \ldots \cdot \sigma(m_\ell \mid \mathcal{I}_\ell) \\
&= \frac{P([(\mathcal{I}_1, m_1)] \mid \pi)}{P([] \mid \pi)} \cdot \frac{P([(\mathcal{I}_1, m_1), (\mathcal{I}_2, m_2)] \mid \pi)}{P([(\mathcal{I}_1, m_1)] \mid \pi)} \cdot \ldots \cdot \frac{P([(\mathcal{I}_1, m_1), \ldots, (\mathcal{I}_\ell, m_\ell)] \mid \pi)}{P([(\mathcal{I}_1, m_1), \ldots, (\mathcal{I}_{\ell-1}, m_{\ell-1})] \mid \pi)} \\
&= \frac{P([(\mathcal{I}_1, m_1), \ldots, (\mathcal{I}_\ell, m_\ell)] \mid \pi)}{P([] \mid \pi)} = P(h_i \mid \pi)
\end{aligned}
$$

□

Example: Simplified Poker

Binmore's Simplified Poker

- Two players, `Ann` and `Bob`, each put $1 into a jackpot.
- They then draw one card from a deck of three cards: $\{1, 2, 3\}$.
- `Ann` can either `check` (pass on), or `raise` (put another $1 into the jackpot).
- Next, `Bob` responds:
  - If `Ann` has `check`ed, then `Bob` must `call`, that is, a showdown happens: Both players show their cards and the player with the higher (number) card receives the jackpot.
  - If `Ann` has `raise`d, then `Bob` can decide between `fold` (withdraw from the game and let `Ann` get the jackpot) or `call` (put another $1 into the jackpot and then have a showdown).

Simplified Poker: Preliminary Analysis

`Nature` shuffles and deals the cards. There are six possible outcomes:

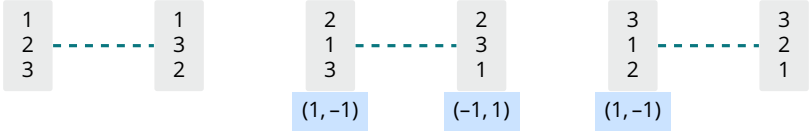

- If `Ann` draws a 3, she will `raise`; if `Bob` draws a 1, he will `fold`.
- If `Bob` draws a 3, he will `call`; if `Ann` draws a 2, she will `check`:
  Were she to `raise`, she would lose 2 if `Bob` has a 3 (as he would `call`), but still only win 1 if `Bob` has a 1 (as he would `fold` then).

What happens in the two remaining cases?

1. Should `Ann` `raise` (i.e. bluff) if she has a 1?
2. Should `Bob` `call` (the bluff) if he has a 2?

Conclusion

Summary

- In **complete information** games, players know the rules, possible outcomes and each other's preferences over outcomes.
- In **perfect information** games, all players know all previous moves.
- An **extensive-form** game may have incomplete or imperfect information.
- Uncertainty of players (due to missing information) can be modelled by **information sets** and **chance nodes** (moves by `Nature`).
- A **behaviour strategy** assigns move probabilities to information sets.
- In a game with **perfect recall**, players remember their own previous moves and the information sets in which they played them.
- **Kuhn's Theorem:** Considering only behaviour strategies "is enough" for playing games with perfect recall.

Goat and Car graphics: Twemoji, Copyright 2020 Twitter, Inc and other contributors (CC-BY 4.0)

Appendix: Probabilities

Recall

- A **probability space** is a finite set $\mathcal{E} = \{e_1, \ldots, e_k\}$ of **atomic events**.
- A **probability distribution** is a mapping $P\colon \mathcal{E} \to [0,1]$, where atomic event e_i occurs with probability $P(e_i)$ and we have $\sum_{i=1}^{k} P(e_i) = 1$.
- An **event** $E \subseteq \mathcal{E}$ has (total) probability $P(E) = \sum_{e \in E} P(e)$.
- For all events $A, B \subseteq \mathcal{E}$ we have the following:
  1. $0 \leq P(A) \leq 1$ with $P(\emptyset) = 0$ and $P(\mathcal{E}) = 1$.
  2. $P(\overline{A}) = 1 - P(A)$ where $\overline{A} := \mathcal{E} \setminus A$ is the event **complementary** to A.
  3. $P(A \cup B) = P(A) + P(B) - P(A \cap B)$.

Example

If all events $e_i \in \mathcal{E}$ have the same probability $\frac{1}{|\mathcal{E}|}$, we have a **uniform distribution**.

Conditional Probabilities

Definition

Let A and B be events with $P(B) > 0$.

1. The **conditional probability** for A to occur under the condition of B occurring is
$$P(A \mid B) := \frac{P(A \cap B)}{P(B)}$$
2. Events A and B are **independent** iff
$$P(A \cap B) = P(A) \cdot P(B)$$

That events A and B are independent is equivalently characterised by each of:

- $P(A \mid B) = P(A)$
- $P(A \mid B) = P(A \mid \overline{B})$
- $P(B \mid A) = P(B)$
- $P(B \mid A) = P(B \mid \overline{A})$

Bayes' Theorem

Theorem (Bayes)

1. If A and B are two events with $P(A) > 0$ and $P(B) > 0$, then
$$P(A) \cdot P(B \mid A) = P(B) \cdot P(A \mid B)$$
2. If A and $B_1, B_2, \ldots, B_\ell$ are events with $P(A) > 0$ and $P(B_i) > 0$ for all $1 \leq i \leq \ell$, where $\bigcup_{i=1}^{\ell} B_i = \mathcal{E}$ is a partition of $\mathcal{E}$, then for every $1 \leq i \leq \ell$:
$$P(B_i \mid A) = \frac{P(A \mid B_i) \cdot P(B_i)}{\sum_{j=1}^{\ell} \big(P(A \mid B_j) \cdot P(B_j)\big)} = \frac{P(A \mid B_i) \cdot P(B_i)}{P(A)}$$

In the second item of the theorem, the **law of total probability** is used:

$$P(A) = \sum_{j=1}^{\ell} P(B_j \cap A) = \sum_{j=1}^{\ell} \big(P(A \mid B_j) \cdot P(B_j)\big)$$

Note that $P(A) = P(\mathcal{E} \cap A) = P((\bigcup_{j=1}^{\ell} B_j) \cap A) = P(\bigcup_{j=1}^{\ell} (B_j \cap A)) = \sum_{j=1}^{\ell} P(B_j \cap A)$.