Repeated Games

CMPUT 654: Modelling Human Strategic Behaviour

S&LB §6.1

Assignment #1

- **Assignment #1 is now available on eclass**
  - Worth approximately 19% of your total mark
  - Don't leave it to the last minute
- Due **Tuesday Feb 6** at 11:59pm (**next week**)

Recap: Imperfect Information Extensive Form Example

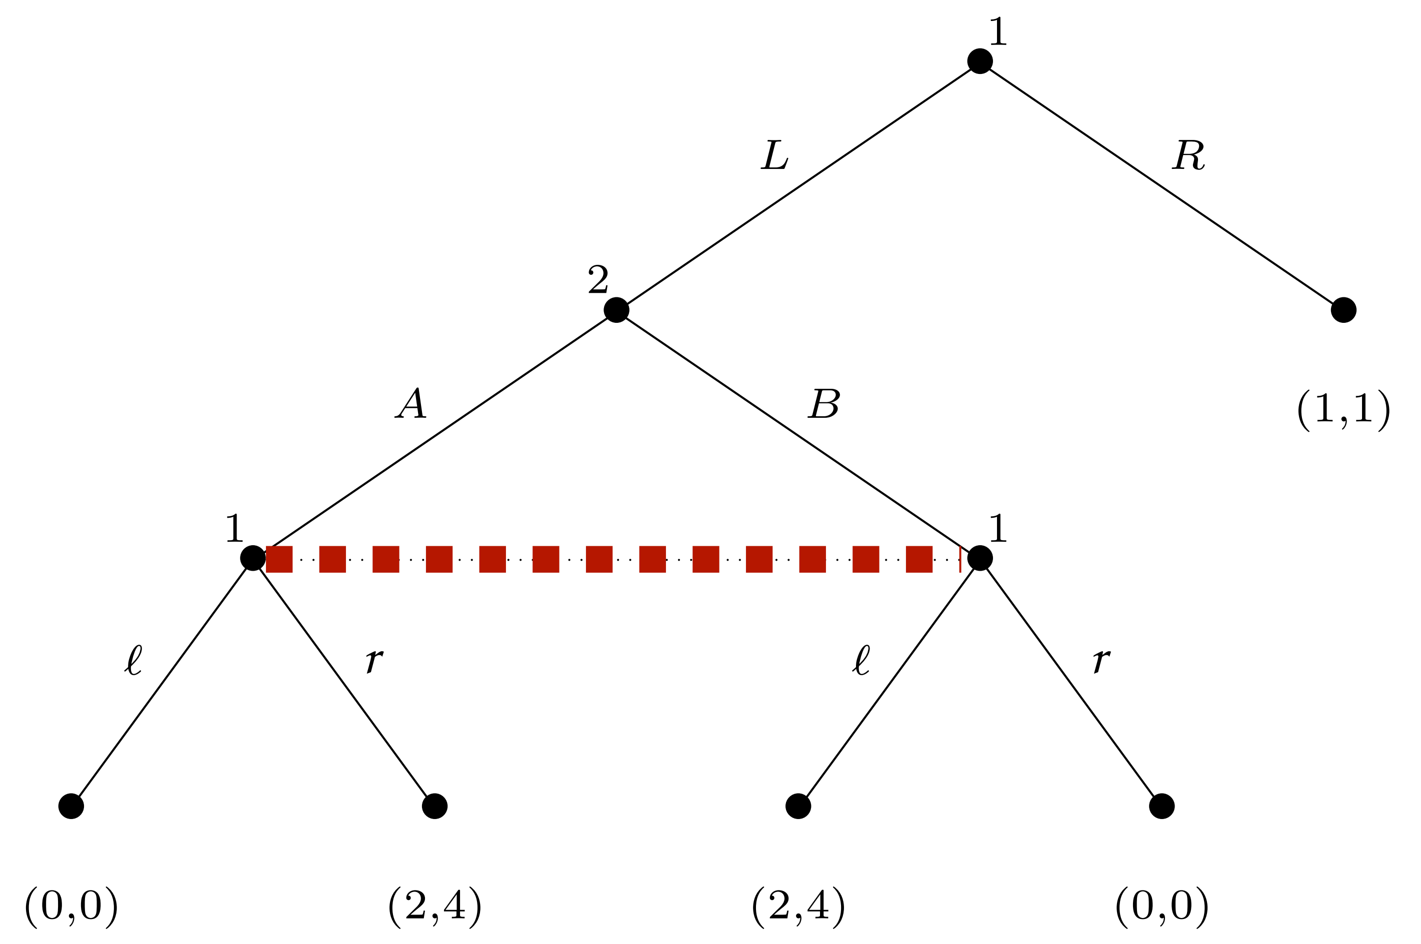

- We represent sequential play using **extensive form games**
- In an **imperfect information** extensive form game, we represent private knowledge by grouping histories into **information sets**
- Players **cannot distinguish** which **history** they are in within an information set

Recap: Behavioural vs. Mixed Strategies

Definition:

A **mixed strategy** $s_i \in \Delta(A^{I_i})$ is any distribution over an agent's **pure strategies**.

Definition:

A **behavioural strategy** $b_i \in [\Delta(A)]^{I_i}$ is a probability distribution over an agent's actions at an **information set**, which is **sampled independently** each time the agent arrives at the information set.

Kuhn's Theorem:

These are **equivalent**(*) in games of **perfect recall**.

Recap:
Normal to Extensive Form

	c	d
C	-1,-1	-4,0
D	0,-4	-3,-3

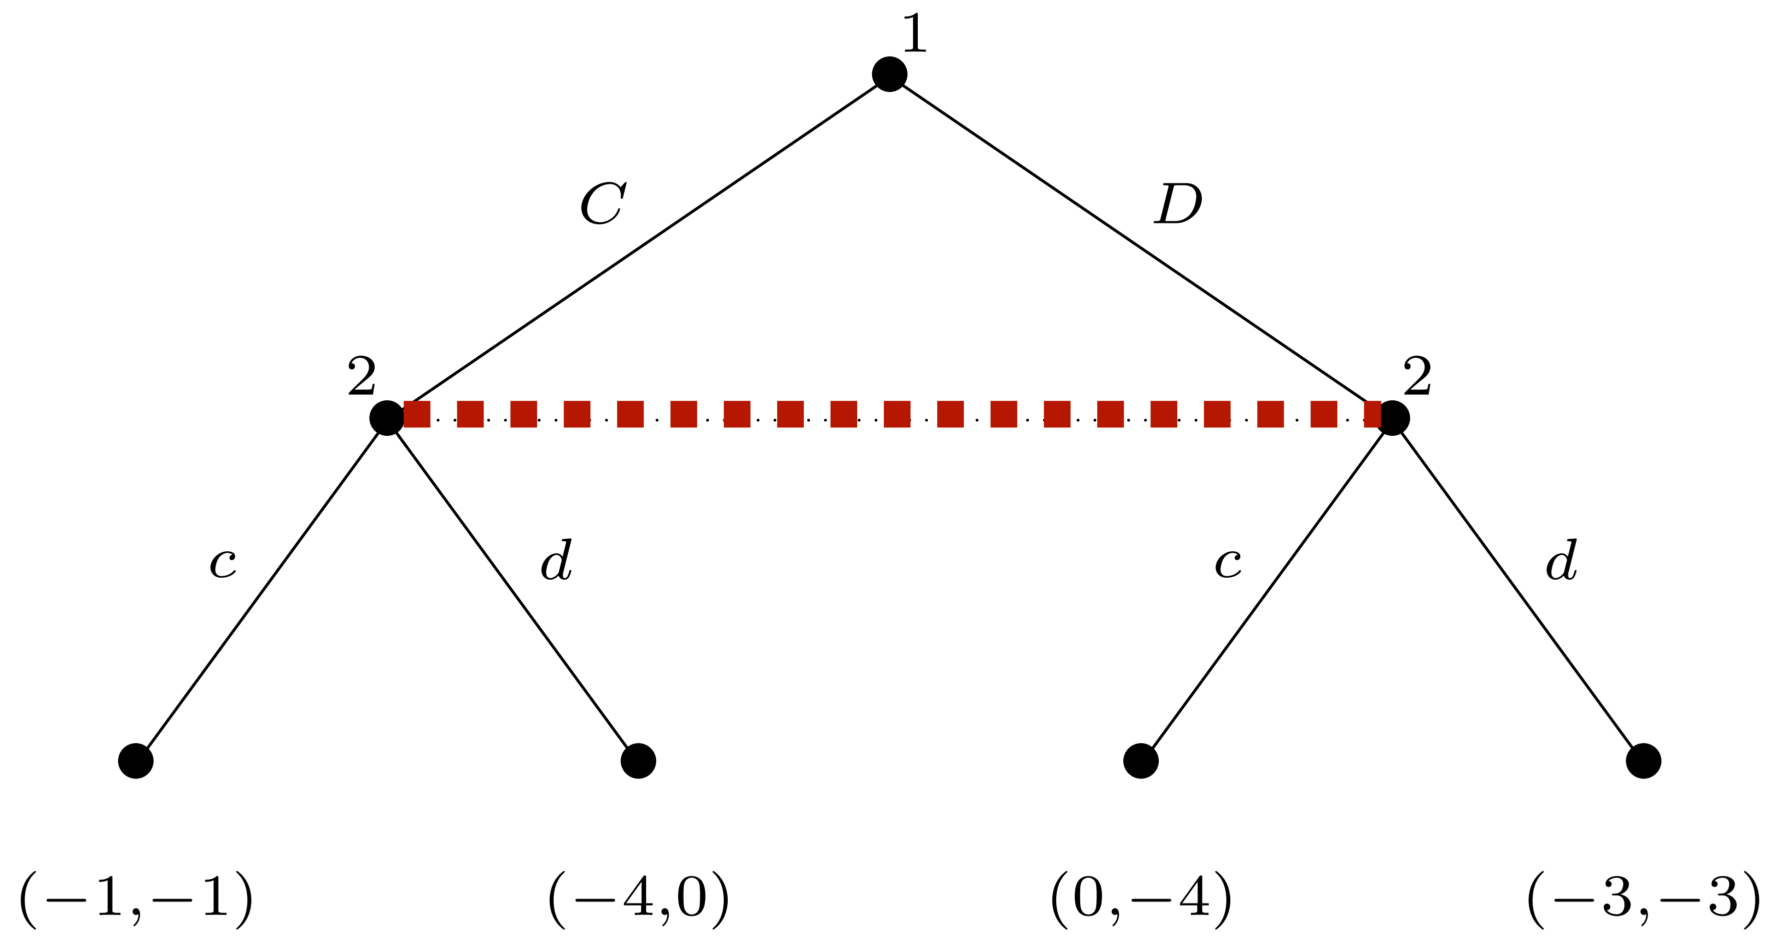

Unlike perfect information games, we can go in the opposite direction and represent **any normal form game** as an **imperfect information extensive form game**

Lecture Outline

1. Recap
2. Repeated Games
3. Infinitely Repeated Games
4. The Folk Theorem

Repeated Game

- Some situations are well-modelled as the **same agents** playing a normal-form game **multiple times**.
  - The normal-form game is the **stage game**; the whole game of playing the stage game repeatedly is a **repeated game.**
  - The stage game can be repeated a **finite** or an **infinite** number of times.
- Questions to consider:
  1. What do agents **observe**?
  2. What do agents **remember**?
  3. What is the agents' **utility** for the whole repeated game?

Finitely Repeated Game

Suppose that n players play a normal form game against each other $k \in \mathbb{N}$ times.

Questions:

1. Do they **observe** the other players' actions? If so, **when**?
2. Do they **remember** what happened in the previous games?
3. What is the **utility** for the whole game?
4. What are the **pure strategies**?

Representing Finitely Repeated Games

- Recall that we can represent normal form games as **imperfect information extensive form** games
- We can do the same for **repeated games**:

	c	d
C	-1,-1	-4,0
D	0,-4	-3,-3

and then

	c	d
C	-1,-1	-4,0
D	0,-4	-3,-3

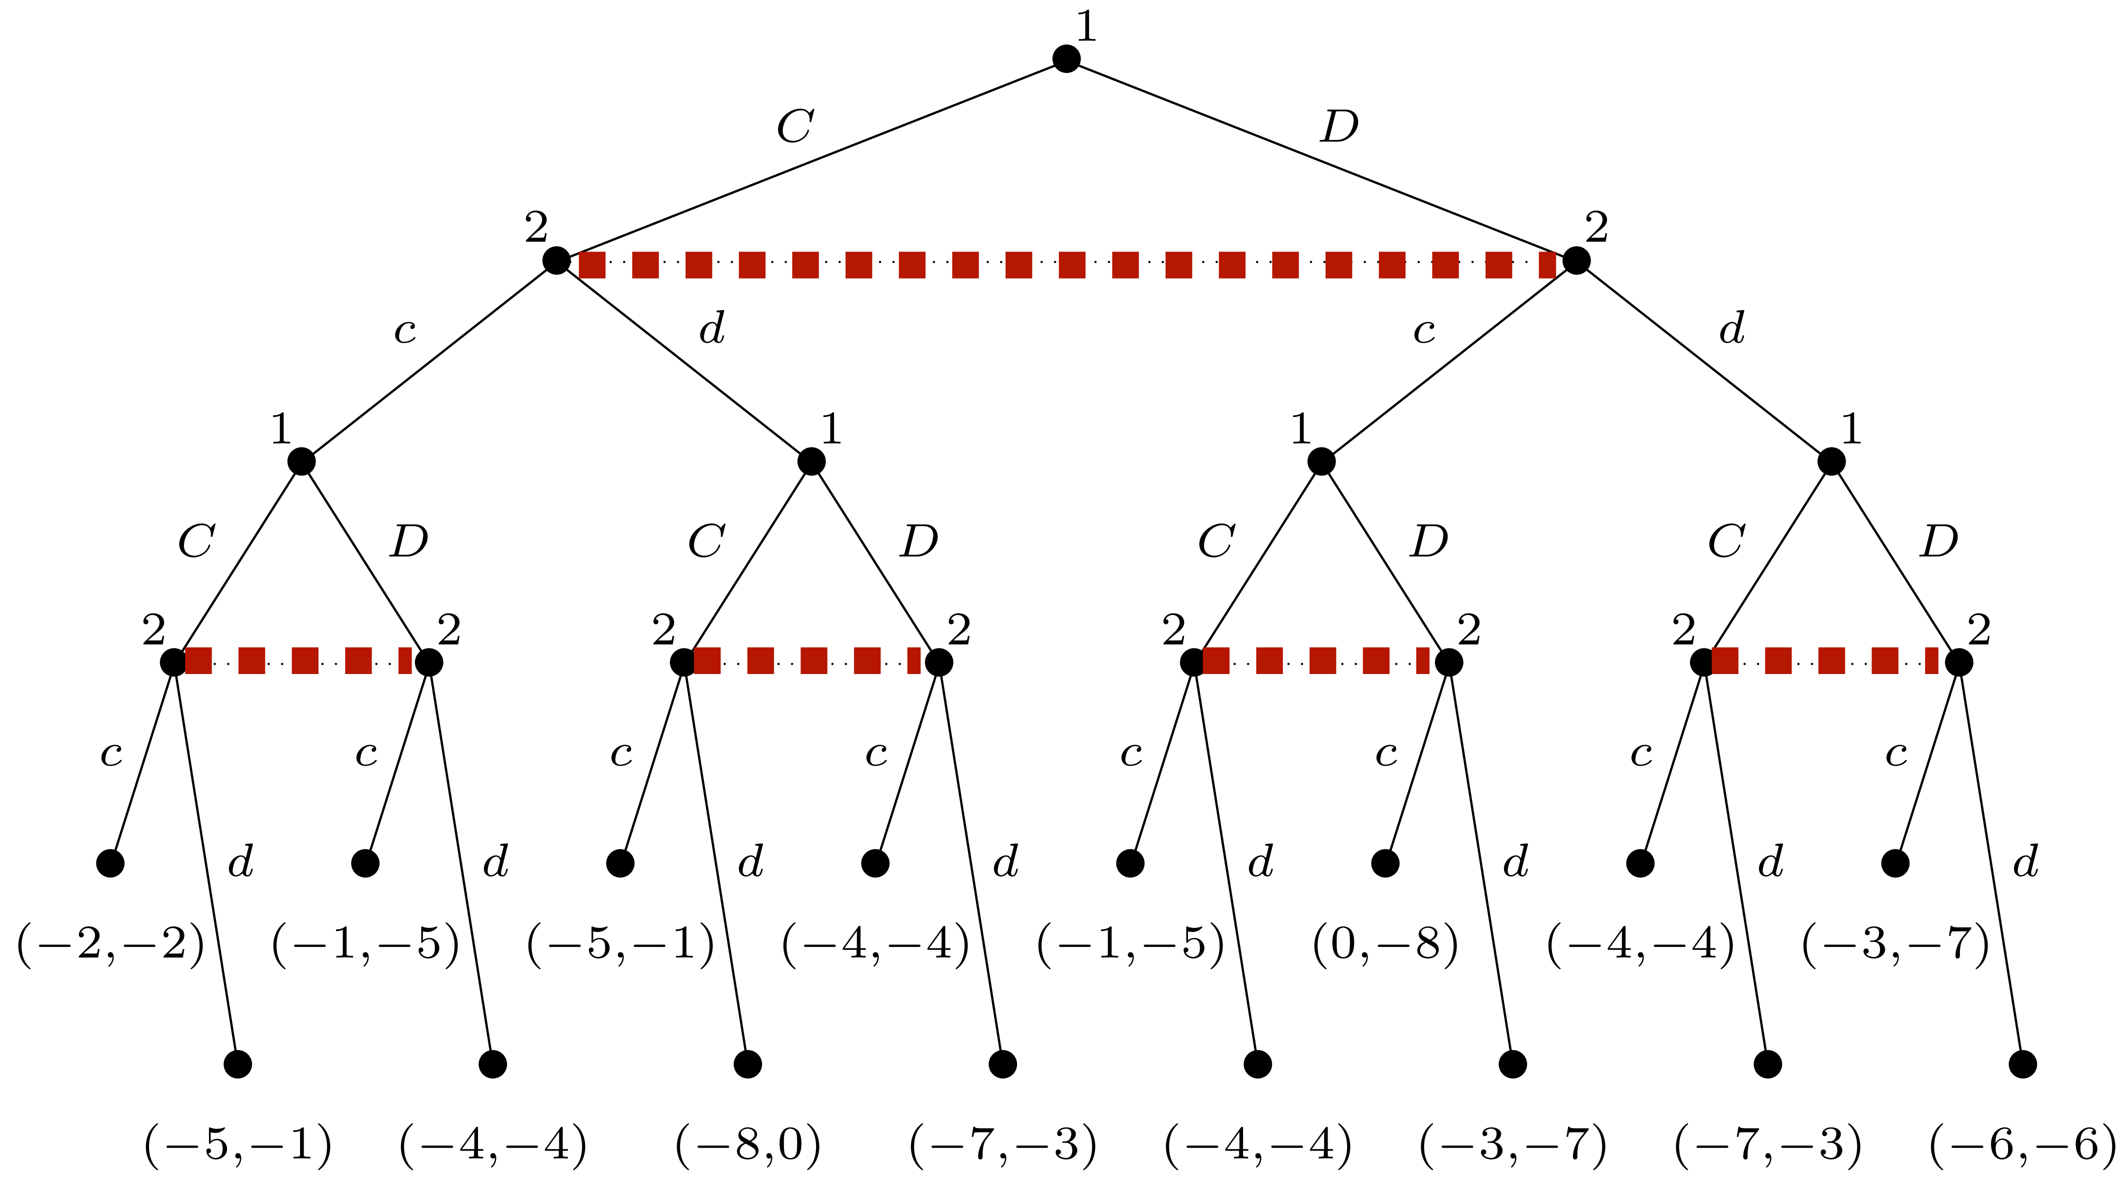

Fun (Repeated) Game

	c	d
C	-1,-1	-4,0
D	0,-4	-3,-3

and then

	c	d
C	-1,-1	-4,0
D	0,-4	-3,-3

and then

	c	d
C	-1,-1	-4,0
D	0,-4	-3,-3

and then

	c	d
C	-1,-1	-4,0
D	0,-4	-3,-3

and then

	c	d
C	-1,-1	-4,0
D	0,-4	-3,-3

- Play the **Prisoner's Dilemma** five times in a row against the **same person**
- Play at least two people

Properties of Finitely Repeated Games

- Playing an **equilibrium of the stage game** at every stage is an equilibrium of the repeated game (**why?**)
    - Instance of a **stationary** strategy
- In general, pure strategies can depend on the **previous history** (**why?**)
- **Question:** When the normal form game has a **strictly dominant strategy**, what can we say about the equilibrium of the finitely repeated game?

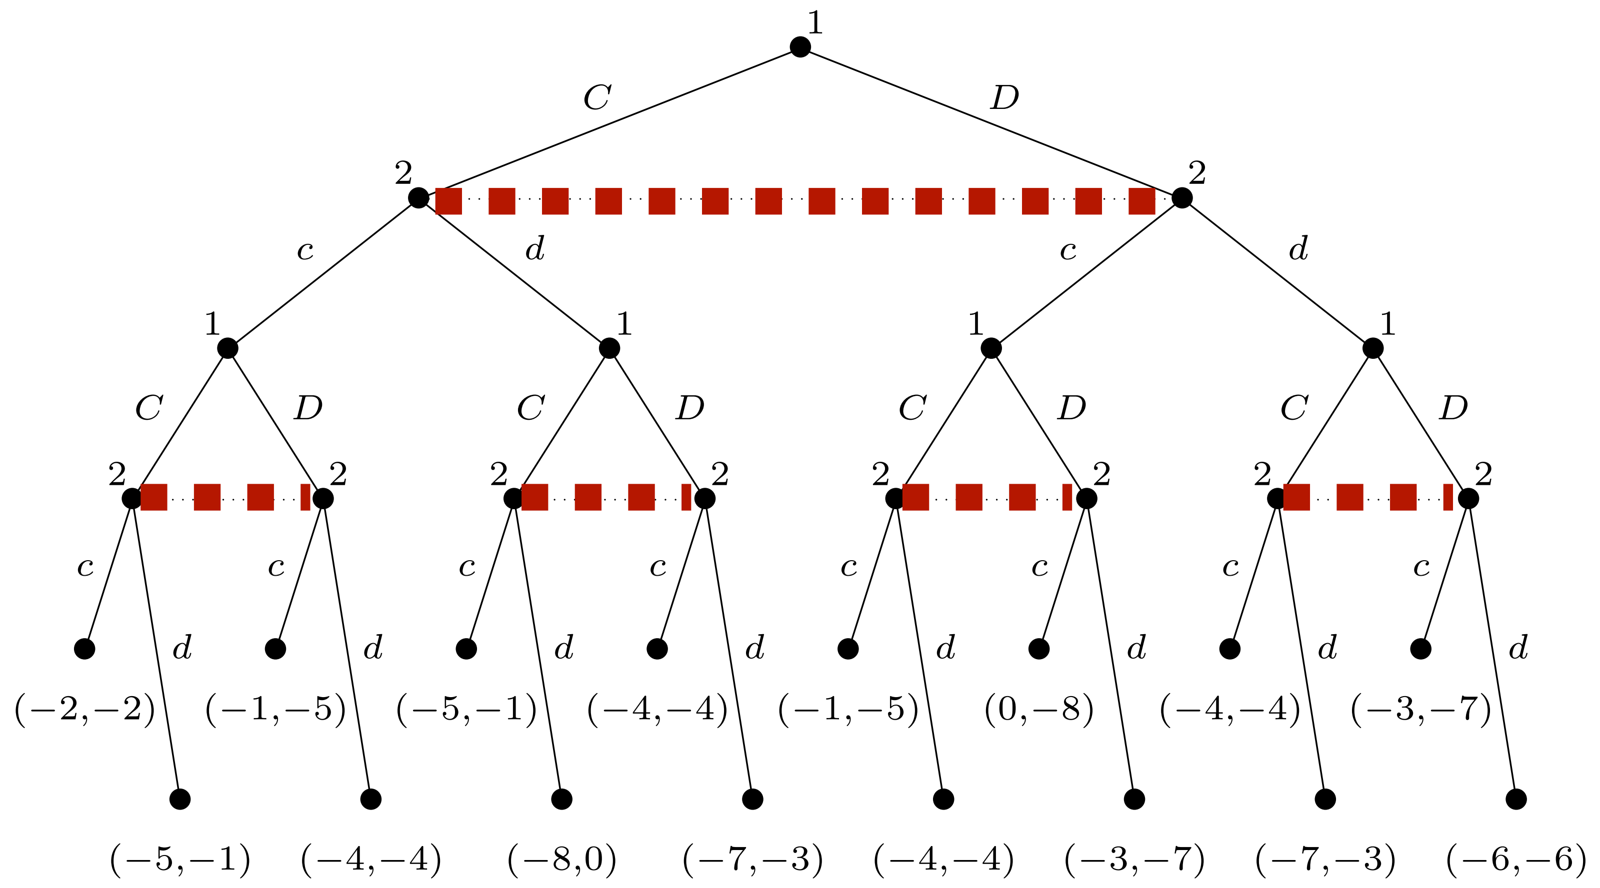

Infinitely Repeated Game

Suppose that n players play a normal form game against each other **infinitely many** times.

Questions:

1. Do they **remember** what happened in the previous games?
2. What is the **utility** for the whole game?
3. What are the **pure strategies**?
4. Can we write these games in the **imperfect information extensive form**?

Payoffs in Infinitely Repeated Games

- **Question:** What are the **payoffs** in an infinitely repeated game?
  - We cannot take the **sum of payoffs** in an infinitely repeated game (**why not?**)
  - We cannot put the overall utility on the **terminal nodes**, because there **aren't any**
- Two possible approaches:
  1. **Average reward:** Take the limit of the average reward to be the overall reward of the game
  2. **Discounted reward:** Apply a **discount factor** to future rewards to guarantee that they will converge

Average Reward

Definition:
Given an infinite sequence of payoffs $r_i^{(1)}, r_i^{(2)}, \ldots$ for player i, the **average reward** of i is

$$\lim_{t \to \infty} \frac{1}{T} \sum_{t=1}^{T} r_i^{(t)}.$$

- Problem: May not **converge** (**why**?)

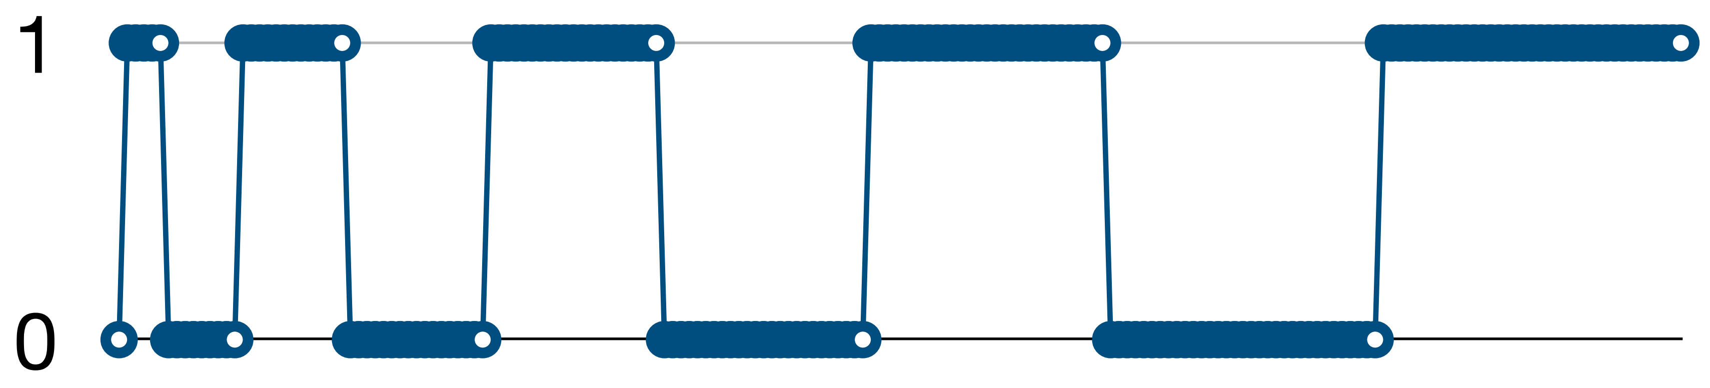

Discounted Reward

Definition:
Given an infinite sequence of payoffs $r_i^{(1)}, r_i^{(2)}, \ldots$ for player i, and a discount factor $0 \leq \beta \leq 1$, the **future discounted reward** of i is

$$\sum_{t=1}^{\infty} \beta^t r_i^{(t)}.$$

- Interpretations:
  1. Agent is **impatient**: cares more about rewards that they will receive earlier than rewards they have to wait for.
  2. Agent cares equally about all rewards, but at any given round the game will **stop** with probability $1 - \beta$.
- The two interpretations have **identical implications** for analyzing the game.

Strategy Spaces in Infinitely Repeated Games

Question: What is a pure strategy in an infinitely repeated game?

Definition:

For a stage game $G = (N, A, u)$, let

$$A^* = \{\emptyset\} \cup A^1 \cup A^2 \cup \cdots = \bigcup_{t=0}^{\infty} A^t$$

be the set of **histories** of the infinitely repeated game.

Then a **pure strategy** of the infinitely repeated game for an agent i is a mapping $s_i : A^* \to A_i$ from histories to player i's actions.

Equilibria in Infinitely Repeated Games

- **Question:** Are infinitely repeated games guaranteed to have **Nash equilibria**?
  - *Recall:* Nash's Theorem only applies to **finite games**
- Can we characterize the set of **equilibria** for an infinitely repeated game?
  - Can't build the induced normal form, there are **infinitely many pure strategies** (**why?**)
  - There could even be infinitely many **pure strategy Nash equilibria**! (**how?**)
- We can characterize the set of **payoff profiles** that are achievable in an equilibrium, instead of characterizing the equilibria themselves.

Enforceable

Definition:

Let $\underline{v}_i = \min_{s_{-i} \in S_{-i}} \max_{s_i \in S_i} u_i(s_i, s_{-i})$ be i's **minmax value** in $G = (N, A, u)$.

Then a payoff profile $r = (r_1, \ldots, r_n)$ is **enforceable** if $r_i \geq \underline{v}_i$ for all $i \in N$.

- A payoff vector is enforceable (on i) if the other agents working together can **ensure** that i's utility is no greater than r_i.

Feasible

Definition:
A payoff profile $r = (r_1, \ldots, r_n)$ is **feasible** if there exist **rational**, non-negative values $\{\alpha_a \mid a \in A\}$ such that for all $i \in N$,

$$r_i = \sum_{a \in A} \alpha_a u_i(a),$$

with $\sum_{a \in A} \alpha_a = 1.$

- A payoff profile is feasible if it is a (rational) **convex combination** of the **outcomes** in G.

Folk Theorem

Theorem:
Consider any n-player normal form game G and payoff profile $r = (r_1, \ldots, r_n)$.

1. If r is the payoff profile for any Nash equilibrium of the infinitely repeated G with average rewards, then r is **enforceable**.
2. If r is both **feasible** and **enforceable**, then r is the payoff profile for some Nash equilibrium of the infinitely repeated G with average rewards.

- Whole family of similar proofs for discounted rewards case, subgame perfect equilibria, **real** convex combinations, etc.

Folk Theorem Proof Sketch:

Nash $\implies$ Enforceable

- Suppose for contradiction that r is **not** enforceable, but r is the payoff profile in a Nash equilibrium s^* of the infinitely repeated game.
- Consider the strategy $s_i'(h) \in BR_i(s_{-i}^*(h))$ for each $h \in A^*$.
- Player i receives **at least** $\underline{v}_i > r_i$ in every stage game by playing strategy s_i' (**why**?)
- So strategy s_i' is a **utility-increasing deviation** from s^*, and hence s^* is not an equilibrium.

Folk Theorem Proof Sketch: Enforceable & Feasible $\implies$ Nash

- Suppose that r is both feasible and enforceable.
- We can construct a strategy profile s^* that visits each action profile a with frequency α_a (since α_a's are all rational).
- At every history where a player i has not played their part of the cycle, **all of the other players** switch to playing the **minmax strategy against *i*** (this is called a GRIM TRIGGER strategy)
  - That makes i's overall utility for the game $\underline{v}_i \leq r_i$ for any deviation s_i'. (**why**?)
  - Thus there is no utility-increasing deviation for i.

Summary

- A **repeated game** is one in which agents play the same normal form game (the **stage game**) multiple times.
- **Finitely repeated:** Can represent as an **imperfect information extensive form** game.
- **Infinitely repeated:** Life gets more complicated
    - Payoff to the game: either **average** or **discounted** reward
    - Pure strategies map from **entire previous history** to action
- **Folk theorem** characterizes which **payoff profiles** can arise in any equilibrium
    - All profiles that are both **enforceable** and **feasible.**