Game Theory for Sequential Interactions

CMPUT 366: Intelligent Systems

S&LB §5.0-5.2.2

Lecture Outline

1. Recap
2. Perfect Information Games
3. Backward Induction
4. Imperfect Information Games

Recap: Game Theory

- Game theory studies the **interactions of rational agents**
  - Canonical representation is the **normal form game**
- Game theory uses **solution concepts** rather than optimal behaviour
  - "Optimal behaviour" is not clear-cut in multiagent settings
  - **Pareto optimal:** no agent can be made better off without making some other agent worse off
  - **Nash equilibrium:** no agent regrets their strategy given the choice of the other agents' strategies

	Ballet	Soccer
Ballet	2, 1	0, 0
Soccer	0, 0	1, 2

Mixed Strategies

Definitions:

- A **strategy** s_i for agent i is any probability distribution over the set A_i, where each action a_i is played with probability $s_i(a_i)$.
  - **Pure strategy:** only a single action is played
  - **Mixed strategy:** randomize over multiple actions
- Set of i's strategies: $S_i \doteq \Delta(A_i)$ ← $\Delta(X)$ = "set of distributions over elements of X"
- Set of **strategy profiles:** $S = S_1 \times S_2 \times \cdots \times S_n$
- **Utility** of a mixed strategy profile:

$$u_i(s) \doteq \sum_{a \in A} u_i(a) \prod_{j \in N} s_j(a_j)$$

Best Response and Nash Equilibrium

Definition:

The set of i's **best responses** to a strategy profile $s \in S$ is

$$BR_i(s_{-i}) \doteq \{a_i^* \in A_i \mid u_i(a_i^*, s_{-i}) \geq u_i(a_i, s_{-i}) \;\; \forall a_i \in A_i\}$$

Definition:

A strategy profile $s \in S$ is a **Nash equilibrium** iff

$$\forall i \in N, a_i \in A_i \quad s_i(a_i) > 0 \implies a_i \in BR_{-i}(s_{-i})$$

- When at least one s_i is mixed, s is a **mixed strategy Nash equilibrium**

Nash's Theorem

Theorem: [Nash 1951]
Every game with a finite number of players and action profiles has at least one Nash equilibrium.

- **Pure strategy** equilibria are *not* guaranteed to exist

Interpreting Mixed Strategy Nash Equilibrium

What does it even mean to say that agents are playing a mixed strategy Nash equilibrium?

- They truly are **sampling a distribution** in their heads, perhaps to **confuse** their opponents (e.g., soccer, other zero-sum games)
- The distribution represents the **other agents' uncertainty** about what the agent will do
- The distribution is the **empirical frequency** of actions in repeated play
- The distribution is the frequency of a pure strategy in a **population** of pure strategies (i.e., every individual plays a pure strategy)

Extensive Form Games

- Normal form games don't have any notion of **sequence**: all actions happen **simultaneously**
- The **extensive form** is a game representation that explicitly includes temporal structure (i.e., a **game tree**)

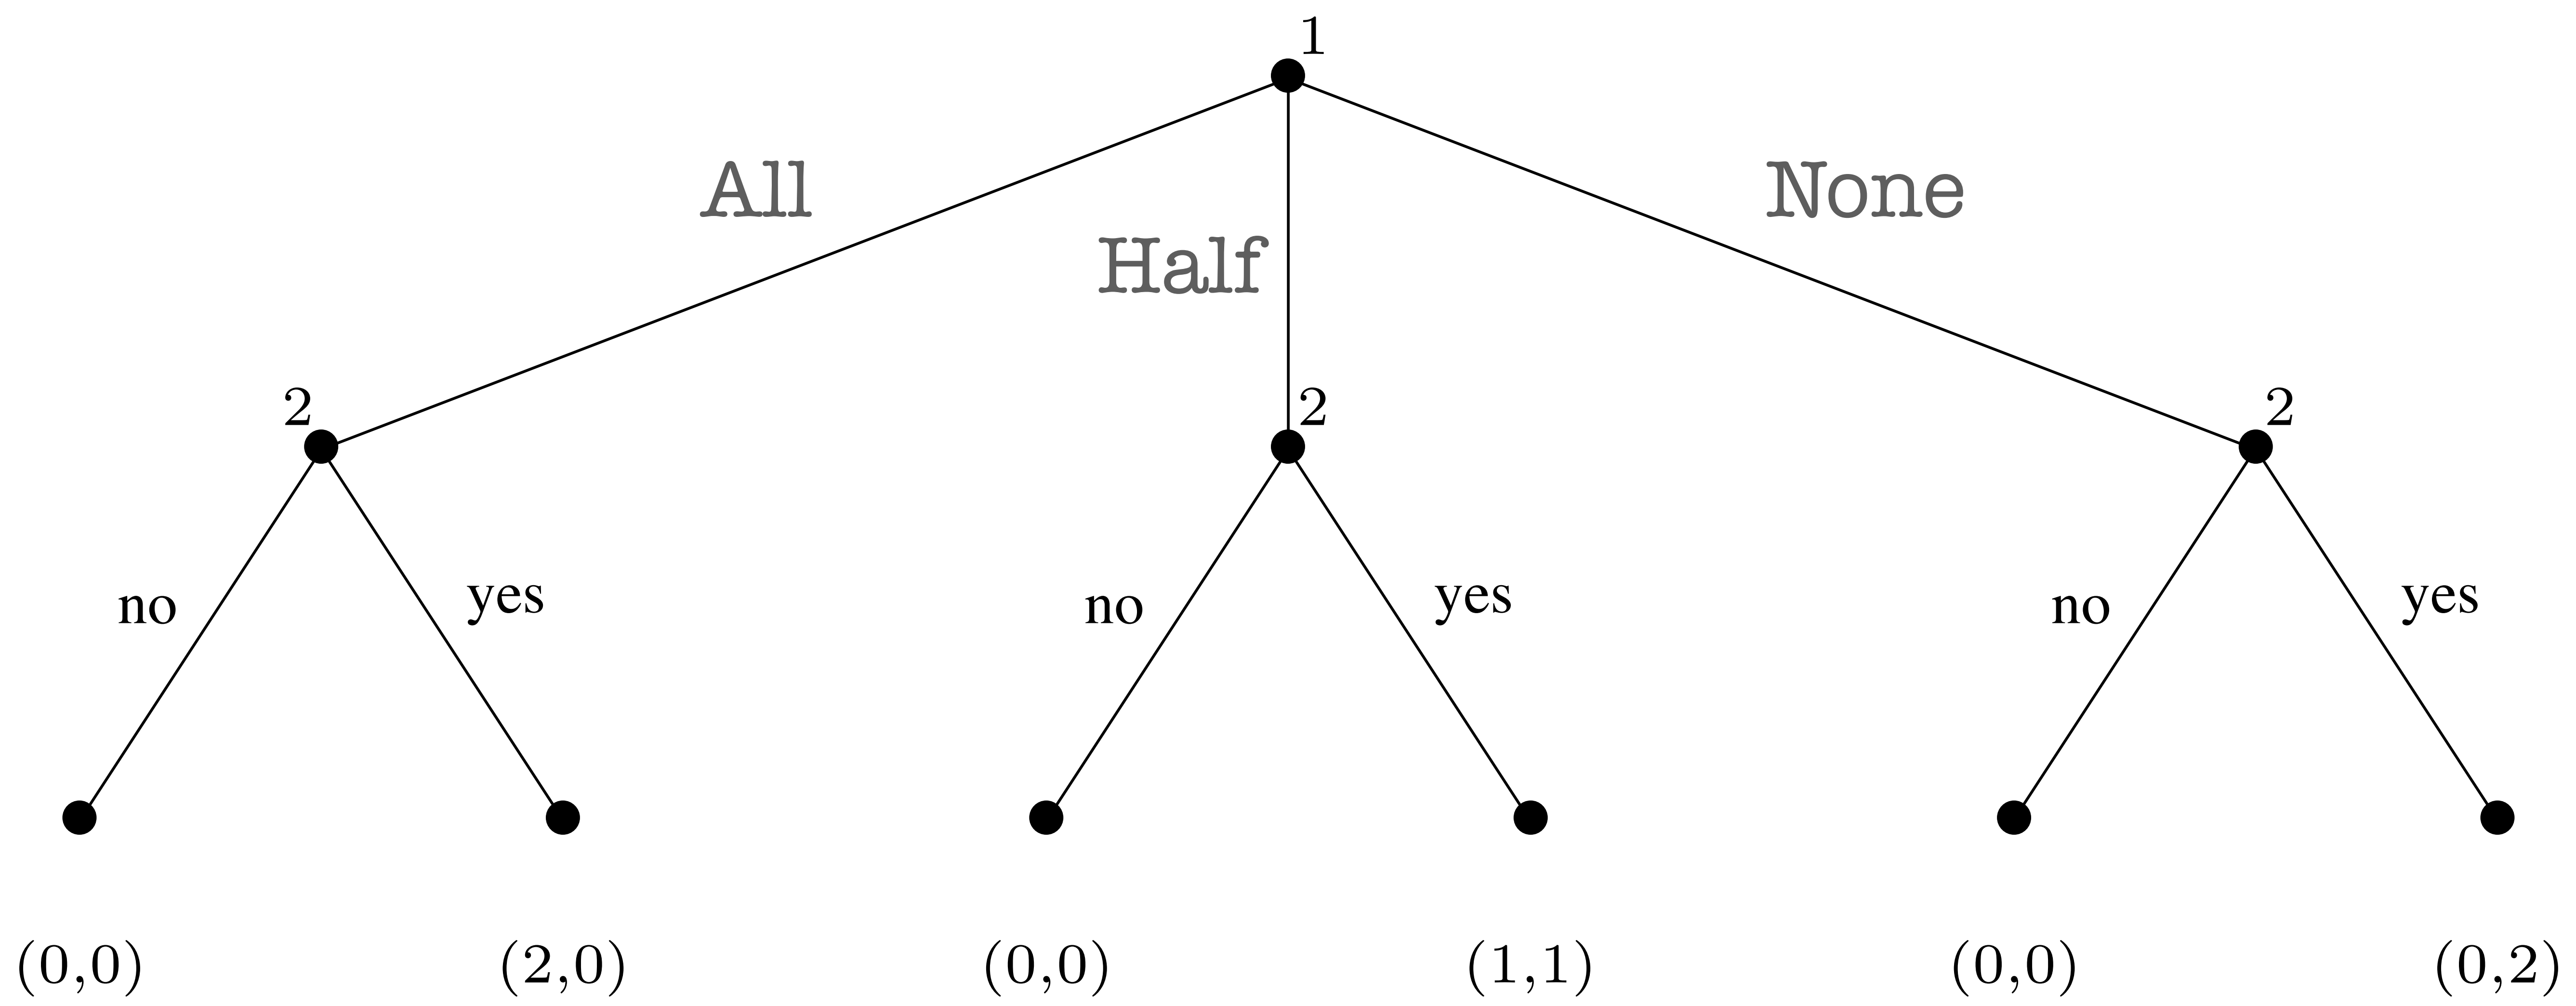

Perfect Information

There are two kinds of extensive form game:

1. **Perfect information:** Every agent **sees all actions** of the other players (including "**Nature**")
    - e.g.: Chess, checkers, Pandemic
2. **Imperfect information:** Some actions are **hidden**
    - Players may not know exactly where they are in the tree
    - e.g.: Poker, rummy, Scrabble

Perfect Information Extensive Form Game

Definition:
A **finite perfect-information game in extensive form** is a tuple $G = (N, A, H, Z, \chi, \rho, \sigma, u)$, where

- N is a set of n **players**,
- A is a single set of **actions**,
- H is a set of nonterminal **choice nodes**,
- Z is a set of **terminal nodes** (disjoint from H),
- $\chi : H \rightarrow 2^A$ is the **action function**,
- $\rho : H \rightarrow N$ is the **player function**,
- $\sigma : H \times A \rightarrow H \cup Z$ is the **successor function**,
- $u = (u_1, u_2, \ldots, u_n)$ is a **utility function** for each player, $u_i : Z \rightarrow \mathbb{R}$

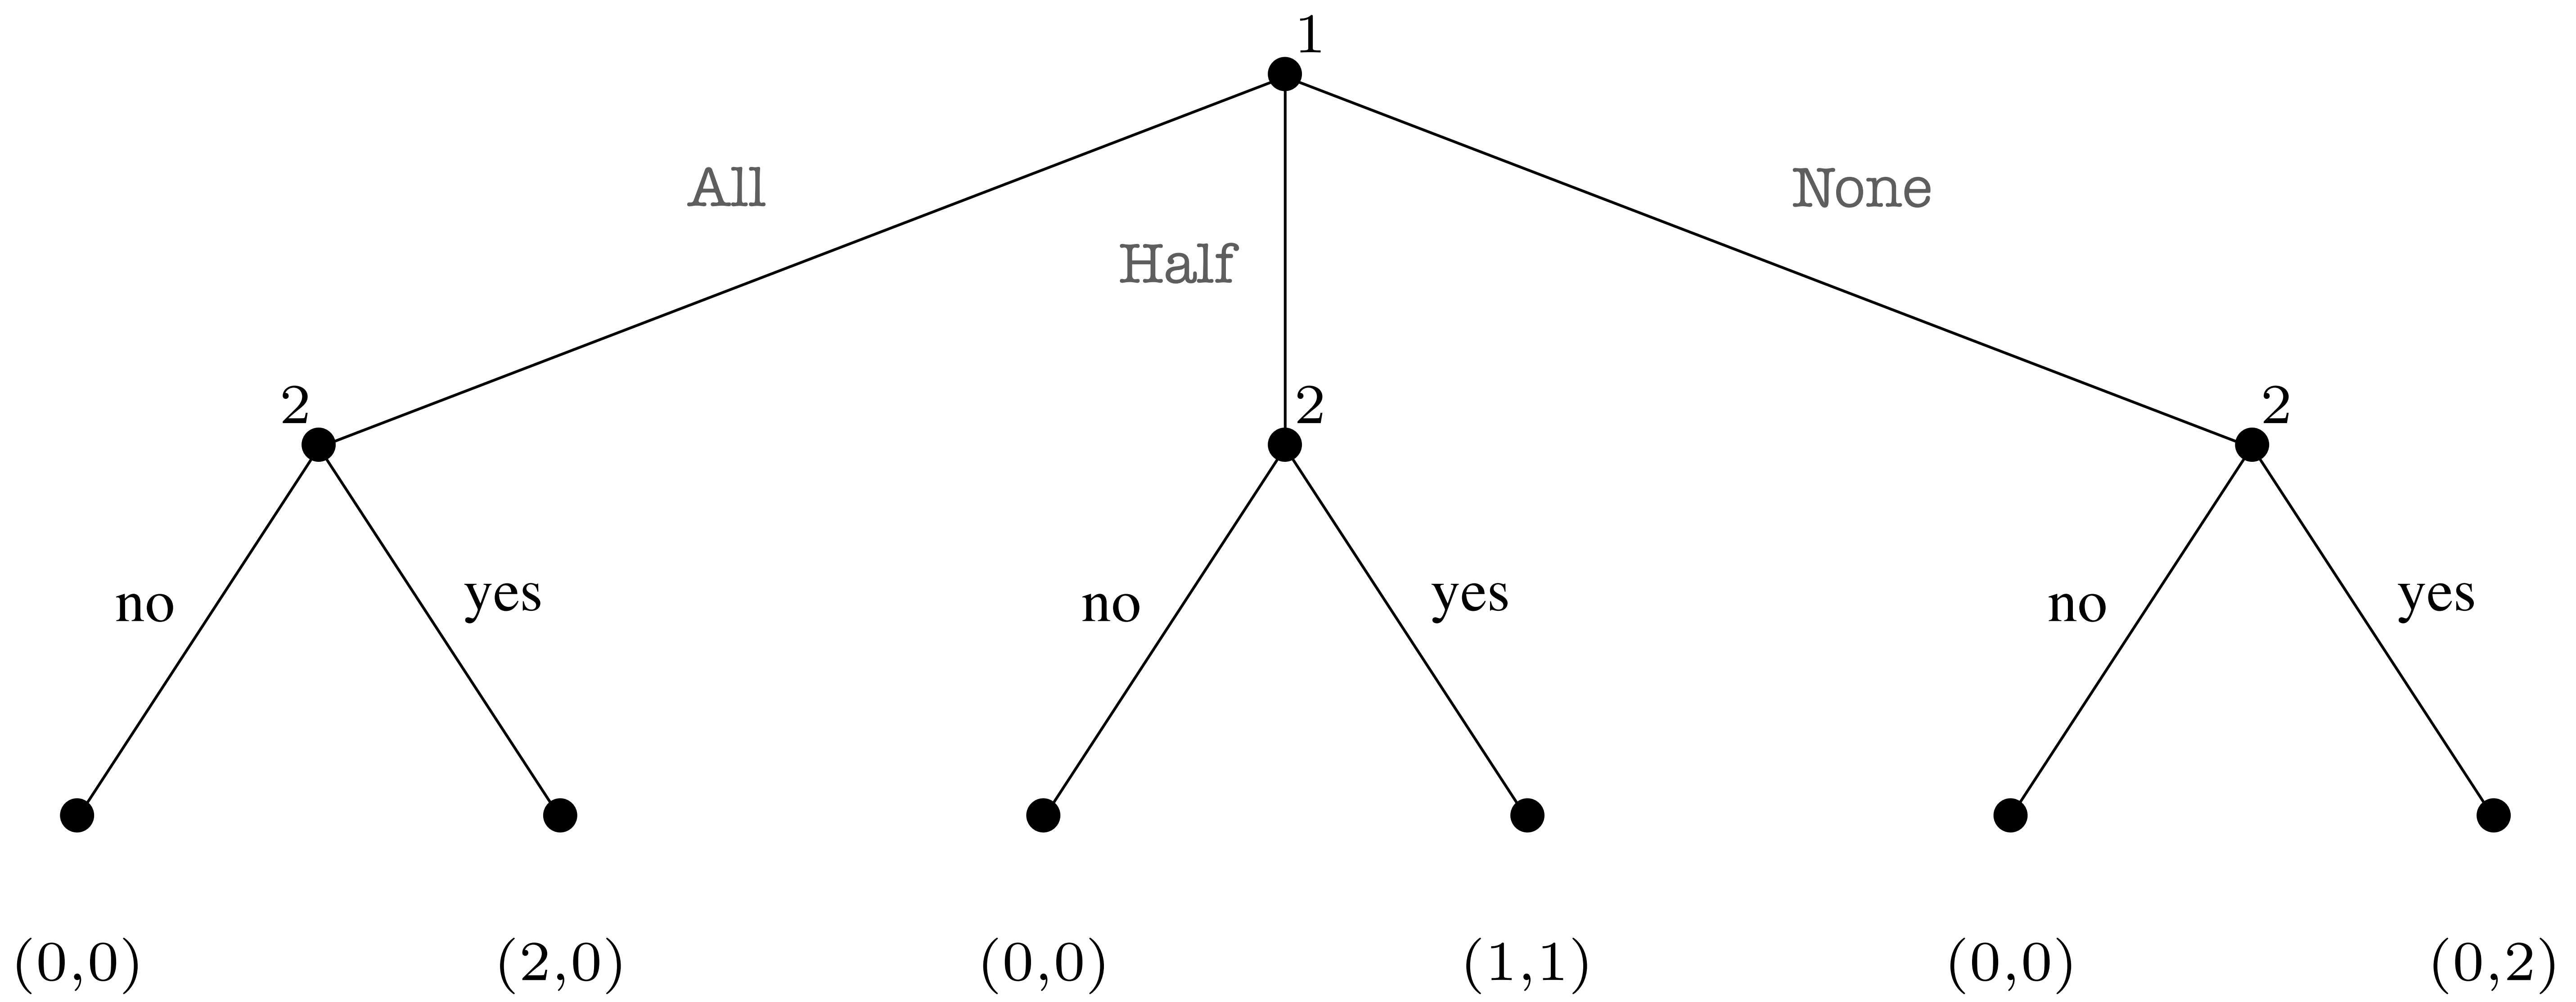

Fun Game:
The Sharing Game

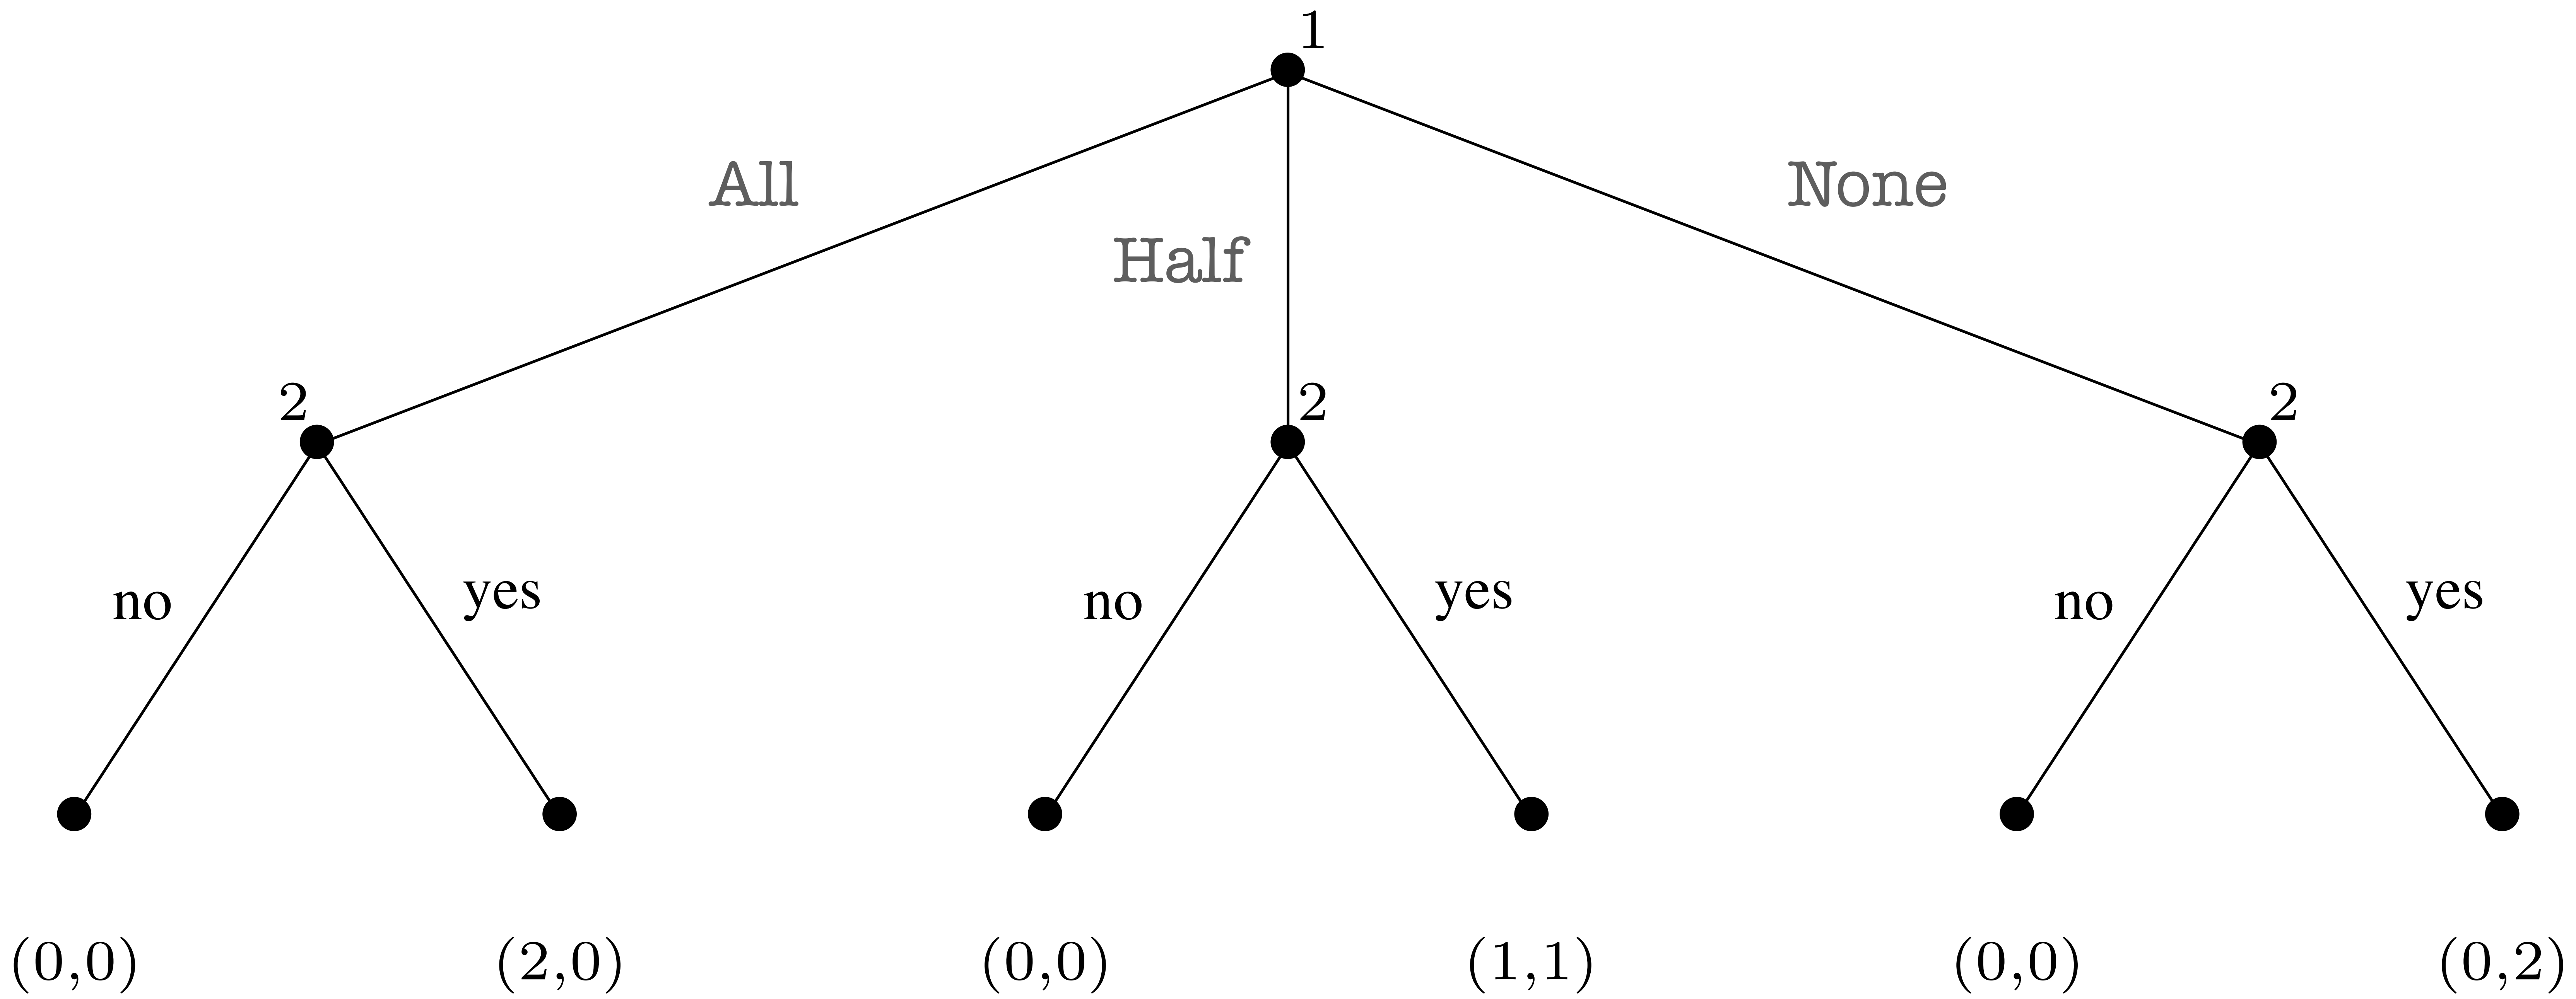

- Two siblings must decide how to share two $100 coins
- Sibling 1 suggests a division, then sibling 2 **accepts** or **rejects**
    - If **rejected**, nobody gets any coins.

Pure Strategies

Question: What are the **pure strategies** in an extensive form game?

Definition:
Let $G = (N, A, H, Z, \chi, \rho, \sigma, u)$ be a perfect information game in extensive form. Then the **pure strategies of player** i consist of the cross product of actions available to player i at each of their choice nodes, i.e.,

$$\prod_{h \in H \mid \rho(h) = i} \chi(h)$$

- A pure strategy associates an action with **each** choice node, even those that will **never be reached**
  - Even nodes that will never be reached as a result of the strategy itself!

Pure Strategies Example

Question: What are the **pure strategies** for **player 2**?

- $\{(C,E),(C,F),(D,E),(D,F)\}$

Question: What are the **pure strategies** for **player 1**?

- $\{(A,G),(A,H),(B,G),(B,H)\}$
- Note that these associate an action with the second choice node even when it can never be reached

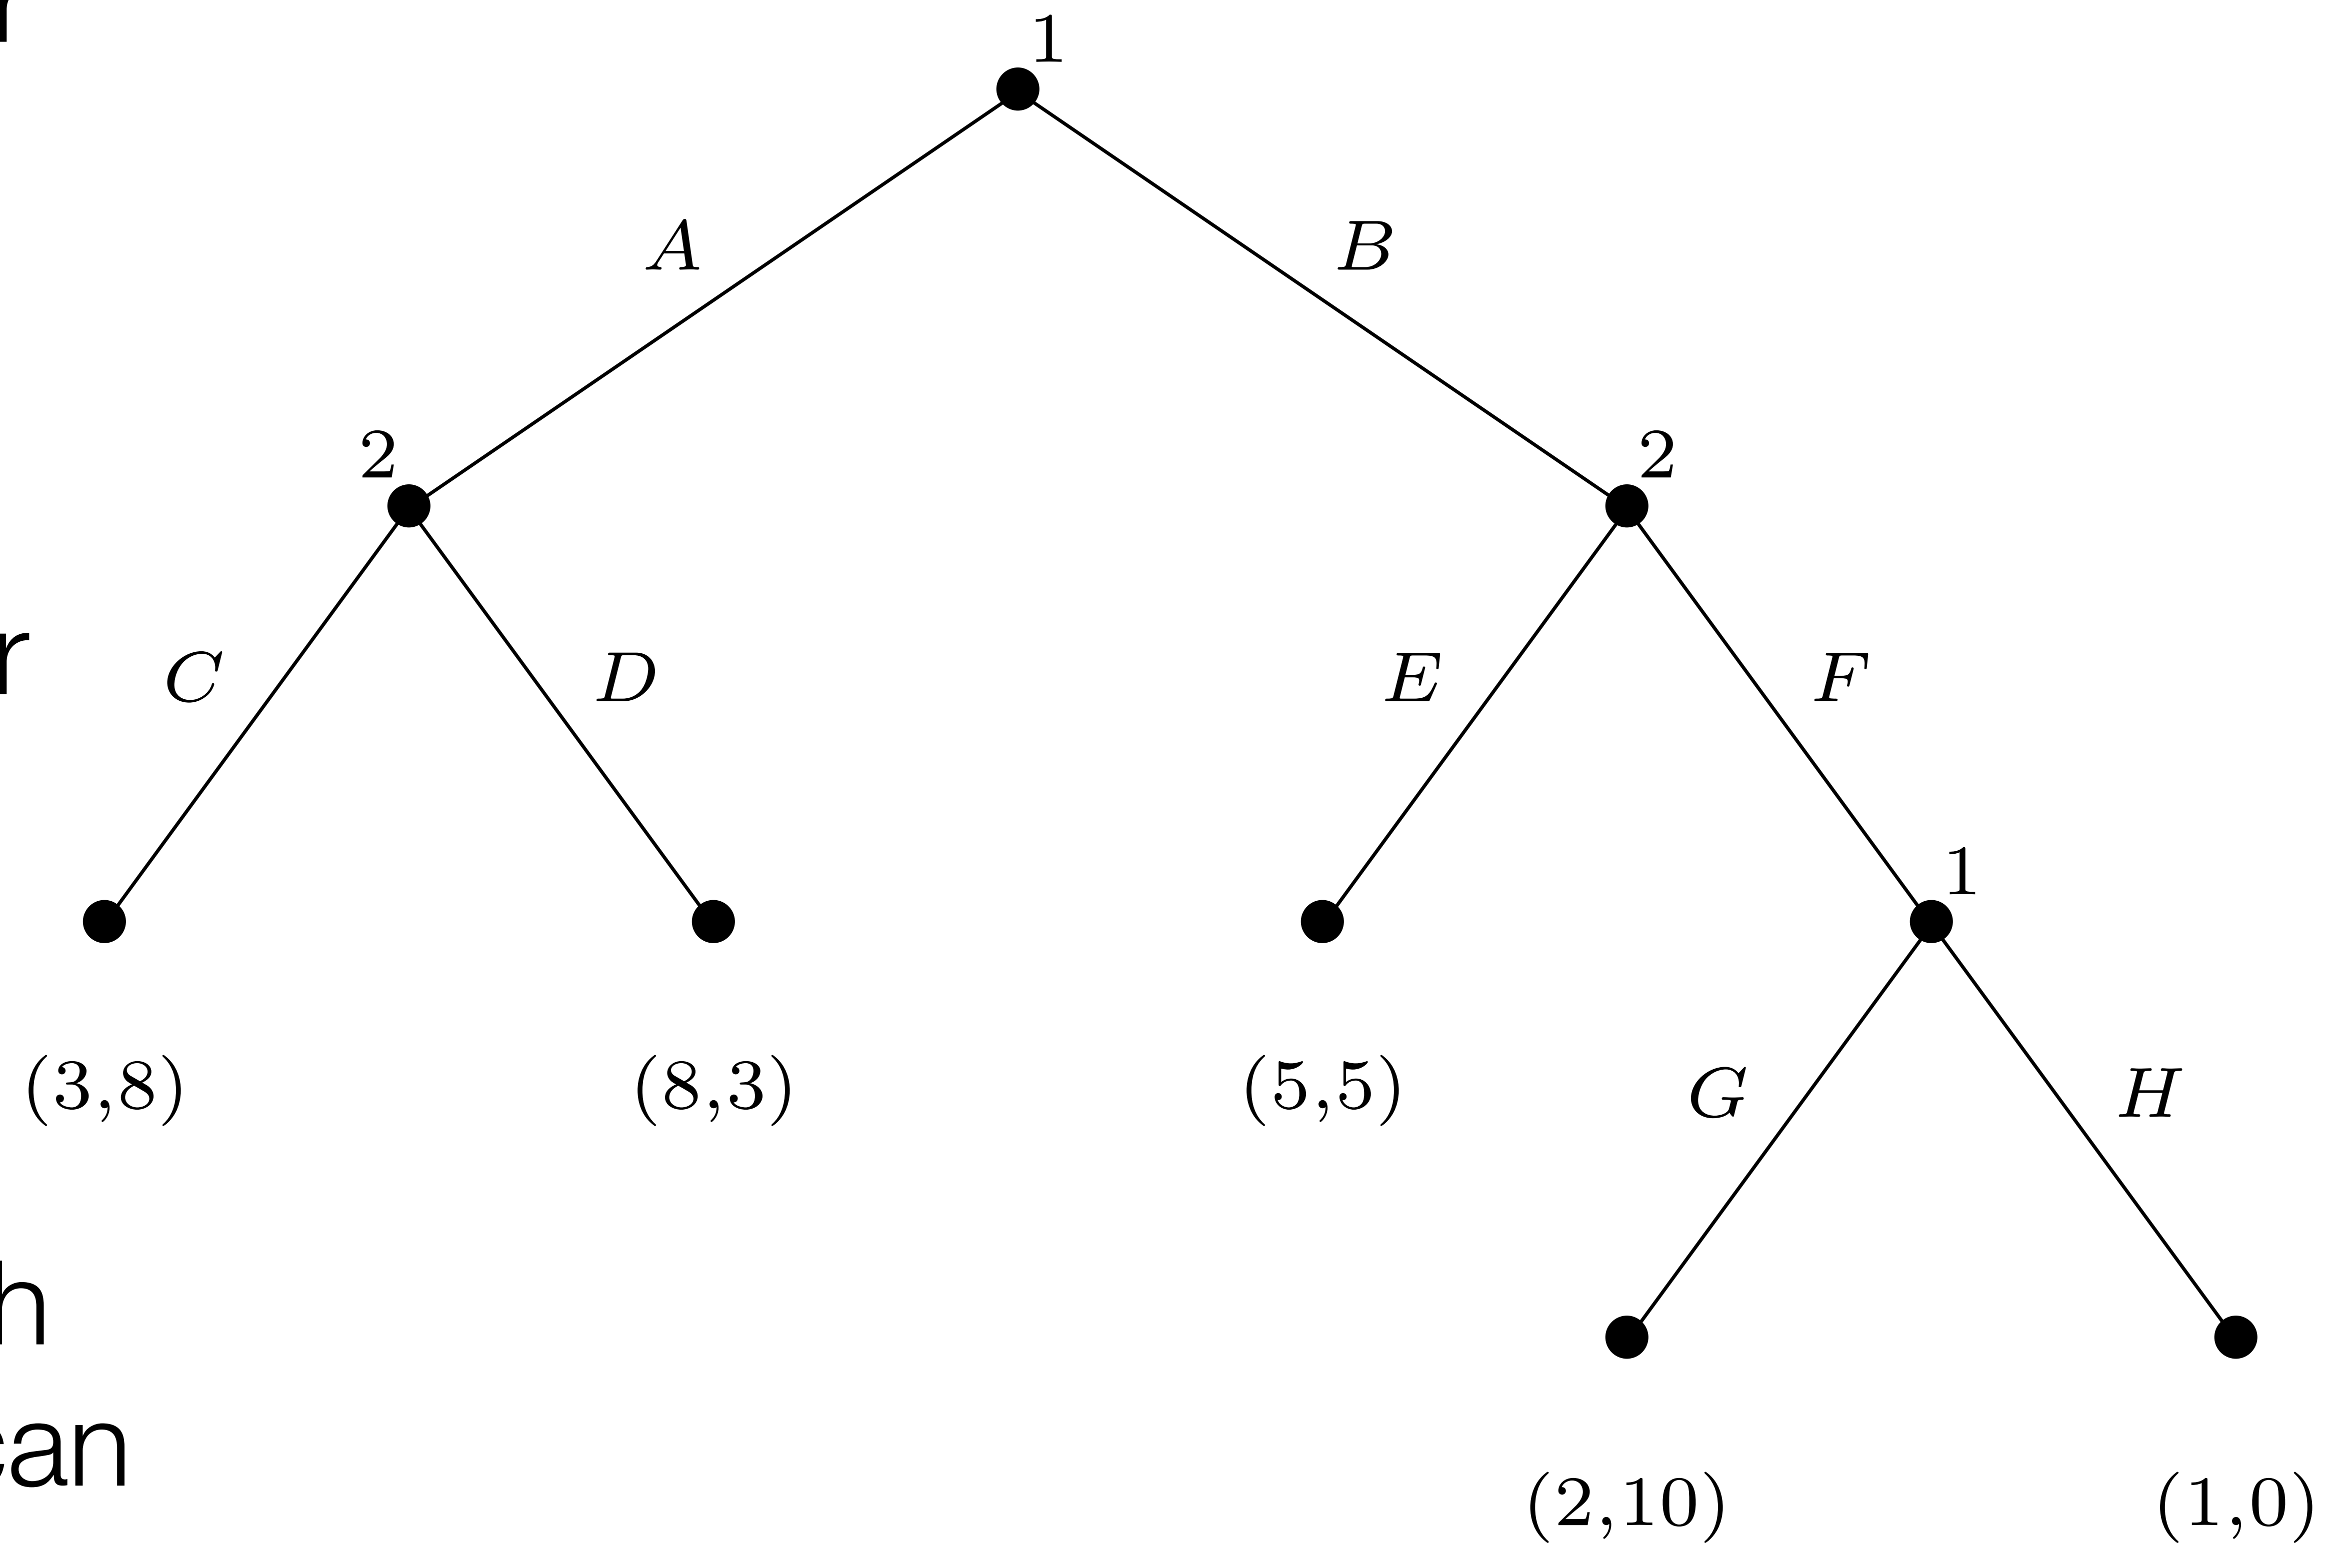

Induced Normal Form

Question:
Which representation is more **compact**?

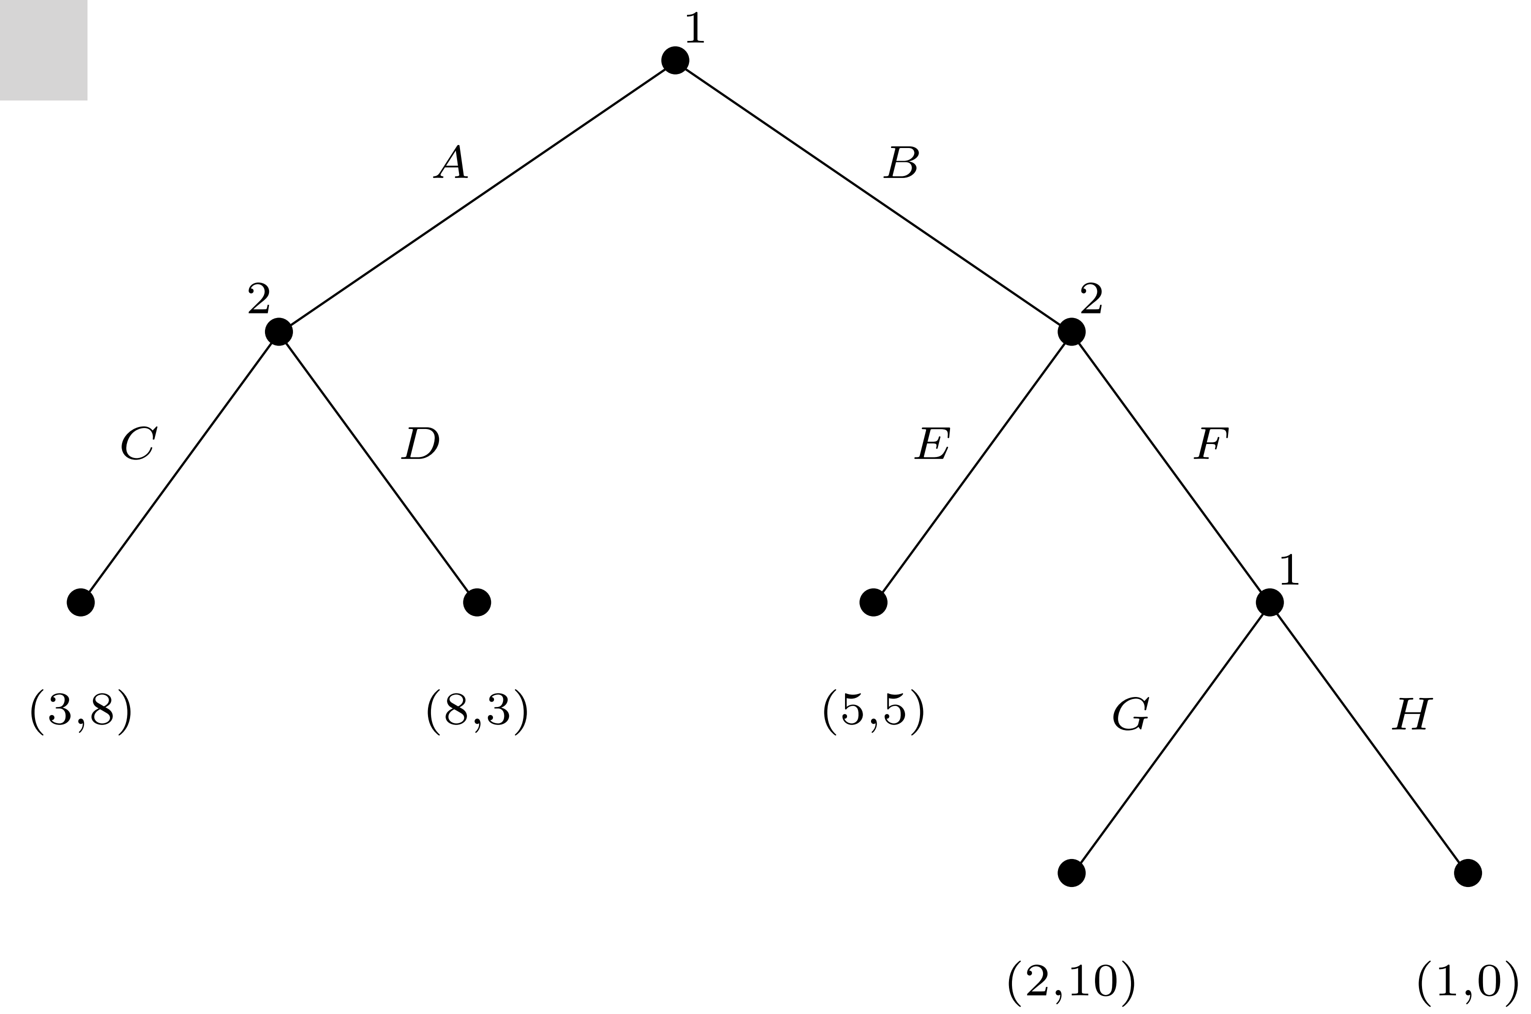

	C,E	C,F	D,E	D,F
A,G	3,8	3,8	8,3	8,3
A,H	3,8	3,8	8,3	8,3
B,G	5,5	2,10	5,5	2,10
B,H	5,5	1,0	5,5	1,0

- Any pair of pure strategies uniquely identifies a **terminal node**, which identifies a **utility** for each agent (**why?**)
- We have now defined a set of **agents**, **pure strategies**, and **utility functions**
- Any extensive form game defines a corresponding **induced normal form game**

Reusing Old Definitions

- We can plug our new definition of **pure strategy** into our existing definitions for:
  - Mixed strategy
  - Best response
  - Nash equilibrium (both pure and mixed strategy)

Question:

What is the definition of a **mixed strategy** in an extensive form game?

Pure Strategy Nash Equilibria

Theorem: [Zermelo, 1913]
Every finite perfect-information game in extensive form has at least one **pure strategy Nash equilibrium**.

- Starting from the bottom of the tree, no agent needs to **randomize**, because they already know the best response
- There might be **multiple** pure strategy Nash equilibria in cases where an agent has multiple best responses at a **single choice node**

Backward Induction

- **Backward induction** is a straightforward algorithm that is guaranteed to compute a pure strategy Nash equilibrium.
- **Idea:** Replace subgames lower in the tree with their equilibrium values

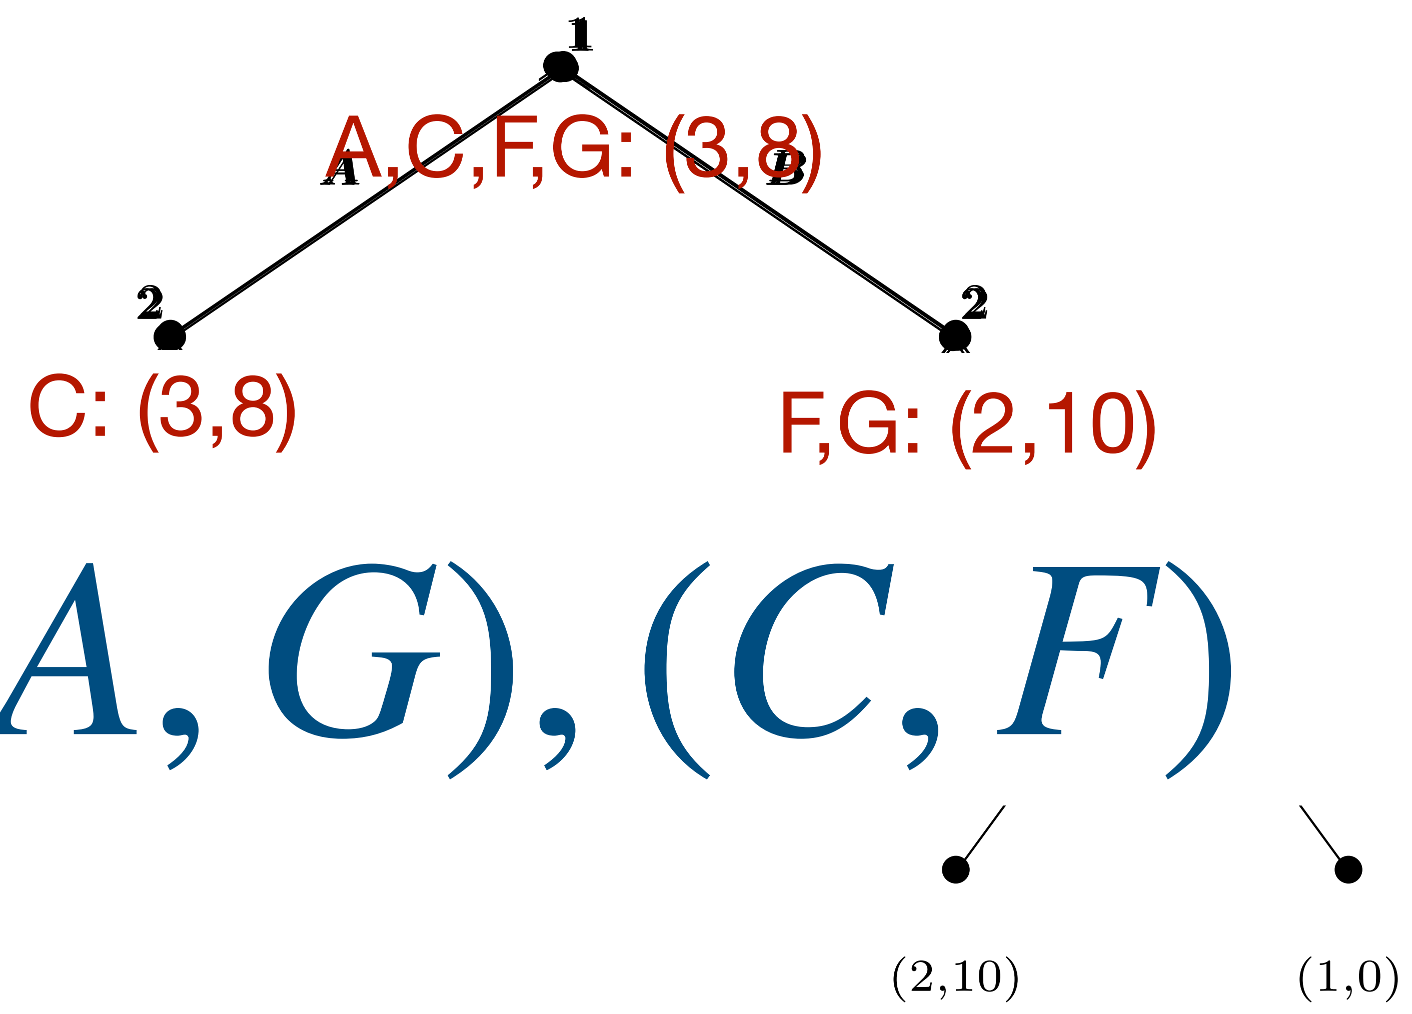

$$(A, G), (C, F)$$

Fun Game: Centipede

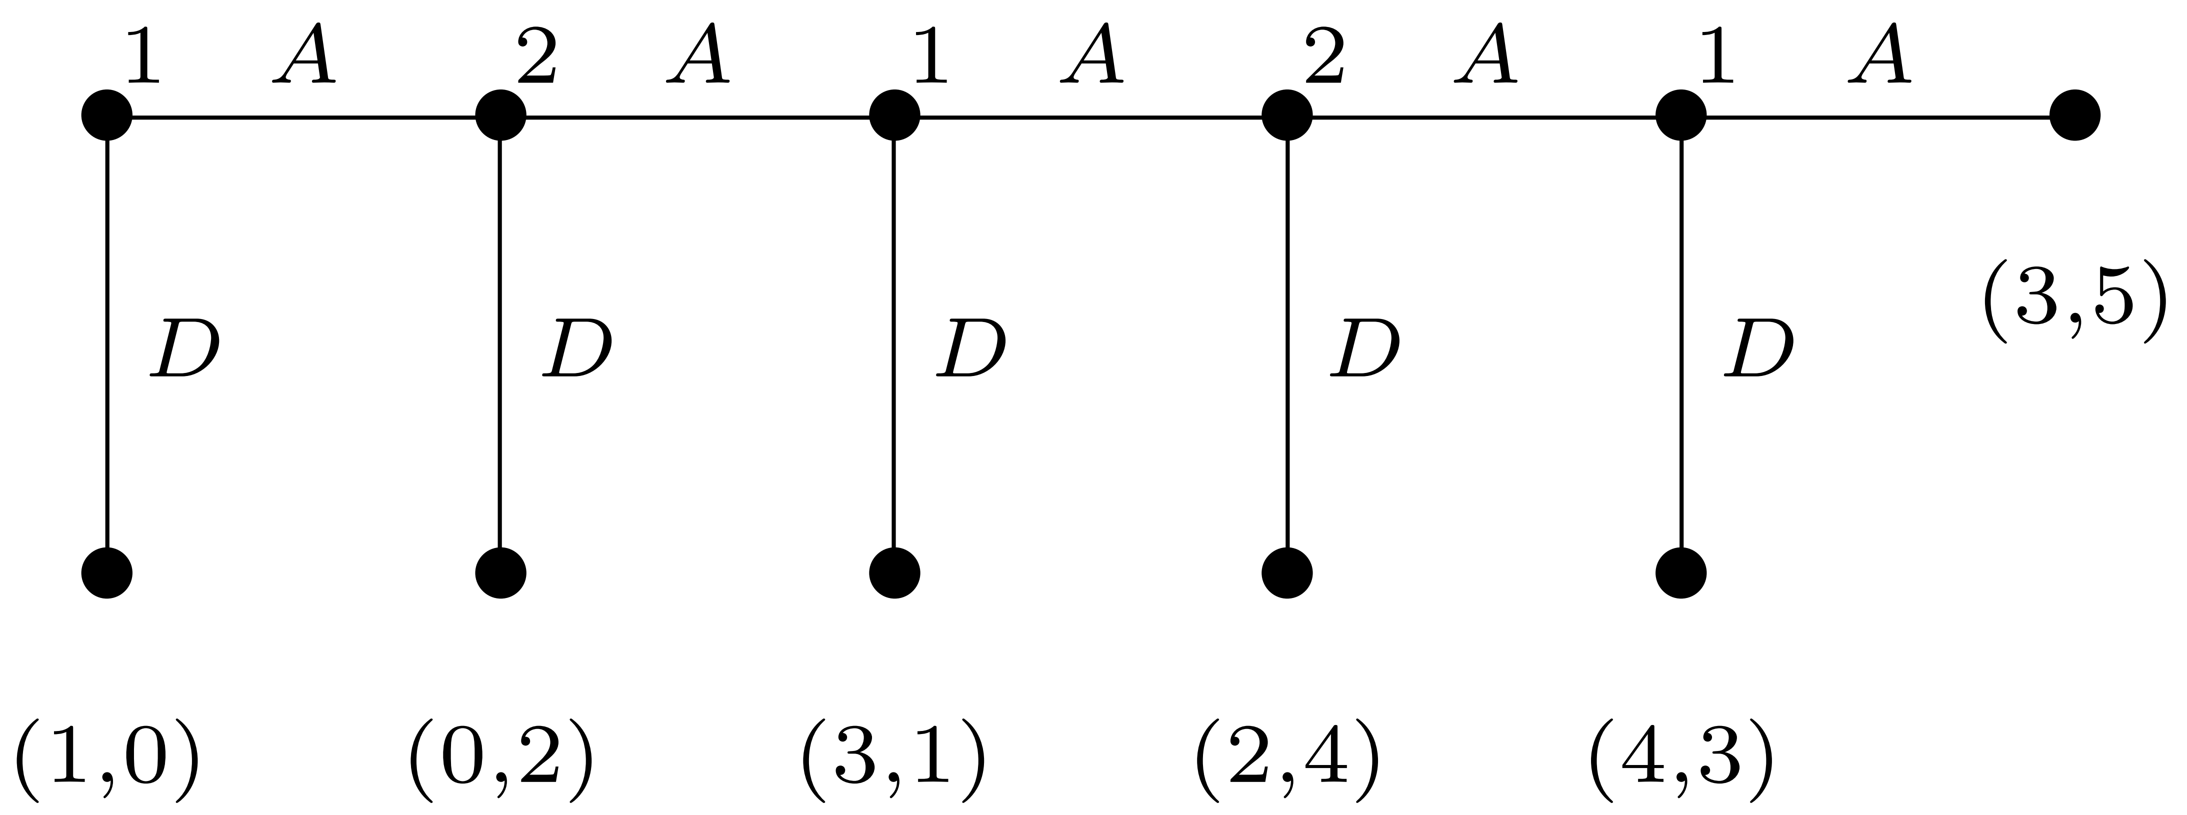

- At each stage, one of the players can go **Across** or **Down**
- If they go **Down**, the game ends.

Backward Induction Criticism

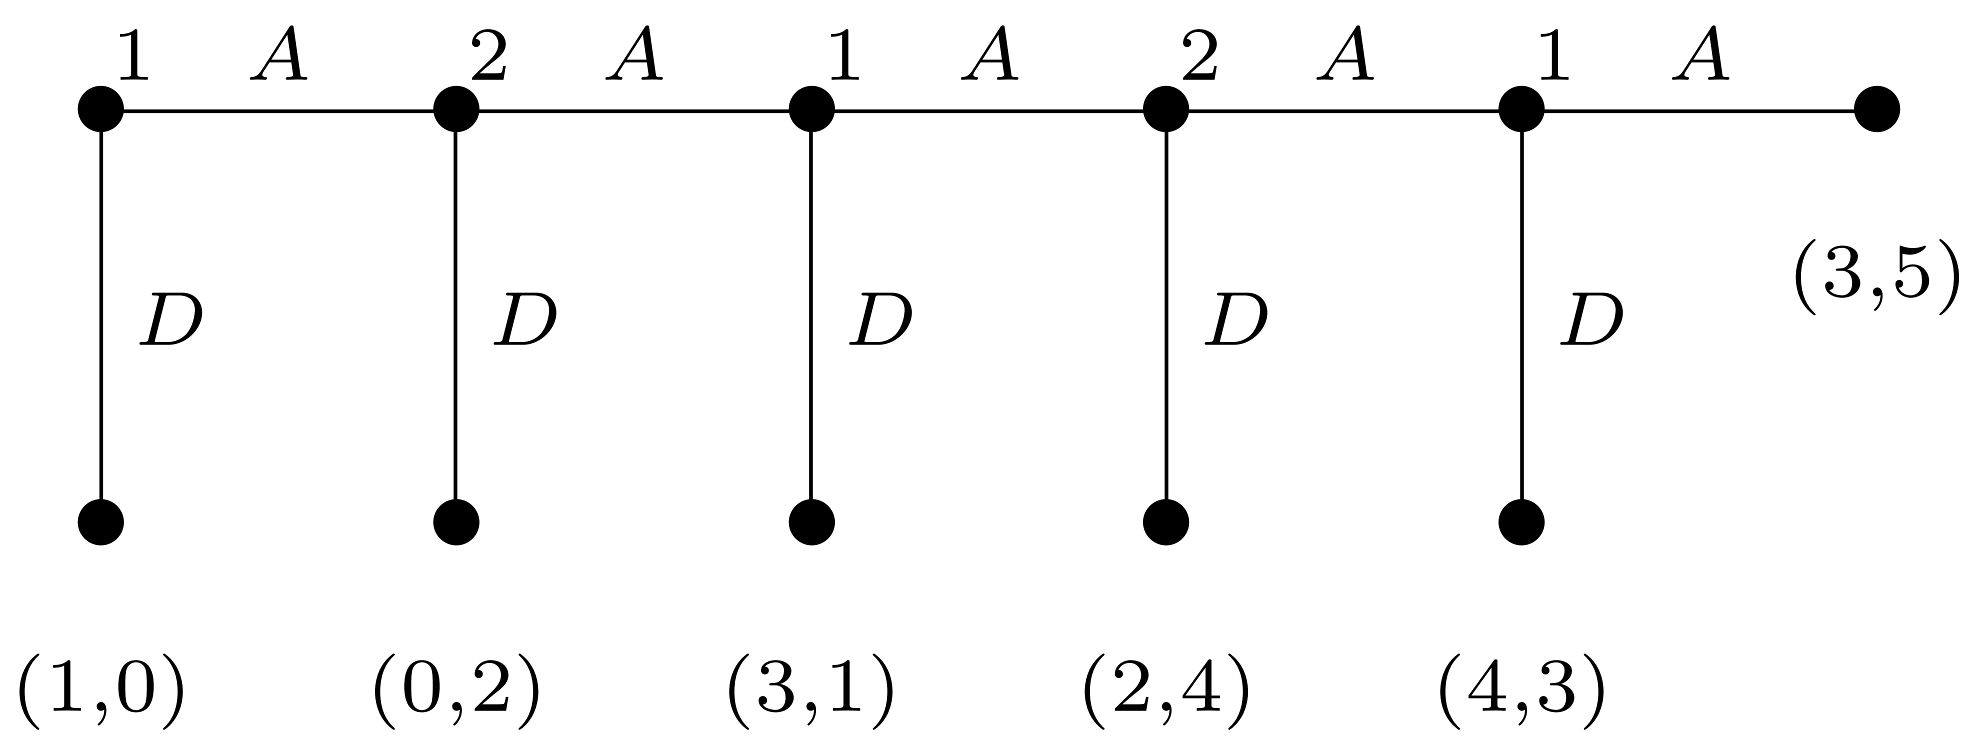

- The **unique** equilibrium is for each player to play **Down** **at the first opportunity.**
- **Empirically**, this is not how real people tend to play!
- **Theoretically**, what should you do if you arrive at an **off-path** node?
  - How do you update your beliefs to account for this probability 0 event?
  - If player 1 knows that you will update your beliefs in a way that causes you not to play **Down**, then playing **Down** is no longer their only rational choice...

Imperfect Information, informally

- **Perfect information** games model **sequential** actions that are **observed by all players**
  - **Randomness** can be modelled by a special **Nature** player with constant utility and known mixed strategy
- But many games involve **hidden** actions
  - Cribbage, poker, Scrabble
  - Sometimes actions of the **players** are hidden, sometimes **Nature's** actions are hidden, sometimes both
- **Imperfect information extensive form games** are a model of games with sequential actions, some of which may be **hidden**

Imperfect Information Extensive Form Game

Definition:
An **imperfect information game in extensive form** is a tuple $G = (N, A, H, Z, \chi, \rho, \sigma, u, I)$, where

- $(N, A, H, Z, \chi, \rho, \sigma, u)$ is a perfect information extensive form game, and
- $I = (I_1, \ldots, I_n)$, where $I_i = (I_{i,1}, \ldots, I_{i,k_i})$ is an **equivalence relation** on (i.e., partition of) $\{h \in H : \rho(h) = i\}$with the property that $\chi(h) = \chi(h')$ and $\rho(h) = \rho(h')$ whenever there exists a $j \in N$ for which $h \in I_{i,j}$ and $h' \in I_{i,j}$.

Imperfect Information Extensive Form Example

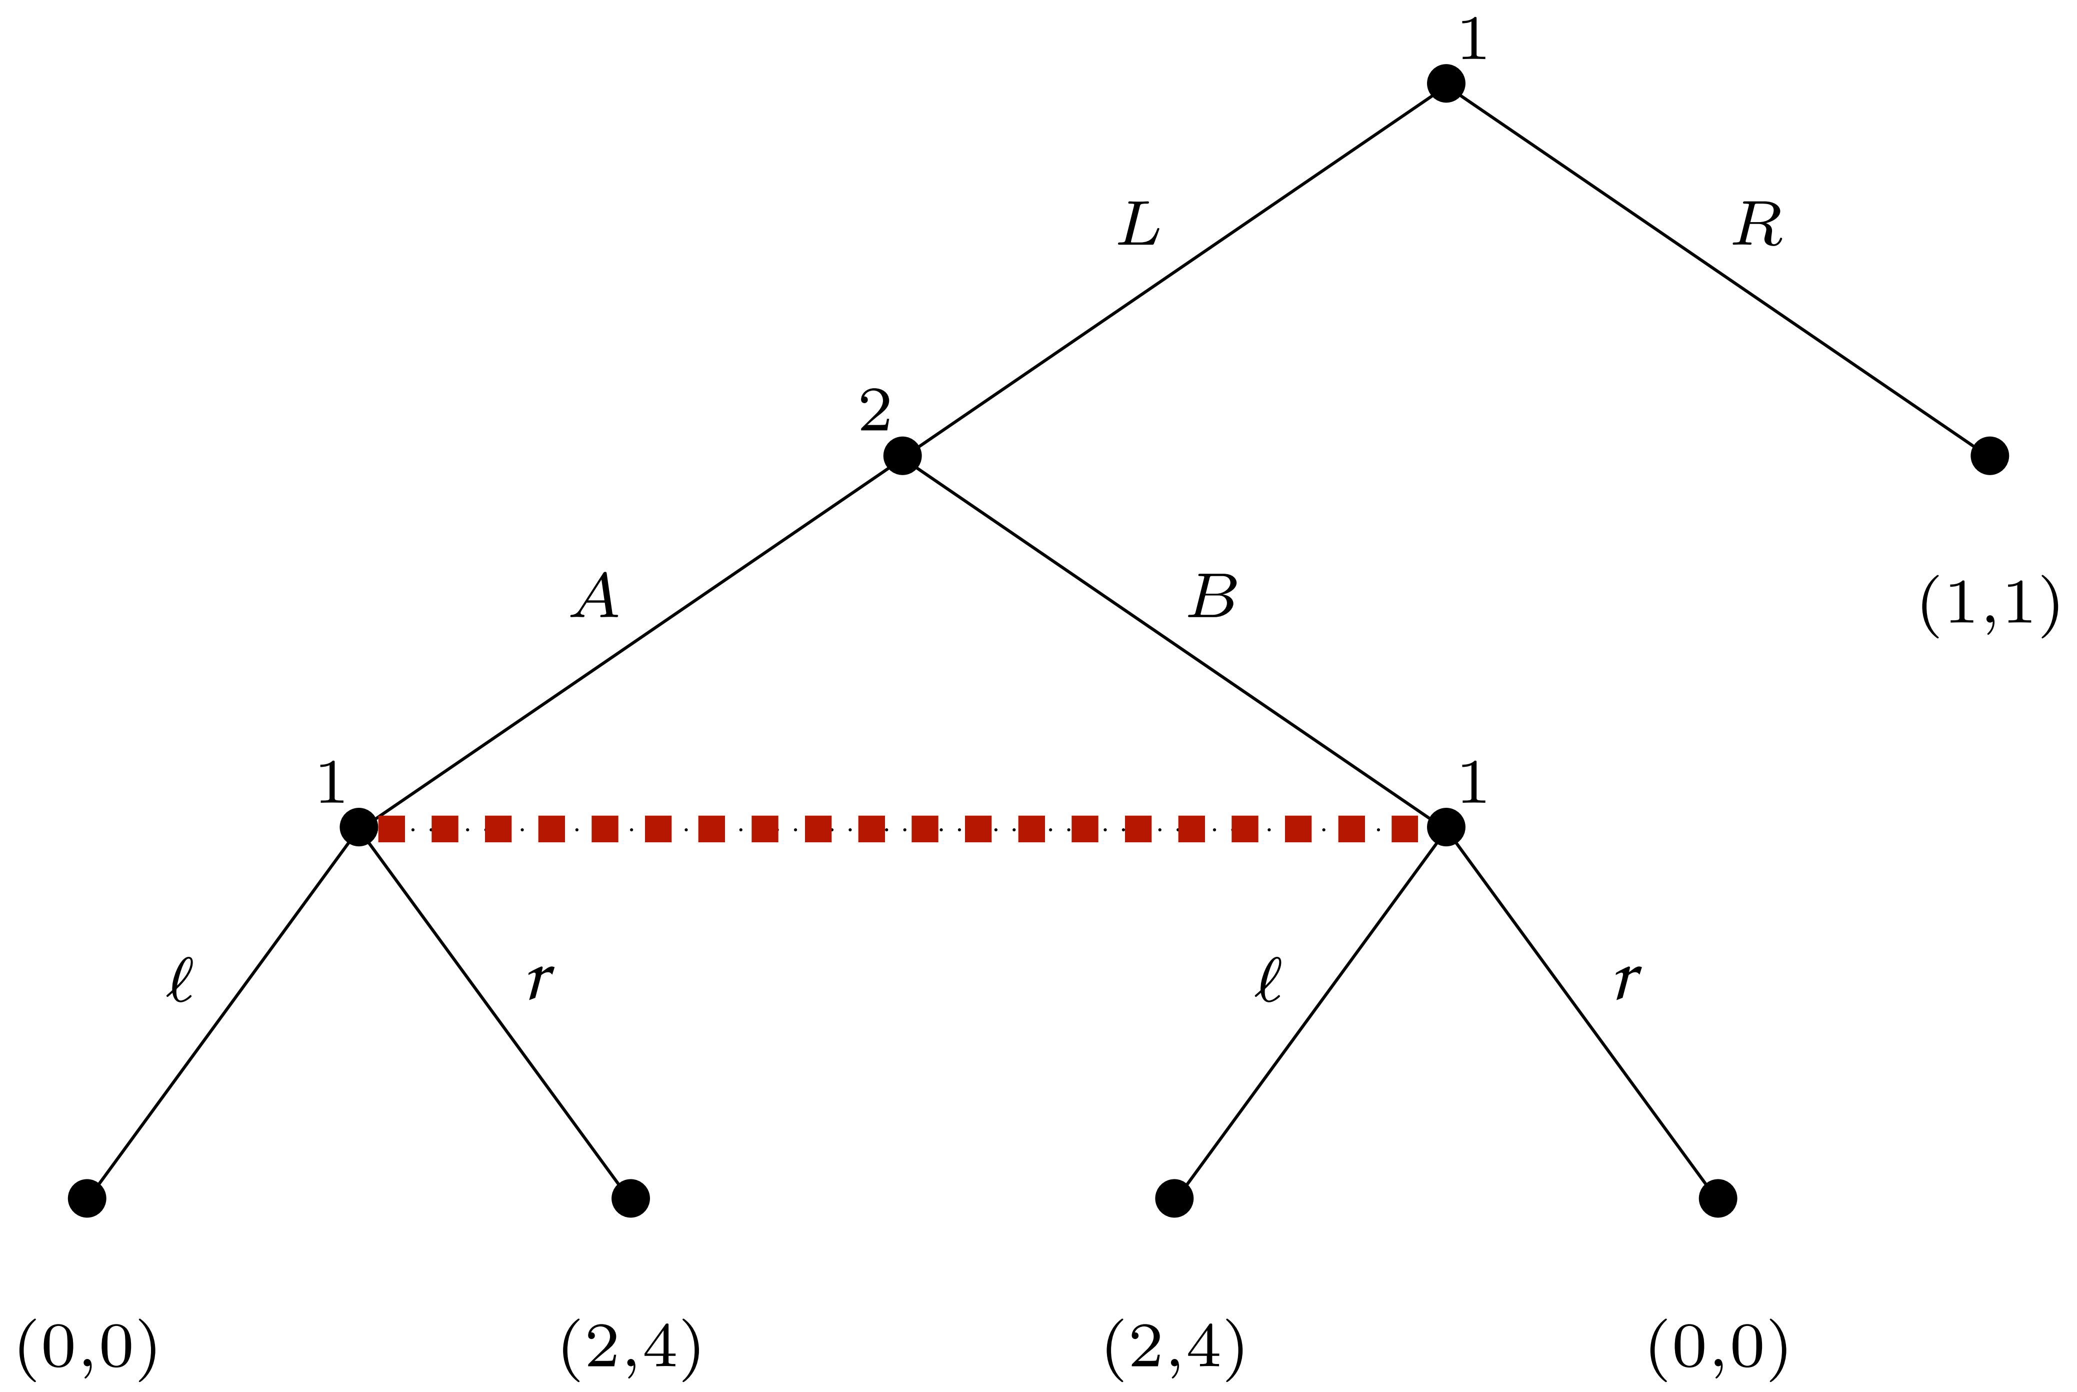

- The members of the equivalence classes are also called **information sets**
- Players **cannot distinguish** which **history** they are in within an information set
- **Question:** What are the information sets for each player in this game?

Pure Strategies

Question: What are the **pure strategies** in an **imperfect information** extensive-form game?

Definition:
Let $G = (N, A, H, Z, \chi, \rho, \sigma, u, I)$ be an imperfect information game in extensive form. Then the **pure strategies of player i** consist of the cross product of actions available to player i at each of their **information sets**, i.e.,

$$\prod_{I_{i,j} \in I_i} \chi(h)$$

- A pure strategy associates an action with **each** information set, even those that will **never be reached**

Questions:

In an imperfect information game:

1. What are the **mixed strategies**?
2. What is a **best response**?
3. What is a **Nash equilibrium**?

Induced Normal Form

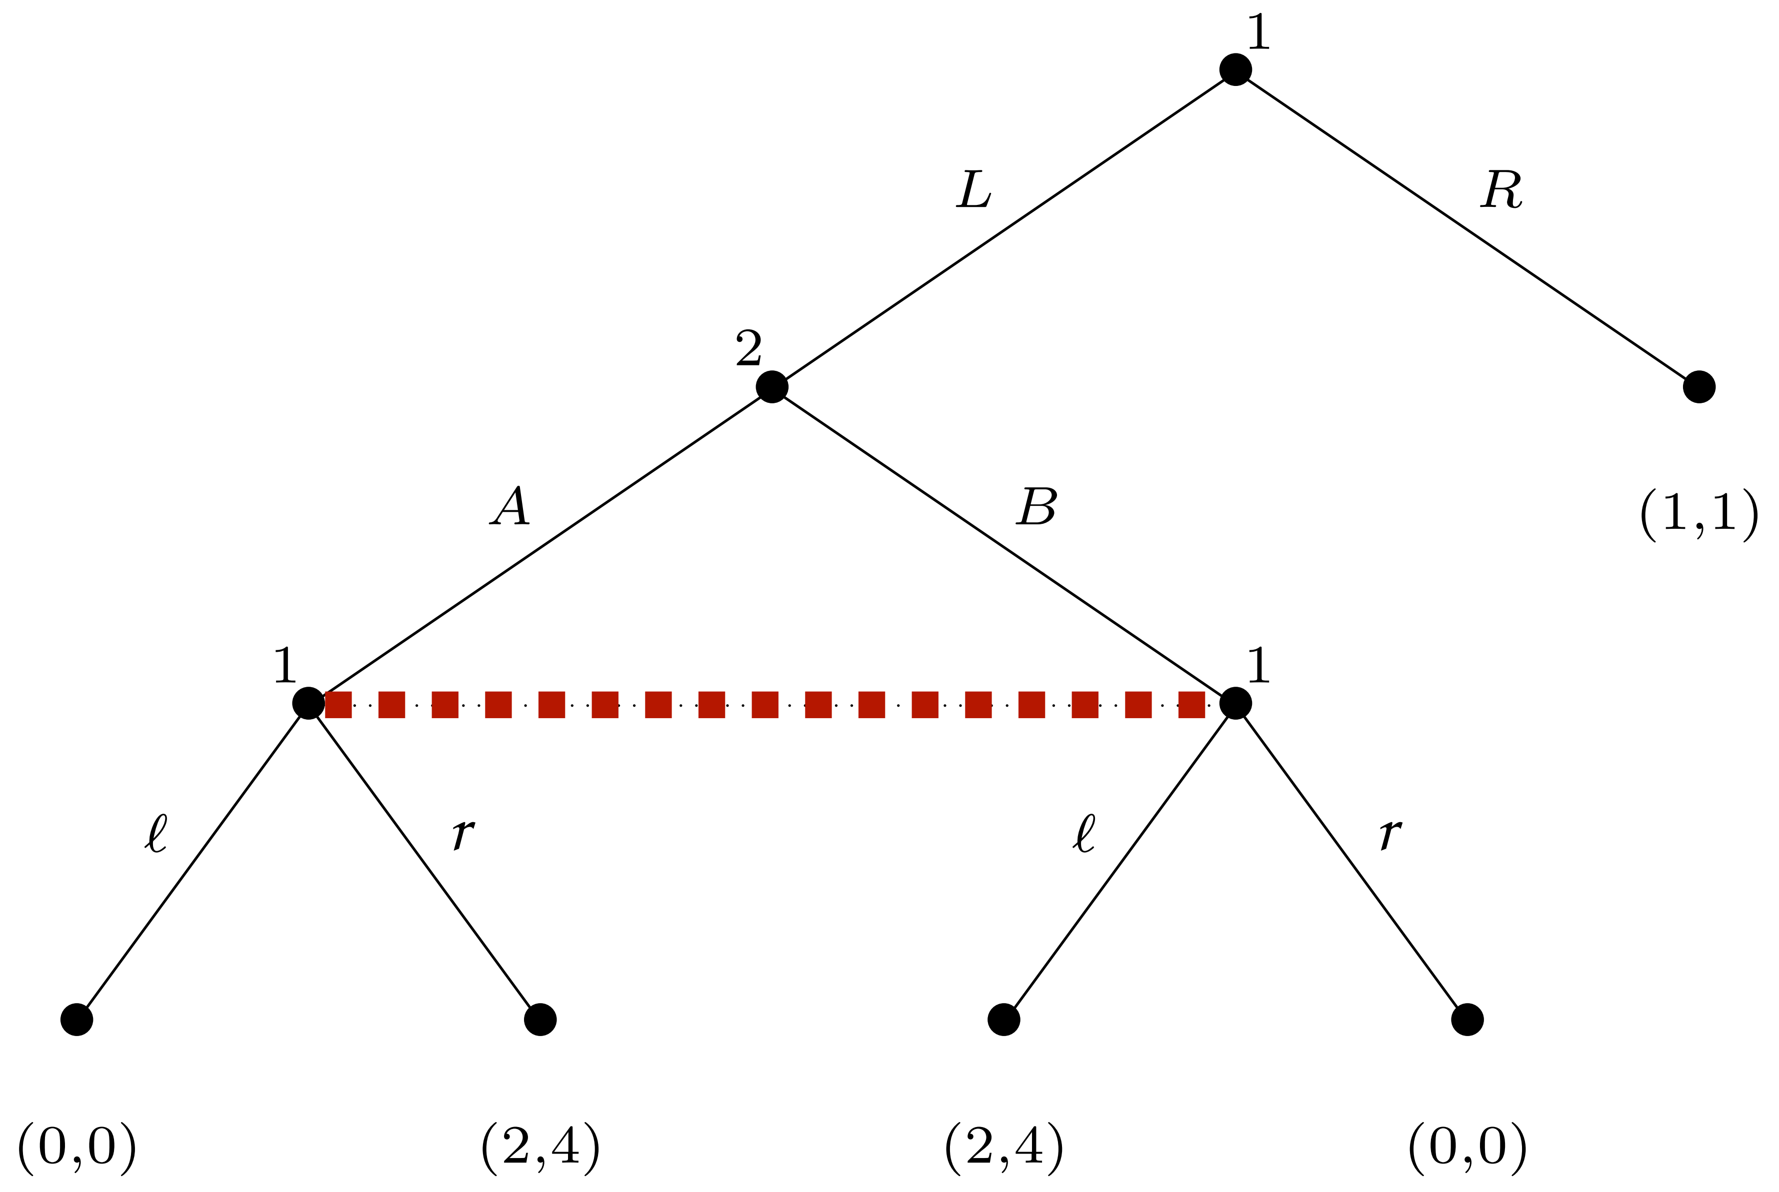

	A	B
L,ℓ	0,0	2,4
L,r	2,4	0,0
R,ℓ	1,1	1,1
R,r	1,1	1,1

Question:
Can you represent an arbitrary **perfect information** extensive form game as an **imperfect information** game?

- Any pair of pure strategies uniquely identifies a **terminal node**, which identifies a **utility** for each agent
- We have now defined a set of **agents**, **pure strategies**, and **utility functions**
- Any extensive form game defines a corresponding **induced normal form game**

Summary

- **Extensive form games** model **sequential** actions
- **Pure strategies** for extensive form games map **choice nodes** to **actions**
  - **Induced normal form:** normal form game with these pure strategies
  - Notions of mixed strategy, best response, etc. **translate directly**
- **Perfect information:** Every agent **sees all actions** of the other players
  - **Backward induction** computes a **pure strategy Nash equilibrium** for any perfect information extensive form game
- **Imperfect information:** Some actions are **hidden**
  - Histories are partitioned into **information sets**; players **cannot distinguish** between histories in the same information set