Crash Course on Game Theory

Kate Larson

Cheriton School of Computer Science

University of Waterloo

Introduction

- In multiagent decision making, the agents need to consider how others will act
  - This influences their own action choices
- We will often take the “**self-interested**” agent perspective
  - Self-interested does not mean adversarial! (A self-interested agent may be cooperative)
  - Self-interest means
    - Agents have their own descriptions of states of the world
    - Agents take actions based on these descriptions

What is Game Theory

The study of games!

- Bluffing in poker
- What move to make in chess
- How to play Rock-Paper-Scissors



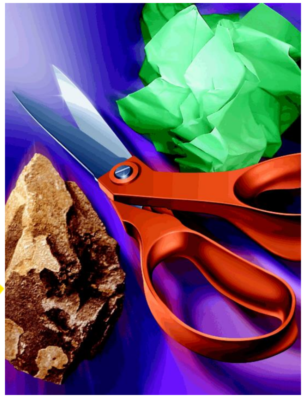

But also

- auction design
- strategic deterrence
- election laws
- coaching decisions
- routing protocols
- …

What is Game Theory

Game theory is a formal way to analyze **interactions** among a **group** of **rational** agents that behave **strategically**

Group: Must have more than 1 decision maker

- Otherwise, you have a decision problem, not a game

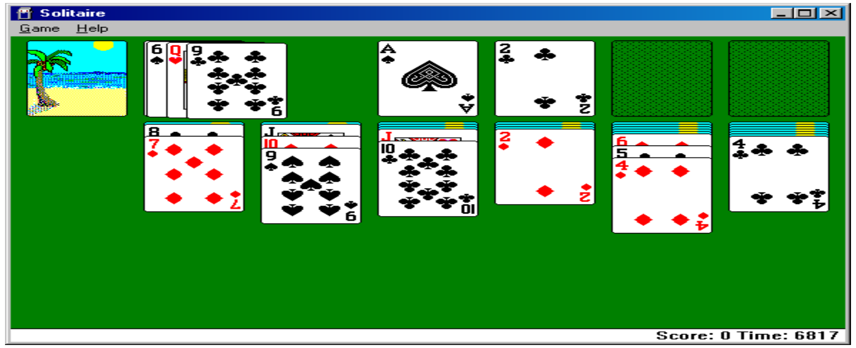

Solitaire is not a game!

What is Game Theory

Game theory is a formal way to analyze **interactions** among a **group** of **rational** agents that behave **strategically**

- **Interaction:** What one agent does directly affects at least one other
- **Strategic**: Agents take into account that their actions influence the game
- **Rational**: Agents chose their best actions

Example

Decision Problem

- Everyone pays their own bill

Game

- Before the meal, everyone decides to split the bill evenly

Strategic Form/Matrix Game/Normal Form

Set of agents: I={1,2,.,,,N}

Set of actions: $A_i=\{a_i^1,\ldots,a_i^m\}$

Outcome of a game is defined by a profile $a=(a_1,\ldots,a_n)$

Agents have preferences over outcomes

Utility functions u_i:A->**R**

Examples

		Agent 2: One	Agent 2: Two
Agent 1	One	2,-2	-3,3
	Two	-3,3	4,-4

Zero-sum game.

$\sum_{i=1}^{n} u_i(o)=0$

$I=\{1,2\}$

$A_i=\{One,Two\}$

A_n outcome is (One, Two)

$U_1((One,Two))=-3$ and $U_2((One,Two))=3$

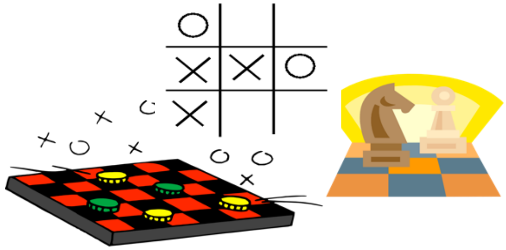

Examples

BoS

	B	S
B	2,1	0,0
S	0,0	1,2

Coordination Game

Chicken

	T	C
T	-1,-1	10,0
C	0,10	5,5

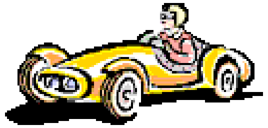

Anti-Coordination Game

Prisoners Dilemma

	Confess	Don't Confess
Confess	-5,-5	0,-10
Don't Confess	-10,0	-1,-1

Playing a Game

- Agents are rational
  - Let p_i be agent *i*'s belief about what its opponents will do
  - **Best response**: $a_i = \operatorname{argmax} \sum_{a_{-i}} u_i(a_i, a_{-i}) p_i(a_{-i})$

Notation Break: $a_{-i} = (a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n)$

Dominated Strategies

a'_i **strictly dominates** strategy a_i if

$$u_i(a'_i, a_{-i}) > u_i(a_i, a_{-i}) \forall a_{-i}$$

A rational agent will never play a dominated strategy!

Example

	Confess	Don't Confess
Confess	-5,-5	0,-10
Don't Confess	-10,0	-1,-1

Strict Dominance Does Not Capture the Whole Picture

	A	B	C
A	0,4	4,0	5,3
B	4,0	0,4	5,3
C	3,5	3,5	6,6

Nash Equilibrium

Key Insight: an agent's best-response depends on the actions of other agents

An action profile a* is a **Nash equilibrium** if no agent has incentive to change given that others do not change

$$\forall i u_i(a_i^*, a_{-i}^*) \geq u_i(a_i', a_{-i}^*) \forall a_i'$$

Nash Equilibrium

Equivalently, a* is a N.E. iff

$$\forall i a_i^* = \arg\max_{a_i} u_i(a_i, a_{-i}^*)$$

	A	B	C
A	0,4	4,0	5,3
B	4,0	0,4	5,3
C	3,5	3,5	6,6

(C,C) is a N.E. because

$$u_1(C,C) = \max \begin{bmatrix} u_1(A,C) \\ u_1(B,C) \\ u_1(C,C) \end{bmatrix}$$

AND

$$u_2(C,C) = \max \begin{bmatrix} u_2(C,A) \\ u_2(C,B) \\ u_2(C,C) \end{bmatrix}$$

Nash Equilibrium

- If (a_1*,a_2*) is a N.E. then player 1 won't want to change its action given player 2 is playing a_2*

- If (a_1*,a_2*) is a N.E. then player 2 won't want to change its action given player 1 is playing a_1*

-5,-5	0,-10
-10,0	-1,-1

	A	B	C
A	0,4	4,0	5,3
B	4,0	0,4	5,3
C	3,5	3,5	6,6

Another Example

	B	S
B	2,1	0,0
S	0,0	1,2

2 Nash Equilibria

Coordination Game

Yet Another Example

		Agent 2	
		One	Two
Agent 1	One	2,-2	-3,3
	Two	-3,3	4,-4

Mixed Strategies

- **(Mixed) Strategy**: s_i is a probability distribution over A_i
- **Strategy profile**: $s=(s_1,...,s_n)$
- **Expected utility**: $u_i(s)=\Sigma_a \Pi_j s(a_j) u_i(a)$

Example

Given strategy profile $s = ((\frac{1}{2}, \frac{1}{2}), (\frac{1}{10}, \frac{9}{10}))$ what is the expected utility of the agents?

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

(Mixed) Nash Equilibria

- **(Mixed) Strategy**: s_i is a probability distribution over A_i
- **Strategy profile**: $s=(s_1,...,s_n)$
- **Expected utility**: $u_i(s)=\Sigma_a \Pi_j s(a_j) u_i(a)$
- **Nash equilibrium**: s* is a (mixed) Nash equilibrium if

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i', s_{-i}^*) \forall s_i'$$

Yet Another Example

		q One	Two
p	One	2,-2	-3,3
	Two	-3,3	4,-4

How do we determine p and q?

Exercise

	B	S
B	2,1	0,0
S	0,0	1,2

This game has 3 Nash Equilibrium (2 pure strategy NE and 1 mixed strategy NE).

Mixed Nash Equilibrium

Theorem (Nash 1950): Every game in which the action sets are finite, has a mixed strategy equilibrium.

John Nash
Nobel Prize in Economics (1994)

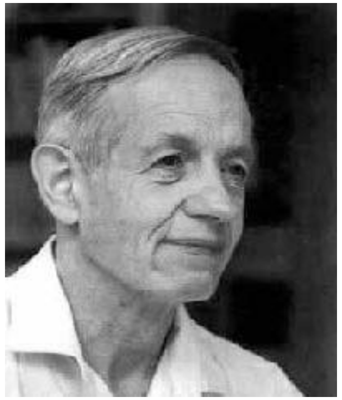

Finding NE

Existence proof is *non-constructive*

Finding equilibria?

- 2 player zero-sum games can be represented as a linear program (polynomial)
- For arbitrary games, the problem is in PPAD
- Finding equilibria with certain properties is often NP-hard

Repeated Games

Recall the Prisonner’s Dilemma. What if the prisoners are **habitual** criminals?

-5,-5	0,-10
-10,0	-1,-1

-5,-5	0,-10
-10,0	-1,-1

-5,-5	0,-10
-10,0	-1,-1

How do we define payoffs?

What is the strategy space?

Repeated Games

Recall the Prisonner's Dilemma. What if the prisoners are **habitual** criminals?

-5,-5	0,-10
-10,0	-1,-1

-5,-5	0,-10
-10,0	-1,-1

-5,-5	0,-10
-10,0	-1,-1

…

How do we define payoffs?

Average reward

Discounted Awards

…

Repeated Games

Recall the Prisonner's Dilemma. What if the prisoners are habitual criminals?

-5,-5	0,-10
-10,0	-1,-1

-5,-5	0,-10
-10,0	-1,-1

-5,-5	0,-10
-10,0	-1,-1

…

Strategy space becomes significantly larger!

S:H→A where H is the **history** of play so far

Can now reward and punish past behaviour, worry about reputation, establish trust,…

Repeated Games

Recall the Prisonner’s Dilemma. What if the prisoners are habitual criminals?

-5,-5	0,-10
-10,0	-1,-1

-5,-5	0,-10
-10,0	-1,-1

-5,-5	0,-10
-10,0	-1,-1

…

Grim Strategy: In first step cooperate. If opponent defects at some point, then defect forever

Tit-for-Tat: In first step cooperate. Copy whatever opponent did in previous stage.

Extensive Form Games

Perfect Information Game: $G = (N, A, H, Z, \alpha, \rho, \sigma, u)$

- N is the player set $|N| = n$
- $A = A_1 \times \ldots \times A_n$ is the action space
- H is the set of non-terminal choice nodes
- Z is the set of terminal nodes
- $\alpha : H \to 2^A$ action function, assigns to a choice node a set of possible actions
- $\rho : H \to N$ player function, assigns a player to each non-terminal node (player who gets to take an action)
- $\sigma : H \times A \to H \cup Z$, successor function that maps choice nodes and an action to a new choice node or terminal node where

  $$\forall h_1, h_2 \in H \text{ and } a_1, a_2 \in A \text{ if } h_1 \neq h_2 \text{ then } \sigma(h_1, a_1) \neq \sigma(h_2, a_2)$$

- $u = (u_1, \ldots, u_n)$ where $u_i : Z \to \mathbb{R}$ is utility function for player i over Z

Tree Representation

- The definition is really a tree
- Each node is defined by its history (sequence of nodes on the path between the root and it)
- Descendents of a node are all choice and terminal nodes in the subtree rooted at the node

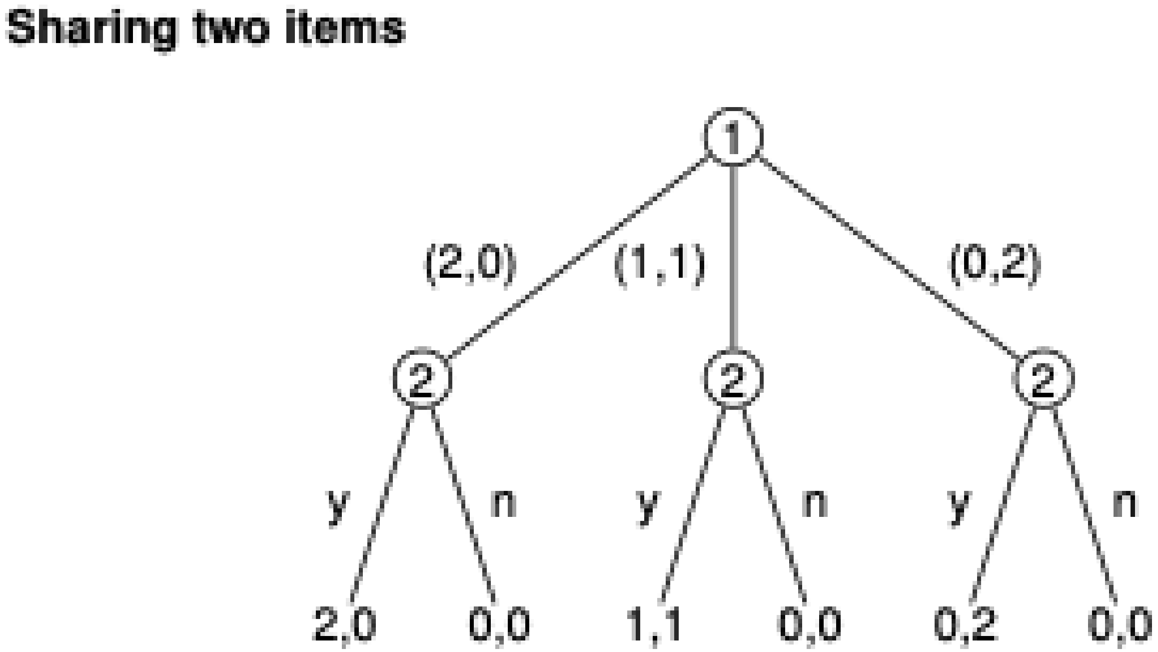

Strategies

- A strategy of a player is a function that assigns an action to each non-terminal history where the agent can take an action
- **Important**: The definition of a strategy requires a decision at each choice node, regardless of whether or not it is possible to reach that node given earlier moves.

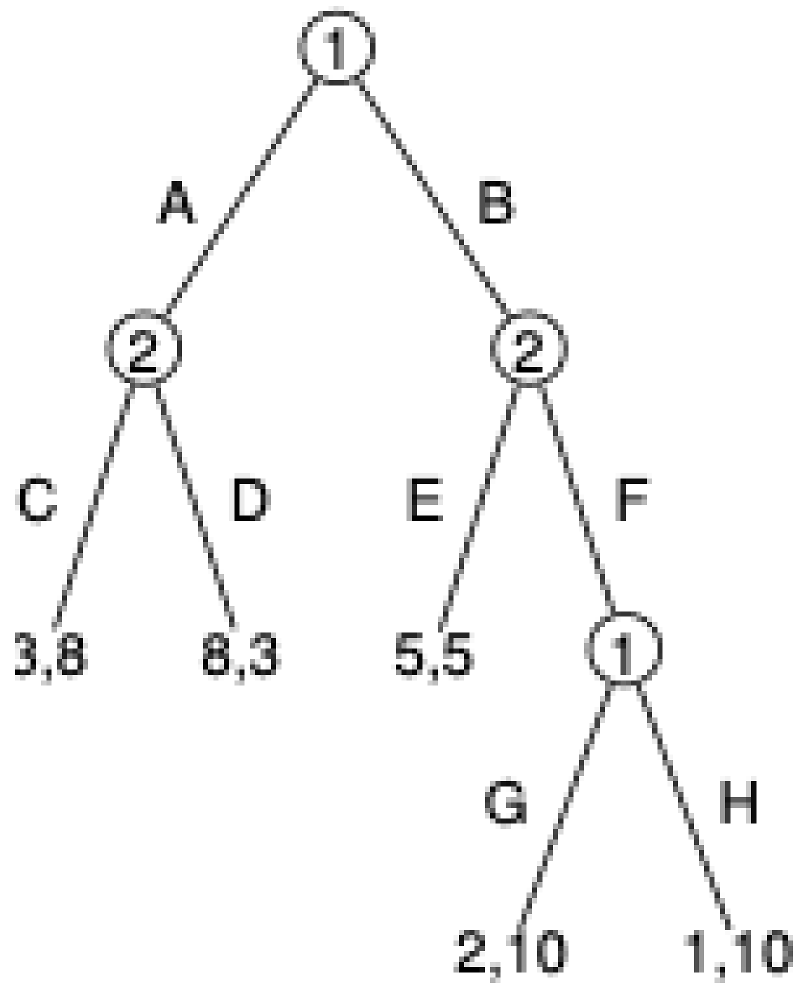

Example

We can transform an extensive form game into a normal form game.

	(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

Now we can just use the standard definition of Nash equilibrium, but....

Consider Subgames

Given a game G, the subgame of G rooted at node n is the restriction of G to n and its descendants.

Definition (Subgame perfect equilibrium)

A strategy profile s^* is a subgame perfect equilibrium if for all $i \in N$, and for all subgames of G, the restriction of s^* to G' (G' is a subgame of G) is a Nash equilibrium in G'. That is

$$\forall i, \forall G', u_i(s_i^*|_{G'}, s_{-i}^*|_{G'}) \geq u_i(s_i'|_{G'}, s_{-i}^*|_{G'}) \forall s_i'$$

Subgame Perfect Equilibria

Thm (Kuhn's Theorem): Every finite extensive form game with perfect information has a subgame perfect equilibrium (SPE).

You can compute SPE by backward induction.

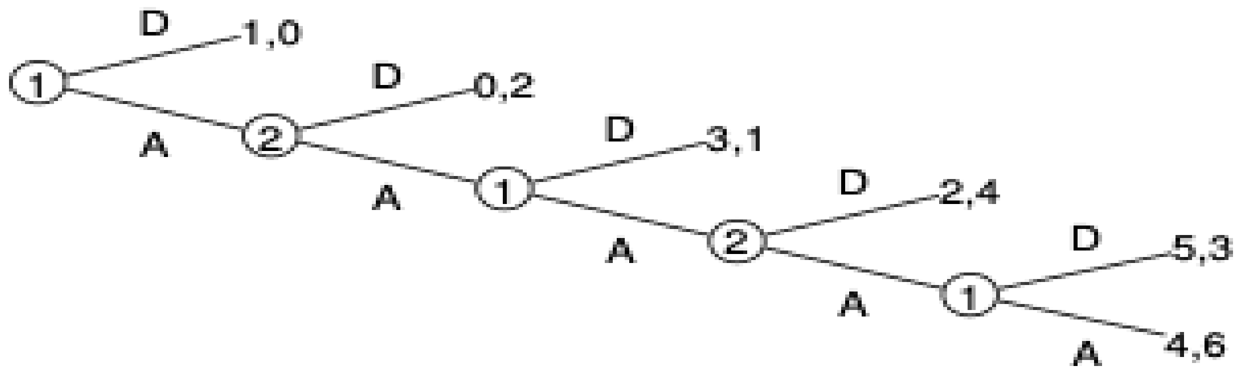

Imperfect Information Games

Sometimes agents have not observed everything, or have forgotten what they have observed

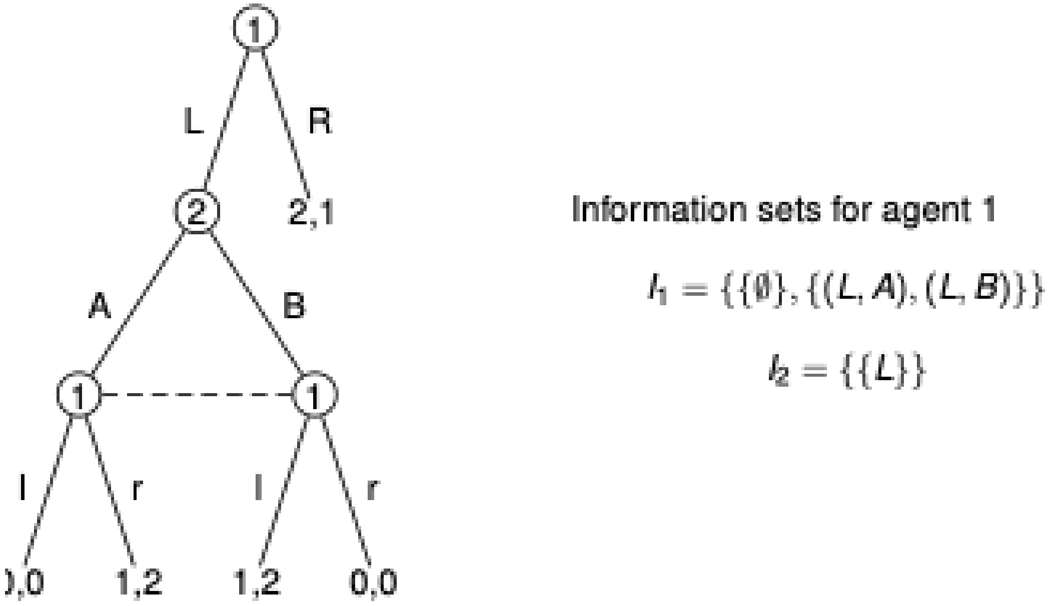

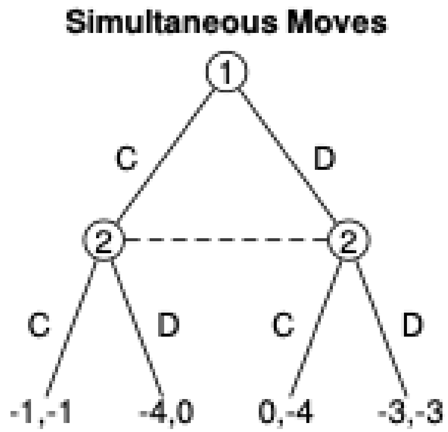

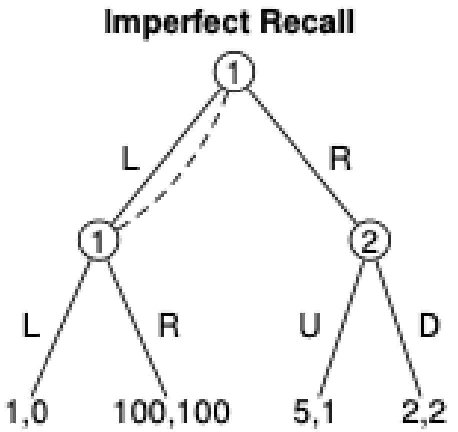

Bayesian Games

Sometime there are uncertainties about the actual game being played (incomplete information)

- Number of players
- Action sets
- Payoffs

	L	R
U	3,?	-2, ?
D	0, ?	6, ?

Bayesian games (games of incomplete information) are used to represent uncertainties about the game being played

Example: Extensive Form with Chance Moves

A special player, Nature, makes probabilistic moves.

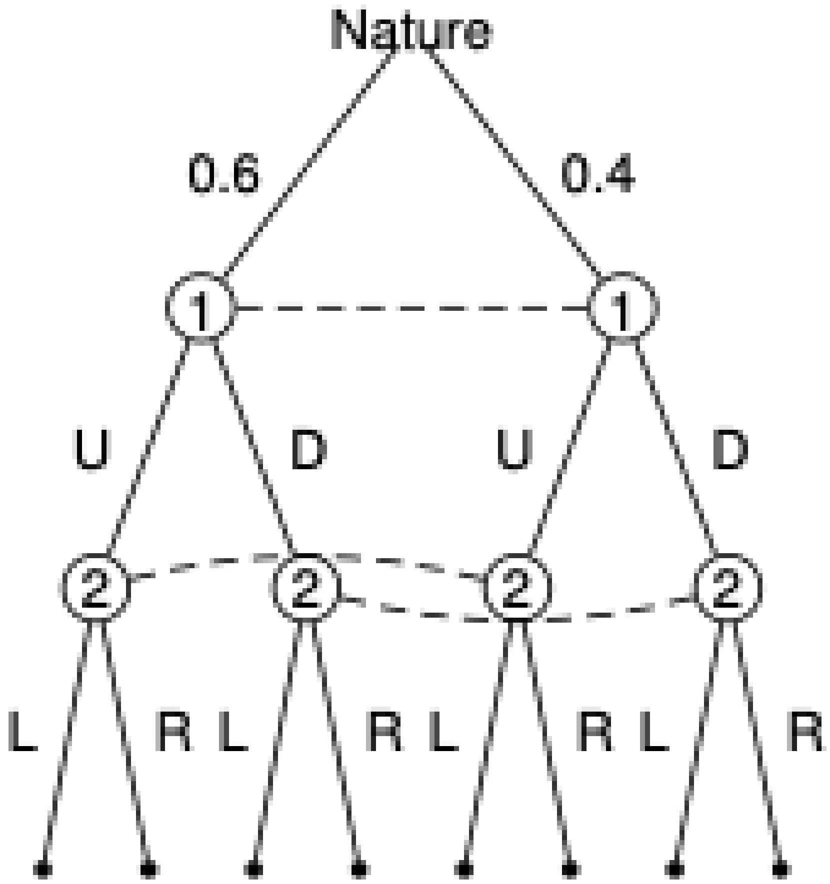

Example: Epistemic Types

BoS

- 2 agents
- $A_1 = A_2 = \{\text{soccer, hockey}\}$
- $\Theta = (\Theta_1, \Theta_2)$ where $\Theta_1 = \{H, S\}$, $\Theta_2 = \{H, S\}$
- Prior: $p_1(H) = 1$, $p_2(H) = \frac{2}{3}$, $p_2(S) = \frac{1}{3}$

Utilities can be captured by matrix-form

$\theta_2 = H$

	H	S
H	2,2	0,0
S	0,0	1,1

$\theta_2 = S$

	H	S
H	2,1	0,0
S	0,0	1,2

Questions