Introduction to Game Theory

Kate Larson

University of Waterloo

What is game theory?

The study of games!

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

•

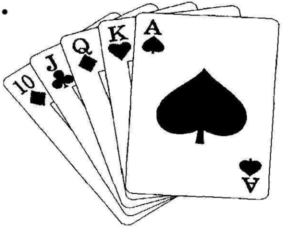

Also study of auction design, strategic deterrence, election laws, coaching decisions, routing protocols,...

What is game theory?

The study of games!

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

•

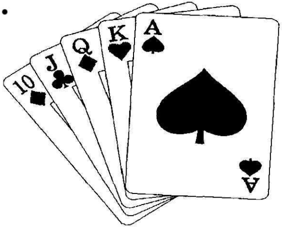

Also study of auction design, strategic deterrence, election laws, coaching decisions, routing protocols,...

What is game theory?

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

What is game theory?

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

Group: Must have more than one decision maker

- Otherwise you have a decision problem, not a game

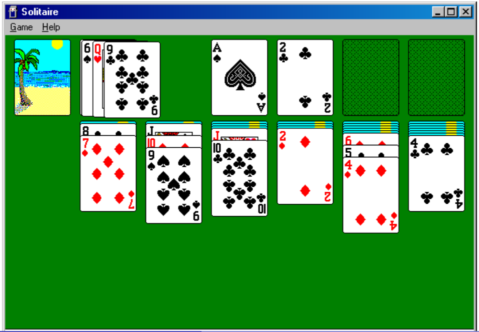

Solitaire is not a game.

What is game theory?

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

Interaction: What one agent does directly affects at least one other agent

Strategic: Agents take into account that their actions influence the game

Rational: An agent chooses its best action (maximizes its expected utility)

Example

Pretend that the entire class is going to go for lunch:

1. Everyone pays their own bill
2. Before ordering, everyone agrees to split the bill equally

Which situation is a game?

Normal Form

A normal form game is defined by

- Finite set of agents (or players) N, $|N| = n$
- Each agent i has an action space A_i
    - A_i is non-empty and finite
- An outcome is defined by an action profile $a = (a_1, \ldots, a_n)$ where a_i is the action taken by agent i
- Each agent has a utility function $u_i : A_1 \times \ldots \times A_n \mapsto \mathbb{R}$

Examples

Prisoners' Dilemma

	C	D
C	a,a	b,c
D	c,b	d,d

$c > a > d > b$

Pure coordination game

$\forall$ action profiles $a \in A_1 \times \ldots \times A_n$ and $\forall i, j$, $u_i(a) = u_j(a)$.

	L	R
L	1,1	0,0
R	0,0	1,1

Agents do not have conflicting interests.

Zero-sum games

$\forall a \in A_1 \times A_2, u_1(a) + u_2(a) = 0.$

Matching Pennies

	H	T
H	1,-1	-1, 1
T	-1,1	1,-1

	H	T
H	1	-1
T	-1	1

Given the utility of one agent, the other's utility is known.

More Examples

Most games have elements of both cooperation and competition.

BoS

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

Hawk-Dove

	D	H
D	3,3	1,4
H	4,1	0,0

Strategies

Notation: Given set X, let ΔX be the set of all probability distributions over X.

Definition

Given a normal form game, the set of mixed strategies for agent i is

$$S_i = \Delta A_i$$

The set of mixed strategy profiles is $S = S_1 \times \ldots \times S_n$.

Definition

A strategy s_i is a probability distribution over A_i. $s_i(a_i)$ is the probability action a_i will be played by mixed strategy s_i.

Strategies

Definition

The support of a mixed strategy s_i is

$$\{a_i | s_i(a_i) > 0\}$$

Definition

A pure strategy s_i is a strategy such that the support has size 1, i.e.

$$|\{a_i | s_i(a_i) > 0\}| = 1$$

A pure strategy plays a single action with probability 1.

Expected Utility

The expected utility of agent i given strategy profile s is

$$u_i(s) = \sum_{a \in A} u_i(a) \Pi_{j=1}^{n} s_j(a_j)$$

Example

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

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

Best-response

Given a game, what strategy should an agent choose?
We first consider only pure strategies.

Definition

Given a_{-i}, the best-response for agent i is $a_i^* \in A_i$ such that

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

Note that the best response may not be unique.
A best-response set is

$$B_i(a_{-i}) = \{a_i \in A_i | u_i(a_i, a_{-i}) \geq u_i(a_i', a_{-i}) \forall a_i' \in A_i\}$$

Nash Equilibrium

Definition

A profile a^* is a Nash equilibrium if $\forall i$, a_i^* is a best response to a_{-i}^*. That is

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

Equivalently, a^* is a Nash equilibrium if $\forall i$

$$a_i^* \in B(a_{-i}^*)$$

Examples

PD

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

BoS

	H	T
H	2,1	0,0
T	0,0	1,2

Matching Pennies

	H	T
H	1,-1	-1,1
T	-1,1	1,-1

Nash Equilibria

We need to extend the definition of a Nash equilibrium. Strategy profile s^* is a Nash equilibrium is for all i

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

Similarly, a best-response set is

$$B(s_{-i}) = \{s_i \in S_i | u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \forall s_i' \in S_i\}$$

Examples

Characterization of Mixed Nash Equilibria

s^* is a (mixed) Nash equilibrium if and only if

- the expected payoff, given s^*_{-i}, to every action to which s^*_i assigns positive probability is the same, and
- the expected payoff, given s^*_{-i} to every action to which s^*_i assigns zero probability is at most the expected payoff to any action to which s^*_i assigns positive probability.

Existence

Theorem (Nash, 1950)

Every finite normal form game has a Nash equilibrium.

Proof: Beyond scope of course.
Basic idea: Define set X to be all mixed strategy profiles. Show that it has nice properties (compact and convex).
Define $f : X \mapsto 2^X$ to be the best-response set function, i.e. given s, $f(s)$ is the set all strategy profiles $s' = (s'_1, \ldots, s'_n)$ such that s'_i is i's best response to s'_{-i}.
Show that f satisfies required properties of a fixed point theorem (Kakutani's or Brouwer's).
Then, f has a fixed point, i.e. there exists s such that $f(s) = s$. This s is mutual best-response – NE!

Existence

Theorem (Nash, 1950)

Every finite normal form game has a Nash equilibrium.

Proof: Beyond scope of course.
Basic idea: Define set X to be all mixed strategy profiles. Show that it has nice properties (compact and convex).
Define $f : X \mapsto 2^X$ to be the best-response set function, i.e. given s, $f(s)$ is the set all strategy profiles $s' = (s'_1, \ldots, s'_n)$ such that s'_i is i's best response to s'_{-i}.
Show that f satisfies required properties of a fixed point theorem (Kakutani's or Brouwer's).
Then, f has a fixed point, i.e. there exists s such that $f(s) = s$. This s is mutual best-response – NE!

Interpretations of Nash Equilibria

- Consequence of rational inference
- Focal point
- Self-enforcing agreement
- Stable social convention
- ...

Dominant and Dominated Strategies

For the time being, let us restrict ourselves to pure strategies.

Definition

Strategy s_i is a strictly dominant strategy if for all $s_i' \neq s_i$ and for all s_{-i}

$$u_i(s_i, s_{-i}) > u_i(s_i', s_{-i})$$

Prisoner's Dilemma

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

Dominant-strategy equilibria

Dominated Strategies

Definition

A strategy s_i is strictly dominated if there exists another strategy s'_i such that for all s_{-i}

$$u_i(s'_i, s_{-i}) > u_i(s_i, s_{-i})$$

Definition

A strategy s_i is weakly dominated if there exists another strategy s'_i such that for all s_{-i}

$$u_i(s'_i, s_{-i}) \geq u_i(s_i, s_{-i})$$

with strict inequality for some s_{-i}.

Example

	L	R
U	1,-1	-1,1
M	-1,1	1,-1
D	-2,5	-3,2

D is strictly dominated

	L	R
U	5,1	4,0
M	6,0	3,1
D	6,4	4,4

U and M are weakly dominated

Iterated Deletion of Strictly Dominated Strategies

Algorithm

- Let R_i be the removed set of strategies for agent i
- $R_i = \emptyset$
- Loop
  - Choose i and s_i such that $s_i \in A_i \setminus R_i$ and there exists s_i' such that
$$u_i(s_i', s_{-i}) > u_i(s_i, s_{-i}) \; \forall s_{-i}$$
  - Add s_i to R_i
  - Continue

Example

	R	C	L
U	3,-3	7,-7	15, -15
D	9,-9	8,-8	10,-10

Some Results

Theorem

If a unique strategy profile s^* survives iterated deletion then it is a Nash equilibrium.

Theorem

If s^* is a Nash equilibrium then it survives iterated elimination.

Weakly dominated strategies cause some problems.

Domination and Mixed Strategies

The definitions of domination (both strict and weak) can be easily extended to mixed strategies in the obvious way.

Theorem

Agent i's pure strategy s_i is strictly dominated if and only if there exists another (mixed) strategy σ_i such that

$$u_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i})$$

for all s_{-i}.

Example

	L	R
U	10,1	0,4
M	4,2	4,3
D	0,5	10,2

Strategy $(\frac{1}{2}, 0, \frac{1}{2})$ strictly dominates pure strategy M.

Theorem

If pure strategy s_i is strictly dominated, then so is any (mixed) strategy that plays s_i with positive probability.

Maxmin and Minmax Strategies

- A **maxmin strategy** of player *i* is one that maximizes its worst case payoff in the situation where the other agent is playing to cause it the greatest harm

$$\arg\max_{s_i}\min_{s_{-i}} u_i(s_i, s_{-i})$$

- A **minmax strategy** is the one that minimizes the maximum payoff the other player can get

$$\arg\min_{s_i}\max s_{-i} u_{-i}(s_i, s_{-i})$$

Example

In 2-player games, maxmin value of one player is equal to the minmax value of the other player.

	L	R
U	2,3	5,4
D	0,1	1,2

Zero-Sum Games

- The maxmin value of one player is equal to the minmax value of the other player
- For both players, the set of maxmin strategies coincides with the set of minmax strategies
- Any maxmin outcome is a Nash equilibrium. These are the only Nash equilibrium.

Solving Zero-Sum Games

Let U_i^* be unique expected utility for player i in equilibrium. Recall that $U_1^* = -U_2^*$.

$$\begin{array}{lll} \text{minimize} & U_1^* & \\ \text{subject to} & \sum_{a_k \in A_2} u_1(a_j, a_k) s_2(a_k) \leq U_1^* & \forall a_j \in A_1 \\ & \sum_{a_k \in A_2} s_2(a_k) = 1 & \\ & s_2(a_k) \geq 0 & \forall a_k \in A_2 \end{array}$$

LP for 2's mixed strategy in equilibrium.

Solving Zero-Sum Games

Let U_i^* be unique expected utility for player i in equilibrium. Recall that $U_1^* = -U_2^*$.

$$\begin{array}{lll} \text{maximize} & U_1^* & \\ \text{subject to} & \sum_{a_j \in A_1} u_1(a_j, a_k) s_1(a_j) \geq U_1^* & \forall a_k \in A_2 \\ & \sum_{a_j \in A_1} s_1(a_j) = 1 & \\ & s_1(a_j) \geq 0 & \forall a_j \in A_1 \end{array}$$

LP for 1's mixed strategy in equilibrium.

Two-Player General-Sum Games

LP formulation does not work for general-sum games since agents' interests are no longer diametrically opposed.

Linear Complementarity Problem (LCP)
Find any solution that satisfies

$$
\begin{array}{ll}
\sum_{a_k \in A_2} u_1(a_j, a_k) s_2(a_k) + r_1(a_j) = U_1^* & \forall a_j \in A_1 \\
\sum_{a_j \in A_1} u_2(a_j, a_k) s_1(a_j) + r_2(a_k) = U_2^* & \forall a_k \in A_2 \\
\sum_{a_j \in A_1} s_1(a_j) = 1 \;\; \sum_{a_k \in A_2} s_2(a_k) = 1 & \\
s_1(a_j) \geq 0, s_2(a_k) \geq 0 & \forall a_j \in A_1, a_k \in A_2 \\
r_1(a_j) \geq 0, r_2(a_k) \geq 0 & \forall a_j \in A_1, a_k \in A_2 \\
r_1(a_j) s_1(a_j) = 0, r_2(a_k) s_2(a_k) = 0 & \forall a_j \in A_1, a_k \in A_2
\end{array}
$$

For $n \geq 3$-player games, formulate a non-linear complementarity problem.

Complexity of Finding a NE

- Characterization is tricky since we do not have a decision problem (i.e. every game has at least one Nash Equilibrium)
- NE is in PPAD: Polynomial parity argument, directed version
  - Given an exponential-size directed graph, with every node having in-degree and out-degree at most one described by a polynomial-time computable function $f(v)$ that outputs the predecessor and successor of v, and a vertex s with a successor but no predecessors, find a $t \neq s$ that either has no successors or predecessors.

Complexity of Finding a NE

- Characterization is tricky since we do not have a decision problem (i.e. every game has at least one Nash Equilibrium)
- NE is in PPAD: Polynomial parity argument, directed version
  - Given an exponential-size directed graph, with every node having in-degree and out-degree at most one described by a polynomial-time computable function $f(v)$ that outputs the predecessor and successor of v, and a vertex s with a successor but no predecessors, find a $t \neq s$ that either has no successors or predecessors.

Extensive Form Games

aka Dynamic Games, aka Tree-Form Games

- Extensive form games allows us to model situations where agents take actions over time
- Simplest type is the perfect information game

Perfect Information Game

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
- $\alpha : H \rightarrow 2^A$ action function, assigns to a choice node a set of possible actions
- $\rho : H \rightarrow N$ player function, assigns a player to each non-terminal node (player who gets to take an action)
- $\sigma : H \times A \rightarrow H \cup Z$, 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 $u_i : Z \rightarrow \mathbb{R}$ is utility function for player i over Z

Tree Representation

- The definition is really a tree description
- Each node is defined by its history (sequence of nodes leading from root to it)
- The descendents of a node are all choice and terminal nodes in the subtree rooted at the node.

Example

Sharing two items

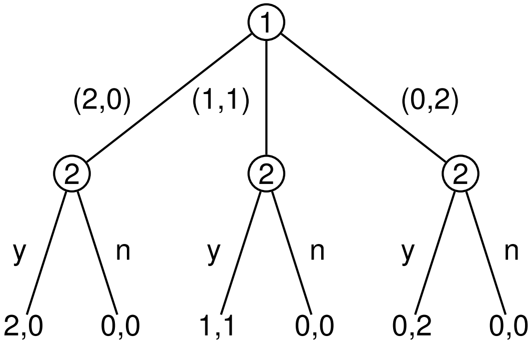

Strategies

- A strategy, s_i of player i is a function that assigns an action to each non-terminal history, at which the agent can move.
- Outcome: $o(s)$ of strategy profile s is the terminal history that results when agents play s
- **Important:** The strategy definition requires a decision at each choice node, regardless of whether or not it is possible to reach that node given earlier moves

Example

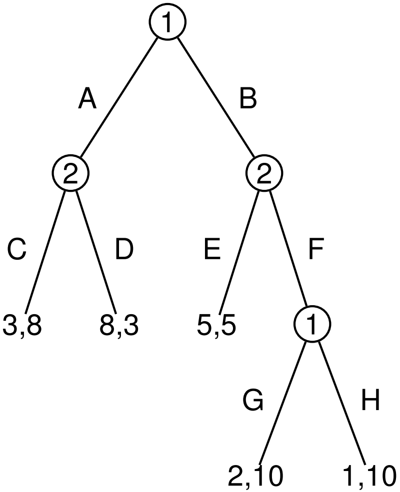

Strategy sets for the agents

$$S_1 = \{(A,G),(A,H),(B,G),(B,H)\}$$

$$S_2 = \{(C,E),(C,F),(D,E),(D,F)\}$$

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

Nash Equilibria

Definition (Nash Equilibrium)

Strategy profile s^* is a Nash Equilibrium in a perfect information, extensive form game if for all i

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

Theorem

Any perfect information game in extensive form has a pure strategy Nash equilibrium.

Intuition: Since players take turns, and everyone sees each move there is no reason to randomize.

Example: Bay of Pigs

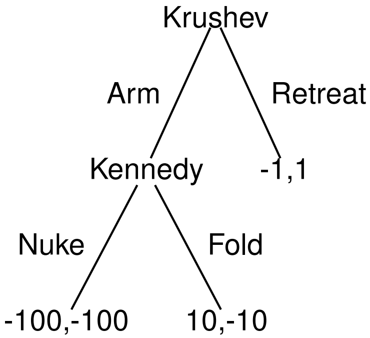

What are the NE?

Subgame Perfect Equilibrium

Nash Equilibrium can sometimes be too weak a solution concept.

Definition (Subgame)

Given a game G, the subgame of G rooted at node j is the restriction of G to its descendents of h.

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

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

Example: Bay of Pigs

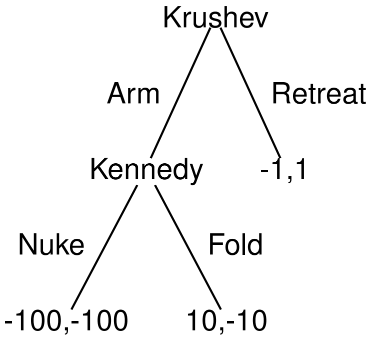

What are the SPE?

Existence of SPE

Theorem (Kuhn's Thm)

Every finite extensive form game with perfect information has a SPE.

You can find the SPE by backward induction.

- Identify equilibria in the bottom-most trees
- Work upwards

Centipede Game

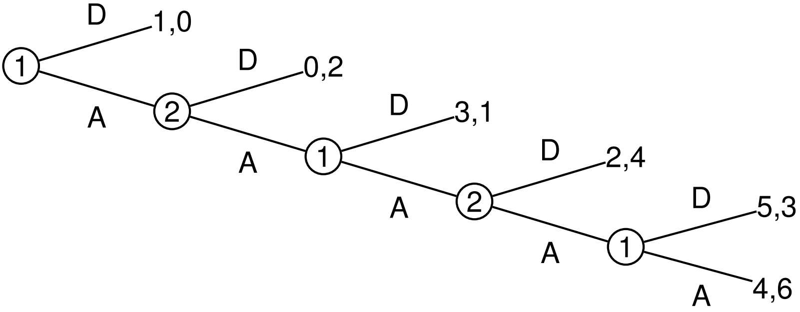

Imperfect Information Games

- Sometimes agents have not observed everything, or else can not remember what they have observed

Imperfect information games: Choice nodes H are partitioned into *information sets*.

- If two choice nodes are in the same information set, then the agent can not distinguish between them.
- Actions available to an agent must be the same for all nodes in the same information set

Example

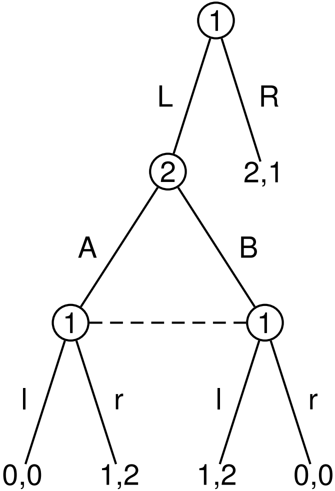

Information sets for agent 1

$$I_1 = \{\{\emptyset\}, \{(L, A), (L, B)\}\}$$

$$I_2 = \{\{L\}\}$$

More Examples

Simultaneous Moves

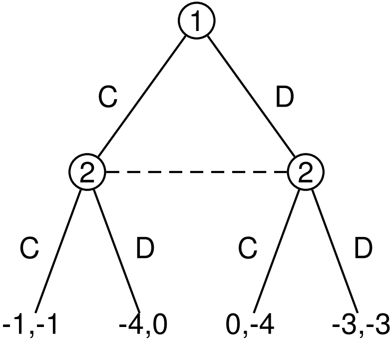

Imperfect Recall

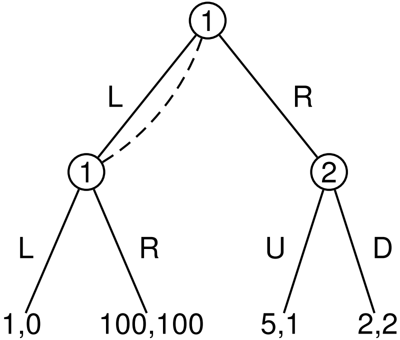

Strategies

- **Pure strategy:** a function that assigns an action in $A_i(I_i)$ to each information set $I_i \in \mathcal{I}_i$
- **Mixed strategy:** probability distribution over pure strategies
- **Behavorial strategy:** probability distribution over actions available to agent i at each of its information sets (independent distributions)

Behavorial Strategies

Definition

Given extensive game G, a behavorial strategy for player i specifies, for every $I_i \in \mathcal{I}_i$ *and action* $a_i \in A_i(I_i)$*, a probability* $\lambda_i(a_i, I_i) \geq 0$ *with*

$$\sum_{a_i \in A_i(I_i)} \lambda(a_i, I_i) = 1$$

Example

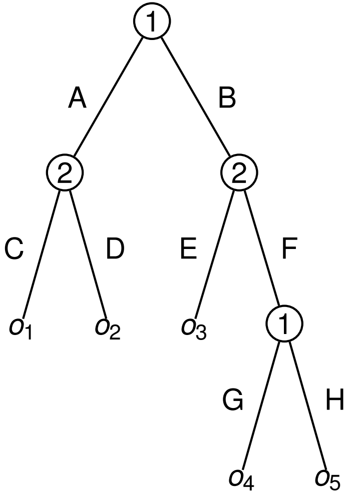

Mixed Strategy:
(0.4(A,G), 0.6(B,H))

Behavorial Strategy:

- Play A with probability 0.5
- Play G with probability 0.3

Mixed and Behavorial Strategies

In general you can not compare the two types of strategies.

But for games with perfect recall

- Any mixed strategy can be replaced with a behavorial strategy
- Any behavorial strategy can be replaced with a mixed strategy

Example

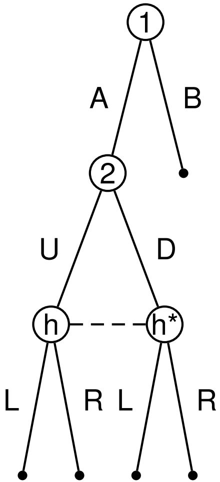

Mixed Strategy:
(<0.3(A,L)>,<0.2(A,R)>,
<0.5(B,L)>)

Behavorial Strategy:

- At I_1: (0.5, 0.5)
- At I_2: (0.6, 0.4)

Bayesian Games

So far we have assumed that all players know what game they are playing

- Number of players
- Actions available to each player
- Payoffs associated with strategy profiles

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

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

Bayesian Games

There are different possible representations.

Information Sets

- N set of agents
- G set of games
  - Same strategy sets for each game and agent
- $\Pi(G)$ is the set of all probability distributions over G
  - $P(G) \in \Pi(G)$ common prior
- $I = (I_1, \ldots, I_n)$ are information sets (partitions over games)

Extensive Form With Chance Moves

A special player, Nature, makes probabilistic moves.

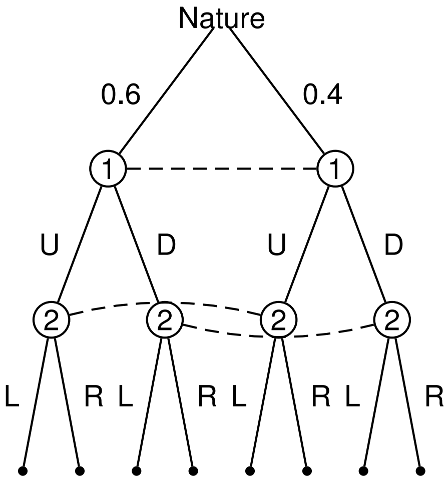

Epistemic Types

Epistemic types captures uncertainty directly over a game's utility functions.

- N set of agents
- $A = (A_1, \ldots, A_n)$ actions for each agent
- $\Theta = \Theta_1 \times \ldots \times \Theta_n$ where Θ_i is *type space* of each agent
- $p : \Theta \to [0, 1]$ is common prior over types
- Each agent has utility function $u_i : A \times \Theta \to \mathbb{R}$

Example

BoS

- 2 agents
- $A_1 = A_2 = \{\text{soccer, hockey}\}$
- $\Theta = (\Theta_1, \Theta_2)$ where $\Theta_1 = \{\text{H, S}\}$, $\Theta_2 = \{\text{H, S}\}$
- 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

Strategies and Utility

- A strategy $s_i(\theta_i)$ is a mapping from Θ_i to A_i. It specifies what action (or what distribution of actions) to take for each type.

Utility: $u_i(s|\theta_i)$

- *ex-ante* EU (know nothing about types)

$$EU = \sum_{\theta_i \in \Theta_i} p(\theta_i) EU_i(s_i|\theta_i)$$

- *interim* EU (know own type)

$$EU = EU_i(s|\theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} p(\theta_{-i}|\theta_i) \sum_{a \in A} \Pi_{j \in N} s_j(a_j, \theta_j)) u_i(a, \theta_{-i}, \theta_i)$$

- *ex-post* EU (know everyones type)

Example

- 2 firms, 1 and 2, competing to create some product.
- If one makes the product then it has to share with the other.
- Product development cost is $c \in (0, 1)$
- Benefit of having the product is known only to each firm
  - Type θ_i drawn uniformly from $[0, 1]$
  - Benefit of having product is θ_i^2

Bayes Nash Equilibrium

Definition (BNE)

Strategy profile s^* is a Bayes Nash equilibrium if $\forall i, \forall\theta_i$

$$EU(s_i^*, s_{-i}^*|\theta_i) \geq EU(s_i', s_{-i}^*|\theta_i) \forall s_i' \neq s_i^*$$

Example Continued

- Let $s_i(\theta_i) = 1$ if i develops product, and 0 otherwise.
- If i develops product

  $$u_i = \theta_i^2 - c$$

  If it does not then

  $$u_i = \theta_i^2 Pr(s_j(\theta_j) = 1)$$

- Thus, develop product if and only if

$$\theta_i^2 - c \geq \theta_i^2 Pr(s_j(\theta_j) = 1) \Rightarrow \theta_i \geq \sqrt{\frac{c}{1 - Pr(s_j(\theta_j) = 1)}}$$

Example Continued

Suppose $\hat{\theta}_1, \hat{\theta}_2 \in (0,1)$ are cutoff values in BNE.

- If so, then $Pr(s_j(\theta_j) = 1) = 1 - \hat{\theta}_j$
- We must have

$$\hat{\theta}_i \geq \sqrt{\frac{c}{\hat{\theta}_j}} \Rightarrow \hat{\theta}_i^2 \hat{\theta}_j = c$$

  and

$$\hat{\theta}_j^2 \hat{\theta}_i = c$$

- Therefore

$$\hat{\theta}_i^2 \hat{\theta}_j = \hat{\theta}_j^2 \hat{\theta}_i$$

  and so

$$\hat{\theta}_i = \hat{\theta}_j = \theta^* = c^{\frac{1}{3}}$$