GAME THEORY

Overview

- Game theory explicitly considers interactions between individuals.
- Thus it seems like a suitable framework for studying agent interactions.
- This lecture provides an introduction to some of the concepts of game theory.
- In particular, this lecture considers zero-sum games.

- Being able to figure out how to interact is important.

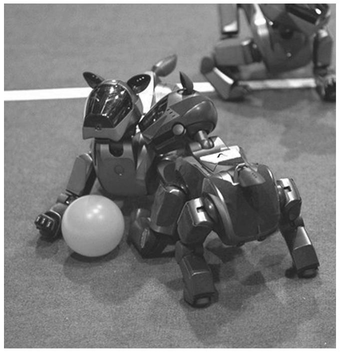

Basic notions

- Game theory is about *games of strategy*.
- When one agent makes a move, another agent repsonds not by chance but by figuring out what is best for it.
- To do this, that agent needs to have some way of knowing what is good for it.
- It also has to have some way of knowing what is good for its *opponent* (note the adversarial langauge) in order to try and second guess it.

- The basic notions of game theory include:
  - players (decision makers);
  - choices (feasible actions);
  - payoffs (benefits, prizes, rewards . . .); and
  - preferences over payoffs (objectives).
- Game theory is concerned with determining when one choice is better than another choice for a particular player.
- These “games” can be *static* or *dynamic*.
- In dynamic games the order of the moves/choices is important.
- Here we will only deal with static games.

- A simple game is this:
  - Player 1 chooses H or T
  - Player 2 chooses H or T (not knowing what Player 1 chooses).
  - If both choose the same Player 2 wins \$1 from Player 1.
  - If they are different,Player 1 wins \$1 from Player 2.
- We can draw this in *extensive form*.

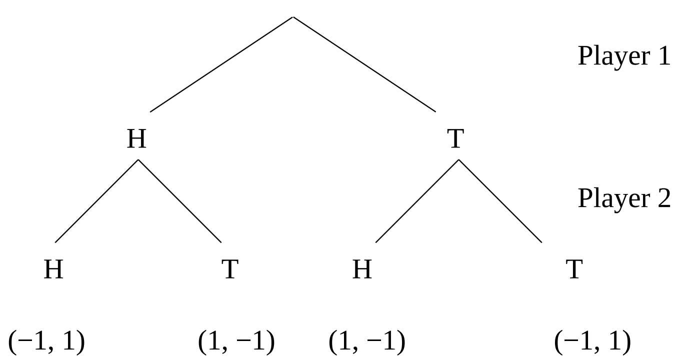

- A *strategy* for a player is a function which determines which choice he makes at every choice point.
- We distinguish games like the one above, in which Player 2 doesn’t know what Player 1 chose, from situations in which Player 2 has *perfect information*.
- The above game is one of perfect information if Player 1 reveals his choice before Player 2 chooses.
- The extensive form for this game is on the next slide

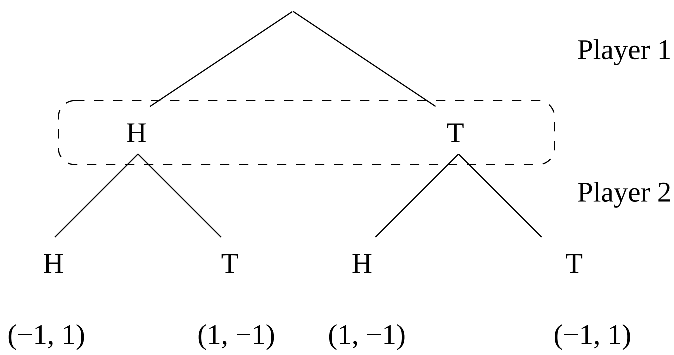

- We can also write games in *strategic form*.
- Here is the matching game:

		Player 2	
		H	T
Player 1	H	(-1, 1)	(1, -1)
	T	(1, -1)	(-1, 1)

- The rows are Player 1's moves, the columns are Player 2's moves.
- The first payoff in each row is that of Player 1, the second is that of Player 2.
- This game is *non-cooperative*
- A game is said to be *zero sum* if and only if the payoffs p_i at each terminal of the extensive form are such that:

$$\sum_i p_i = 0$$

Two Person Zero Sum Games

- We can write two person zero sum games in *normal form*
- An example:

$$A = \begin{bmatrix} -1 & -3 & -3 & -2 \\ 0 & 1 & -2 & -1 \\ 2 & -2 & 0 & 1 \end{bmatrix}$$

- As with strategic form the rows are the moves of P1 and the columns those of P2
- The entries a_{ij} represent the payoff vector $(a_{ij}, -a_{ij})$.
- How should the players behave?

- One thing that P1 might do is to ask "for each move I might make, what is the worst thing that P2 can do?".
- Thus he looks for:

$$\alpha_i = \min_j a_{ij}$$

- He then looks for the move which makes this as good as possible choosing i^* such that:

$$a_{i^*,j} = arg \max_i \min_j a_{ij}$$

- In this case $i^* = \{2, 3\}$.
- Similarly P2 could analyse looking for the move which will minimise his loss given that P1 will try to make this as big as possible choosing j^*:

$$a_{j^*} = arg \min_j \max_i a_{ij}$$

- In this case, $j^* = \{3\}$.

- Here both agents are trying to do their best to hurt the other since this is the same as profiting as much as they can.
- The value $\lfloor(A) = \min_j a_{i^*,j}$ is called the *gain floor* of the game.
- The value $\lceil(A) = \max_j a_{i,j^*}$ is called the *loss ceiling* of the game.
- Now consider:

$$A = \begin{bmatrix} -4 & 0 & 1 \\ 0 & 1 & -3 \\ -1 & -2 & -1 \end{bmatrix}$$

- P1 should take $i^* = 3$ and P2 should consider $j^* = 1$
- However, if P1 knows P2 will choose 1, then he should choose 2.
- But if P2 knows P1 will choose 2, then he should choose 3.
- and so on.

- What we have here is an unstable solution.
- A solution is *stable* if no player wants to unilaterally move away from the solution.
- A solution is inadmissible if there are solutions that produce better payoffs for all players than the given solution.
- What we want is a way of identifying stable solutions.
- It is easy to see that both players will settle on (i^*, j^*) if $\lfloor(A) = \lceil(A)$.
- In this case:

$$\min_j a_{i^*j} = \lfloor(A) = \lceil(A) = \max_i a_{ij^*}$$

- If $\lfloor(A) = \lceil(A)$ then:
  - A has a *saddle point*
  - The *value* for the game is $V = \lfloor(A) = \lceil(A)$
- This works fine for games which do have a saddle point, however, what happens if:

$$\lfloor(A) < \lceil(A)$$

  as in the game:

$$A = \begin{bmatrix} 3 & -1 \\ 0 & 1 \end{bmatrix}$$

- Here P1 has $\lfloor(A) = 0$ and $i^* = 2$.
- For P2, $\lceil(A) = 1$ and $j^* = 2$

Mixed Strategies

- What we want is a "spy-proof" strategy.
- This is one which works even if the other player knows what the strategy is.
- We manage this by moving from a *pure* strategy in which a player makes a definite choice of move. . .
- . . . to a *mixed* strategy in which a player makes a random choice across a set of pure strategies.

- More formally, P1 picks a vector of probabilities:

$$x = (x_1, x_2)$$

where

$$\sum_i x_i = 1$$

and

$$x_i \geq 0$$

- P1 then picks strategy i with probability x_i.
- To determine the strategy, P1 needs then to compute the best values of x_i and x_j.
- These will be the values which give P1 the highest expected payoff for his mixed strategy.

- P1's analysis would be something like this:

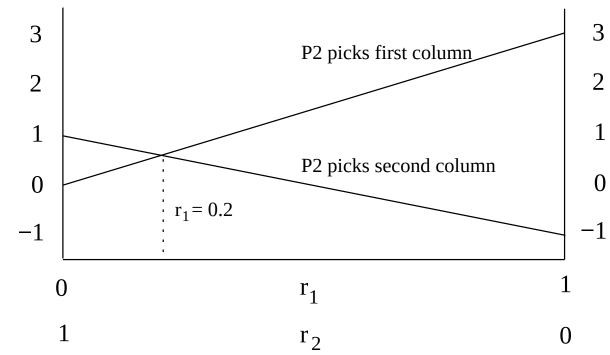

- P2 can analayse the problem in terms of a probability vector

$$y = (y_1, y_2)$$

and come up with a similar picture:

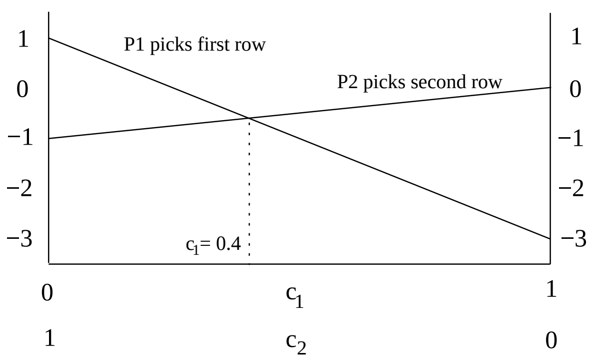

- Now, let's consider the payoff's the players will expect.
- With P1 having mixed strategy (x_1, x_2) and P2 having (y_1, y_2), the value of the game will be:

$$\begin{aligned} V &= 3x_1y_1 + 0(1 - x_1)y_1 \\ &\quad - x_1(1 - y_1) + (1 - x_1)(1 - y_1) \\ &= 5x_1y_1 - y_1 - 2x_1 + 1 \end{aligned}$$

- Now, let's assume that P1 uses $x_1^* = 0.2$ as calculated above. Then:

$$\begin{aligned} V &= 5(0.2y_1) - y_1 - 2(0.2) + 1 \\ &= 0.6 \end{aligned}$$

- Similarly, if P2 picks $y_1^* = 0.4$ then:

$$V = 0.6$$

- The neat thing is that the expected value for one player does not depend upon the strategy of the other player.
- This result generalises.
- Von Neumann's Minimax Theoreom shows that you can always find a pair of mixed strategies x^* and y^* which result in P1 and P2 have the same expected value for the game.

- This is important because it means we have something similar to:

$$\lfloor(A) = \lceil(A)$$

- In other words, there is a kind of stability.
- It is also possible to prove that either player can do no better using a pure strategy than he can using a mixed strategy.
- This makes it possible for one player to know that the other player is going to use a mixed strategy.
- This is the key to stability.

Summary

- This lecture has introduced some of the basic ideas of game theory;
- It has covered the notion of a stable solution to a game; and
- It has covered pure strategy and mixed strategy solutions.