

Multiagent Systems: Spring 2006

Ulle Endriss

Institute for Logic, Language and Computation

University of Amsterdam

Game Theory

- Last week we have looked into the problem of collective decision making from a *social* point of view — what kind of decision would be good for society?
- Today we are going to analyse the behaviour of *individual* agents in the context of making collective decisions.
- *Game Theory* is the branch of Economic Sciences that studies the strategic behaviour of rational agents in the context of interactive decision-making problems.
- Given the *rules of the “game”* (the *negotiation mechanism*, the *protocol*), what *strategy* should a *rational* agent adopt?

Plan for Today

This is an introduction to Game Theory. In particular, we'll discuss:

- Introductory examples: *Prisoners Dilemma*, *Game of Chicken*, ...
- Distinguishing *dominant strategies* and *equilibrium strategies*
- Distinguishing *pure* and *mixed Nash equilibria*
- *Existence* of mixed Nash equilibria
- *Computing* mixed Nash equilibria

We are going to concentrate on *non-cooperative* (rather than cooperative) *strategic* (rather than extensive) games with *perfect* (rather than imperfect) information.

We'll see later what these distinctions actually mean.

Prisoner's Dilemma

Two partners in crime, A and B , are separated by police and each one of them is offered the following deal:

- only you confess \rightsquigarrow go free
- only the other one confesses \rightsquigarrow spend 5 years in prison
- both confess \rightsquigarrow spend 3 years in prison
- neither one confesses \rightsquigarrow get 1 year on remand

u_A/u_B	B confesses	B does not
A confesses	2/2	5/0
A does not	0/5	4/4

(utility = 5 – years in prison)

► What would be a *rational* strategy?

Dominant Strategies

- A strategy is called (strictly) *dominant* iff, independently of what any of the other players do, following that strategy will result in a larger payoff than any other strategy.
- Prisoner's Dilemma: both players have a dominant strategy, namely to confess:
 - from A 's point of view:
 - * if B confesses, then A is better off confessing as well
 - * if B does not confess, then A is also better off confessing
 - similarly for B
- Terminology: For games of this kind, we say that each player may either *cooperate* with its opponent (e.g. by not confessing) or *defect* (e.g. by confessing).

Battle of the Sexes

Ann (A) and Bob (B) have different preferences as to what to do on a Saturday night ...

u_A/u_B	Bob: theatre	Bob: football
Ann: theatre	2/1	0/0
Ann: football	0/0	1/2

Nash Equilibria

- A *Nash equilibrium* is a set of strategies, one for each player, such that no player could improve their payoff by unilaterally deviating from their assigned strategy (named so after John F. Nash, Nobel Prize in Economic Sciences in 1994; Academy Award in 2001).
- Battle of the Sexes: two Nash equilibria
 - Both Ann and Bob go to the theatre.
 - Both Ann and Bob go to see the football match.
- In cases where there are no dominant strategies, a set of equilibrium strategies is the next best thing.
- Discussion: Games with a Nash equilibrium are of great interest to MAS, because you do not need to keep your strategy secret and you do not need to waste resources on trying to find out about other agents' strategies. Naturally, a *unique* equilibrium is better.

Back to the Prisoner's Dilemma

- Unique Nash equilibrium, namely when both players confess:
 - if A changes strategy unilaterally, she will do worse
 - if B changes strategy unilaterally, she will also do worse
- Discussion: Our analysis shows that it would be *rational* to confess. However, this seems counter-intuitive, because both players would be better off if both of them were to remain silent.
- So there's a conflict: the *stable* solution given by the equilibrium is not *efficient*, because the outcome is not Pareto optimal.
- Iterated Prisoner's Dilemma:
 - In each round, each player can either cooperate or defect.
 - Because the other player could retaliate in the next round, it is rational to cooperate.
 - But it does not work if the number of rounds is fixed ...

Game of Chicken

James and Marlon are driving their cars towards each other at top speed. Whoever swerves to the right first is a “chicken”.

u_J/u_M	M drives on	M turns
J drives on	0/0	8/1
J turns	1/8	5/5

Analysing the Game of Chicken

- No dominant strategy (best move depends on the other player)
- Two Nash equilibria:
 - James drives on and Marlon turns
 - * if James deviates (and turns), he will be worse off
 - * if Marlon deviates (and drives on), he will be worse off
 - Marlon drives on and James turns (similar argument)
- If you have reason to believe your opponent will turn, then you should drive on. If you have reason to believe your opponent will drive on, then you should turn.

How many Nash equilibria?

Keep in mind that the first player chooses the row (T/B) and the second player chooses the column (L/R) ...

	L	R
T	2/2	2/1
B	1/3	3/2

	L	R
T	2/2	2/2
B	2/2	2/2

	L	R
T	1/2	2/1
B	2/1	1/2

Notation and Formal Definition

A *strategic game* consists of a set of *players*, a set of *actions* for each player, and a *preference relation* over action profiles for each player.

- Players: $i \in \{1, \dots, n\}$
- Actions: each player i has a set A_i of possible actions
- Action profiles: $a = (a_1, a_2, \dots, a_n)$ for players $1, \dots, n$
- Preferences: represented by utilities $u_i : A_1 \times \dots \times A_n \rightarrow \mathbb{R}$

Write (a_{-i}, a'_i) for the action profile that is like a , except that player i chooses a'_i rather than a_i .

► Then a *Nash equilibrium* is an action profile a such that $u_i(a) \geq u_i(a_{-i}, a'_i)$ for every player i and every action a'_i of player i .

Remarks

- As we have seen, there are games that have *no* Nash equilibrium.
- Observe that while we use utilities for ease of presentation, only *ordinal preferences matter* (cardinal intensities are irrelevant).
- Here we only model *one-off decisions*. In MAS, in particular, it is however more likely that following a given protocol requires taking a *sequence of decisions*.

But we can map an agent's decision making capability to a single strategy encoding what the agent would do in any given situation. Hence, the game theoretical-models do apply here as well (see also extensive games).

Surprise Exam

Suppose a newspaper announces the following competition:

- ▶ *Every reader may send in a (rational) number between 0 and 100. The winner is the player whose number is closest to $\frac{2}{3}$ times the arithmetic mean of all submissions (in case of a tie the prize money is split equally amongst those with the best guesses).*

What number would you submit (and why)?

A Game without Nash Equilibria

Recall that the following game does not have a Nash equilibrium:

	L	R
T	1/2	2/1
B	2/1	1/2

Whichever action the row player chooses, the column player can react in such a way that the row player would have rather chosen the other way. And so on ...

► Idea: Use a probability distribution over all actions as your strategy.

Mixed Strategies

A *mixed strategy* p_i of a player i is a probability distribution over the actions A_i available to i .

Example: Suppose player 1 has three actions: T, M and B; and suppose their order is clear from the context. Then the mixed strategy to play T with probability $\frac{1}{2}$, M with probability $\frac{1}{6}$, and B with probability $\frac{1}{3}$, is written as $p_1 = (\frac{1}{2}, \frac{1}{6}, \frac{1}{3})$.

The *expected payoff* of a profile p of mixed strategies:

$$E_i(p) = \underbrace{\sum_{a \in A_1 \times \dots \times A_n}}_{\text{sum over all action profiles } a} \left(\overbrace{u_i(a)}^{\text{payoff for } a} \right) \times \underbrace{\prod_{i \in \{1, \dots, n\}} p_i(a_i)}_{\text{probability of choosing } a}$$

Discussion

- Earlier, the numbers in a game matrix represented *ordinal* preferences. In particular, many different sets of numbers would represent the same preference relation.
- Ordinal preferences alone would not allow us to compare “lotteries”:
*I like **appeltaart** more than I like **bitterballen** more than I like a **pistoletje gezond** from the cantine in Euclides ... but this is not enough information to compare bitterballen with a 50-50 chance to win either an appeltaart or a pistoletje.*
- So in the context of mixed strategies, we take the numbers to represent utility functions over deterministic outcomes; and we assume that the preferences of players over alternative mixed strategy profiles are representable by the *expected payoffs* wrt. these utility functions.

Mixed Nash Equilibrium

Write (p_{-i}, p'_i) for the mixed strategy profile that is like p , except that player i chooses p'_i rather than p_i .

► A mixed strategy profile p is a *mixed Nash equilibrium* iff $E_i(p) \geq E_i(p_{-i}, p'_i)$ for every player i and every mixed strategy p'_i for i .

Informally: A mixed Nash equilibrium is a set of mixed strategies, one for each player, so that no player has an incentive to unilaterally deviate from their assigned strategy.

Example

Recall our game without a (pure) Nash equilibrium:

	L	R
T	1/2	2/1
B	2/1	1/2

In this case, guessing probabilities for a mixed Nash equilibrium is easy:

- The row player should play T and B with probability $\frac{1}{2}$ each.
- The column player should play L and R with probability $\frac{1}{2}$ each.

Given the strategy of the column player, the row player has no incentive to deviate (expected payoff is 1.5 either way), and vice versa.

Existence of Mixed Equilibria

We will not prove this central result here:

Theorem 1 (Nash, 1950) *Every finite strategic game has got at least one mixed Nash equilibrium.*

J.F. Nash. *Equilibrium Points in n -Person Games*. Proc. National Academy of Sciences of the United States of America, 36:48–49, 1950.

Computing Mixed Nash Equilibria

Recall the Game of Chicken, now in a more abstract form . . .

	L	R
T	0/0	8/1
B	1/8	5/5

We've already seen that this game has got two pure Nash equilibria. Does it also have a (truly) mixed equilibrium? How can we compute such an equilibrium?

► Note that $((\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, \frac{1}{2}))$ does not work this time (why?).

Best Response of Player 1

Let p (q) be the probability that player 1 (player 2) plays T (L):

	L	R
T	0/0	8/1
B	1/8	5/5

	L	R
T	$p \cdot q$	$p \cdot (1 - q)$
B	$(1 - p) \cdot q$	$(1 - p) \cdot (1 - q)$

Expected payoff for 1 playing T given q : $E_1(\text{T}, q) = q \cdot 0 + (1 - q) \cdot 8$

Expected payoff for 1 playing B given q : $E_1(\text{B}, q) = q \cdot 1 + (1 - q) \cdot 5$

Solving $E_1(\text{T}, q) \geq E_1(\text{B}, q)$ yields $q \leq \frac{3}{4}$.

► The *best response* p of player 1 is given by the following function:

$$p \in \text{best}_1(q) = \begin{cases} \{1\} & \text{if } E_1(\text{T}, q) > E_1(\text{B}, q), \text{ i.e. if } q < \frac{3}{4} \\ [0, 1] & \text{if } E_1(\text{T}, q) = E_1(\text{B}, q), \text{ i.e. if } q = \frac{3}{4} \\ \{0\} & \text{if } E_1(\text{T}, q) < E_1(\text{B}, q), \text{ i.e. if } q > \frac{3}{4} \end{cases}$$

Computing Mixed Nash Equilibria (cont.)

The same kind of reasoning can be used to compute the best response function of player 2 as well (payoffs happen to be symmetric here):

$$q \in best_2(p) = \begin{cases} \{1\} & \text{if } E_2(L, p) > E_2(R, p), \text{ i.e. if } p < \frac{3}{4} \\ [0, 1] & \text{if } E_2(L, p) = E_2(R, p), \text{ i.e. if } p = \frac{3}{4} \\ \{0\} & \text{if } E_2(L, p) < E_2(R, p), \text{ i.e. if } p > \frac{3}{4} \end{cases}$$

Each *intersection* of the two curves corresponds to a mixed Nash equilibrium $((p, 1 - p), (q, 1 - q))$:

$((1, 0), (0, 1))$: player 1 plays T and player 2 plays R [pure]

$((0, 1), (1, 0))$: player 1 plays B and player 2 plays L [pure]

$((\frac{3}{4}, \frac{1}{4}), (\frac{3}{4}, \frac{1}{4}))$: player 1 (2) plays T (L) with probability $\frac{3}{4}$

Complexity of Computing Nash Equilibria

We have just seen a general method for computing all mixed Nash equilibria for a given two-player game with two possible actions each.

In general, computing Nash equilibria is a very difficult problem, but it is not quite clear *how* difficult. According to Papadimitriou (2001),

“... [this] is a most fundamental computational problem whose complexity is wide open.”

It appears to be somewhere “between” P and NP: having guaranteed existence would be untypical for NP-hard problems, but it is not at all clear how to set up a polynomial algorithm either.

STOP PRESS: Now I’ve just seen that there are some papers from late 2005 by Papadimitriou and colleagues that seem to establish some complexity results ... definitely a possible topic for the term paper!

C.H. Papadimitriou. *Algorithms, Games, and the Internet*. Proc. STOC-2001.

Summary

This has been an introduction to Game Theory. You should now know about *dominant strategies* and both *pure* and *mixed equilibrium strategies*. You should also be able to *compute* the mixed Nash equilibria of a simple game.

- We have covered *non-cooperative* rather than *cooperative* games.
 - Cooperative game theory studies competition amongst coalitions of players rather than amongst individuals ...
- We have covered *strategic* rather than *extensive* games.
 - Extensive games model interactions as trees ...
- We have covered games *with perfect information*.
 - Games with imperfect information are used to model situations where the players do not know each others' preferences ...

References

What we have discussed today would be covered by most textbooks on game theory, including these:

- M.J. Osborne. *An Introduction to Game Theory*. Oxford University Press, 2004.
- M.J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
- R.B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1991.
- K. Binmore. *Fun and Games*. Heath, 1992.

The book by Osborne is the most introductory of these, and it has been my main reference for the preparation of this lecture.

What next?

We are going to apply some of the concepts we have learned about to different negotiation scenarios in multiagent systems:

- Negotiation between two agents
- Auctions