CS 798: Multiagent Systems

Normal Form Games

Kate Larson

Cheriton School of Computer Science
University of Waterloo

January 11, 2010

Outline

1. Review
2. Pareto Optimality
3. Nash Equilibria

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
- Outcomes are defined by action profiles, $a = (a_1, \ldots, a_n)$, 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. There sole challenge is to coordinate on an action which is good for all.

Zero-sum games

$\forall a \in A_1 \times A_2, u_1(a) + u_2(a) = 0$. That is, one player gains at the other player's expense.

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

Analyzing Games

We have defined some games, but so far we have no way of discussing what a *good* outcome of a game is.

- Sometimes one outcome o is at least as good for every agent as another outcome o' and there is some agent who strictly prefers o to o'.
  - It seem reasonable to say that o is better than o'
  - We say that o *Pareto dominates* o'.
- An outcome o^* is *Pareto-optimal* if there is no other outcome that Pareto-dominates it.

Analyzing Games

We have defined some games, but so far we have no way of discussing what a *good* outcome of a game is.

- Sometimes one outcome o is at least as good for every agent as another outcome o' and there is some agent who strictly prefers o to o'.
  - It seem reasonable to say that o is better than o'
  - We say that o *Pareto dominates* o'.
- An outcome o^* is *Pareto-optimal* if there is no other outcome that Pareto-dominates it.

Analyzing Games

We have defined some games, but so far we have no way of discussing what a *good* outcome of a game is.

- Sometimes one outcome o is at least as good for every agent as another outcome o' and there is some agent who strictly prefers o to o'.
  - It seem reasonable to say that o is better than o'
  - We say that *o Pareto dominates o'*.
- An outcome o^* is *Pareto-optimal* if there is no other outcome that Pareto-dominates it.

Analyzing Games

We have defined some games, but so far we have no way of discussing what a *good* outcome of a game is.

- Sometimes one outcome o is at least as good for every agent as another outcome o' and there is some agent who strictly prefers o to o'.
  - It seem reasonable to say that o is better than o'
  - We say that *o Pareto dominates o'*.
- An outcome o^* is *Pareto-optimal* if there is no other outcome that Pareto-dominates it.

Examples

Prisoners' Dilemma

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

Pure coordination game

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

Examples

Matching Pennies

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

BoS

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

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 $s = ((\frac{1}{2}, \frac{1}{2}), (\frac{1}{10}, \frac{9}{10}))$

$$u_1 = -1(\frac{1}{2})(\frac{1}{10}) - 4(\frac{1}{2})(\frac{9}{10}) - 3(\frac{1}{2})(\frac{9}{10}) = -3.2$$

$$u_2 = -1(\frac{1}{2})(\frac{1}{10}) - 4(\frac{1}{2})(\frac{1}{10}) - 3(\frac{1}{2})(\frac{9}{10}) = -1.6$$

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!

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!

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!

Interpreting Mixed Strategy Equilibria

What does it mean to play a mixed strategy?

- Randomize to confuse your opponent
- Randomize when you are uncertain about the other's action
- Mixed strategies are a description of what might happen in repeated play
- Mixed strategies describe population dynamics

Interpreting Mixed Strategy Equilibria

What does it mean to play a mixed strategy?

- Randomize to confuse your opponent
- Randomize when you are uncertain about the other's action
- Mixed strategies are a description of what might happen in repeated play
- Mixed strategies describe population dynamics

Interpreting Mixed Strategy Equilibria

What does it mean to play a mixed strategy?

- Randomize to confuse your opponent
- Randomize when you are uncertain about the other's action
- Mixed strategies are a description of what might happen in repeated play
- Mixed strategies describe population dynamics

Interpreting Mixed Strategy Equilibria

What does it mean to play a mixed strategy?

- Randomize to confuse your opponent
- Randomize when you are uncertain about the other's action
- Mixed strategies are a description of what might happen in repeated play
- Mixed strategies describe population dynamics

Interpreting Mixed Strategy Equilibria

What does it mean to play a mixed strategy?

- Randomize to confuse your opponent
- Randomize when you are uncertain about the other's action
- Mixed strategies are a description of what might happen in repeated play
- Mixed strategies describe population dynamics

Interpretations of Nash Equilibria

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