

Introduction to Game Theory from Multi-Agent System Perspective

Michal Pechoucek,
based on slides from Kevin Leyton Brown

Agent Technology Group, Czech Technical University in Prague

October 24, 2010

Content of the course

- 1 Intro to utility theory, basics noncooperative game theory, representation of games in normal form, examples of games, dominance and Pareto efficiency
- 2 Nash equilibrium, mixed strategies, MINIMAX, solution concepts, finding a solution
- 3 Mechanism design, auctions, combinatorial auction and voting
- 4 Negotiation protocol, monotonic concession protocol, contract-net-protocol

Self-interested agents

- What does it mean to say that an agent is **self-interested**?
 - not that they want to harm other agents
 - not that they only care about things that benefit them
 - that the agent has its own description of states of the world that it likes, and that its actions are motivated by this description

Self-interested agents

- What does it mean to say that an agent is **self-interested**?
 - not that they want to harm other agents
 - not that they only care about things that benefit them
 - that the agent has its own description of states of the world that it likes, and that its actions are motivated by this description
- Utility theory:
 - **quantifies** degree of preference across alternatives
 - understand the impact of **uncertainty** on these preferences
 - **utility function**: a mapping from states of the world to real numbers, indicating the agent's level of happiness with that state of the world
 - **Decision-theoretic rationality**: take actions to maximize expected utility.

Example: friends and enemies

- Alice has three options: club (c), movie (m), watching a video at home (h)
- On her own, her utility for these three outcomes is 100 for c , 50 for m and 50 for h
- However, Alice also cares about Bob (who she hates) and Carol (who she likes)
 - Bob is at the club 60% of the time, and at the movies otherwise
 - Carol is at the movies 75% of the time, and at the club otherwise
- If Alice runs into Bob at the movies, she suffers disutility of 40; if she sees him at the club she suffers disutility of 90.
- If Alice sees Carol, she enjoys whatever activity she's doing 1.5 times as much as she would have enjoyed it otherwise (taking into account the possible disutility caused by Bob)

Example: friends and enemies

- Alice has three options: club (c), movie (m), watching a video at home (h)
- On her own, her utility for these three outcomes is 100 for c , 50 for m and 50 for h
- However, Alice also cares about Bob (who she hates) and Carol (who she likes)
 - Bob is at the club 60% of the time, and at the movies otherwise
 - Carol is at the movies 75% of the time, and at the club otherwise
- If Alice runs into Bob at the movies, she suffers disutility of 40; if she sees him at the club she suffers disutility of 90.
- If Alice sees Carol, she enjoys whatever activity she's doing 1.5 times as much as she would have enjoyed it otherwise (taking into account the possible disutility caused by Bob)
- What should Alice do (show of hands)?

What activity should Alice choose?

	$B = c$	$B = m$
$C = c$	15	150
$C = m$	10	100
	$A = c$	

	$B = c$	$B = m$
$C = c$	50	10
$C = m$	75	15
	$A = m$	

What activity should Alice choose?

	$B = c$	$B = m$
$C = c$	15	150
$C = m$	10	100
	$A = c$	

	$B = c$	$B = m$
$C = c$	50	10
$C = m$	75	15
	$A = m$	

- Alice's expected utility for c :

$$0.25(0.6 \cdot 15 + 0.4 \cdot 150) + 0.75(0.6 \cdot 10 + 0.4 \cdot 100) = 51.75.$$
- Alice's expected utility for m :

$$0.25(0.6 \cdot 50 + 0.4 \cdot 10) + 0.75(0.6(75) + 0.4(15)) = 46.75.$$
- Alice's expected utility for h : 50.

Alice prefers to go to the club (though Bob is often there and Carol rarely is), and prefers staying home to going to the movies (though Bob is usually not at the movies and Carol almost always is).

Lecture Overview

- 1 Self-interested agents
- 2 Utility Theory**
- 3 Game Theory
- 4 Example Matrix Games

Why utility?

- Why would anyone argue with the idea that an agent's preferences could be described using a utility function as we just did?

Why utility?

- Why would anyone argue with the idea that an agent's preferences could be described using a utility function as we just did?
 - why should a single-dimensional function be enough to explain preferences over an arbitrarily complicated set of alternatives?
 - Why should an agent's response to uncertainty be captured purely by the *expected value* of his utility function?
- It turns out that the claim that an agent has a utility function is substantive.

Preferences Over Outcomes

If o_1 and o_2 are outcomes

- $o_1 \succeq o_2$ means o_1 is at least as desirable as o_2 .
 - read this as “the agent **weakly prefers** o_1 to o_2 ”
- $o_1 \sim o_2$ means $o_1 \succeq o_2$ and $o_2 \succeq o_1$.
 - read this as “the agent is **indifferent** between o_1 and o_2 .”
- $o_1 \succ o_2$ means $o_1 \succeq o_2$ and $o_2 \not\succeq o_1$
 - read this as “the agent **strictly prefers** o_1 to o_2 ”

Lotteries

- An agent may not know the outcomes of his actions, but may instead only have a probability distribution over the outcomes.

Definition (lottery)

A **lottery** is a probability distribution over outcomes. It is written

$$[p_1 : o_1, p_2 : o_2, \dots, p_k : o_k]$$

where the o_i are outcomes and $p_i > 0$ such that

$$\sum_i p_i = 1$$

- The lottery specifies that outcome o_i occurs with probability p_i .
- We will consider lotteries to be outcomes.

Preference Axioms: Completeness

Definition (Completeness)

A preference relationship must be defined between every pair of outcomes:

$$\forall o_1 \forall o_2 \quad o_1 \succeq o_2 \text{ or } o_2 \succeq o_1$$

Preference Axioms: Transitivity

Definition (Transitivity)

Preferences must be transitive:

if $o_1 \succeq o_2$ and $o_2 \succeq o_3$ then $o_1 \succeq o_3$

- This makes good sense: otherwise $o_1 \succeq o_2$ and $o_2 \succeq o_3$ and $o_3 \succ o_1$.
- An agent should be prepared to pay some amount to swap between an outcome they prefer less and an outcome they prefer more
- Intransitive preferences mean we can construct a “money pump”!

Preference Axioms

Definition (Monotonicity)

An agent prefers a larger chance of getting a better outcome to a smaller chance:

- If $o_1 \succ o_2$ and $p > q$ then

$$[p : o_1, 1 - p : o_2] \succ [q : o_1, 1 - q : o_2]$$

Preference Axioms

Let $P_\ell(o_i)$ denote the probability that outcome o_i is selected by lottery ℓ . For example, if $\ell = [0.3 : o_1; 0.7 : [0.8 : o_2; 0.2 : o_1]]$ then $P_\ell(o_1) = 0.44$ and $P_\ell(o_3) = 0$.

Definition (Decomposability (“no fun in gambling”))

If $\forall o_i \in O, P_{\ell_1}(o_i) = P_{\ell_2}(o_i)$ then $\ell_1 \sim \ell_2$.

Preference Axioms

Definition (Substitutability)

If $o_1 \sim o_2$ then for all sequences of one or more outcomes o_3, \dots, o_k and sets of probabilities p, p_3, \dots, p_k for which

$$p + \sum_{i=3}^k p_i = 1,$$

$$[p : o_1, p_3 : o_3, \dots, p_k : o_k] \sim [p : o_2, p_3 : o_3, \dots, p_k : o_k].$$

Preference Axioms

Definition (Continuity)

Suppose $o_1 \succ o_2$ and $o_2 \succ o_3$, then there exists a $p \in [0, 1]$ such that $o_2 \sim [p : o_1, 1 - p : o_3]$.

Preferences and utility functions

Theorem (von Neumann and Morgenstern, 1944)

If an agent's preference relation satisfies the axioms Completeness, Transitivity, Decomposability, Substitutability, Monotonicity and Continuity then there exists a function $u : O \rightarrow [0, 1]$ with the properties that:

- 1 $u(o_1) \geq u(o_2)$ iff the agent prefers o_1 to o_2 ; and
- 2 when faced about uncertainty about which outcomes he will receive, the agent prefers outcomes that maximize the expected value of u .

Proof idea:

- define the utility of the best outcome $u(\bar{o}) = 1$ and of the worst $u(\underline{o}) = 0$
- now define the utility of each other outcome o as the p for which $o \sim [p : \bar{o}; (1 - p) : \underline{o}]$.

Agent's Rationality

- An agent that violates any of the Axioms is not acting rationally.
- Important Note: The axioms do not mention the utility function!
 - They define rationality by placing constraints on preferences
 - The assumption is that all agents have some mechanism for computing/acting on preferences

Agent's Rationality

- An agent that violates any of the Axioms is not acting rationally.
- Important Note: The axioms do not mention the utility function!
 - They define rationality by placing constraints on preferences
 - The assumption is that all agents have some mechanism for computing/acting on preferences
- The agents rationality is given by the choice of actions based on expected utility of the outcome of the action. Action can be seen as a choice of the appropriate lottery. The rational agent selects to chose such an action a that executes the lottery l that provides the maximal expected outcome:

$$a = \arg \max_{l \in \mathcal{L}} \sum_{p_i: o_i \in l} p_i u(o_i)$$

Agent's Rationality

- Bounded Rationality: capability of the agent to perform rational decision (to choose the lottery providing maximal expected outcome) given bounds on computational resources:
 - bounds on time complexity
 - bounds on memory requirements
- Calculative Rationality: capability to perform rational choice earlier than a fastest change in the environment can occur.

Agent's Rationality

Let us have a community of agents $A_j \in \mathcal{A}$ each choosing to play an action a_j , executing the lottery l_j . providing the agents with the utility $u(a_i)$.

- **Self-interested rational agent:** selects the action that optimizes its individual utility

$$a = \arg \max_{l \in \mathcal{L}} \sum_{p_i: o_i \in l} p_i U(o_i)$$

- **Cooperative rational agent:** selects the action that optimizes collective utility of the whole team:

$$a = \arg \max_{l \in \mathcal{L}} \sum_{\forall a_j \in \mathcal{A} - a} \sum_{p_{i,j}: o_{i,j} \in l_j} p_{i,j} u(o_{i,j}) + \sum_{p_i: o_i \in l} p_i u(o_i)$$

Game Theory Intro

Lecture 3

Lecture Overview

- 1 What is Game Theory?
- 2 Example Matrix Games

Non-Cooperative Game Theory

- What is it?

Non-Cooperative Game Theory

- What is it?
 - mathematical study of interaction between **rational**, **self-interested** agents

Non-Cooperative Game Theory

- What is it?
 - mathematical study of interaction between **rational**, **self-interested** agents
- Why is it called non-cooperative?


Non-Cooperative Game Theory

- What is it?
 - mathematical study of interaction between **rational**, **self-interested** agents
- Why is it called non-cooperative?
 - while it's most interested in situations where agents' interests conflict, it's not restricted to these settings
 - the key is that the individual is the basic modeling unit, and that individuals pursue their own interests
 - cooperative/coalitional game theory has teams as the central unit, rather than agents

TCP Backoff Game



TCP Backoff Game



Should you send your packets using correctly-implemented TCP (which has a “backoff” mechanism) or using a defective implementation (which doesn’t)?

- Consider this situation as a two-player game:
 - **both use a correct implementation:** both get 1 ms delay
 - **one correct, one defective:** 4 ms delay for correct, 0 ms for defective
 - **both defective:** both get a 3 ms delay.

TCP Backoff Game



Should you send your packets using correctly-implemented TCP (which has a “backoff” mechanism) or using a defective implementation (which doesn’t)?

- Consider this situation as a two-player game:
 - **both use a correct implementation:** both get 1 ms delay
 - **one correct, one defective:** 4 ms delay for correct, 0 ms for defective
 - **both defective:** both get a 3 ms delay.
- Play this game with someone near you. Then find a new partner and play again. Play five times in total.

TCP Backoff Game

- Consider this situation as a two-player game:
 - **both use a correct implementation:** both get 1 ms delay
 - **one correct, one defective:** 4 ms delay for correct, 0 ms for defective
 - **both defective:** both get a 3 ms delay.
- Questions:
 - What **action** should a player of the game take?
 - Would all users behave **the same** in this scenario?
 - What global **patterns of behaviour** should the system designer expect?
 - Under what **changes to the delay numbers** would behavior be the same?
 - What effect would **communication** have?
 - **Repetitions?** (finite? infinite?)
 - Does it matter if I believe that my opponent is **rational**?

Defining Games

- Finite, n -person game: $\langle N, A, u \rangle$:
 - N is a finite set of n **players**, indexed by i
 - $A = A_1 \times \dots \times A_n$, where A_i is the **action set** for player i
 - $a \in A$ is an **action profile**, and so A is the space of action profiles
 - $u = \langle u_1, \dots, u_n \rangle$, a **utility function** for each player, where $u_i : A \mapsto \mathbb{R}$
- Writing a 2-player game as a **matrix**:
 - row player is player 1, column player is player 2
 - rows are actions $a \in A_1$, columns are $a' \in A_2$
 - cells are outcomes, written as a tuple of utility values for each player

Games in Matrix Form

Here's the **TCP Backoff Game** written as a matrix (“normal form”).

	<i>C</i>	<i>D</i>
<i>C</i>	-1, -1	-4, 0
<i>D</i>	0, -4	-3, -3

Lecture Overview

- 1 What is Game Theory?
- 2 Example Matrix Games

More General Form

Prisoner's dilemma is any game

	<i>C</i>	<i>D</i>
<i>C</i>	a, a	b, c
<i>D</i>	c, b	d, d

with $c > a > d > b$.

Games of Pure Competition

Players have **exactly opposed** interests

- There must be precisely two players (otherwise they can't have exactly opposed interests)
- For all action profiles $a \in A$, $u_1(a) + u_2(a) = c$ for some constant c
 - Special case: zero sum
- Thus, we only need to store a utility function for one player
 - in a sense, it's a one-player game

Matching Pennies

One player wants to **match**; the other wants to **mismatch**.

	Heads	Tails
Heads	1	-1
Tails	-1	1

Matching Pennies

One player wants to **match**; the other wants to **mismatch**.

	Heads	Tails
Heads	1	-1
Tails	-1	1

Play this game with someone near you, repeating five times.

Rock-Paper-Scissors

Generalized matching pennies.

	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0

...Believe it or not, there's an annual international competition for this game!

Games of Cooperation

Players have **exactly the same** interests.

- no conflict: all players want the same things
- $\forall a \in A, \forall i, j, u_i(a) = u_j(a)$
- we often write such games with a single payoff per cell
- why are such games “noncooperative”?

Coordination Game

Which **side of the road** should you drive on?

	Left	Right
Left	1	0
Right	0	1

Coordination Game

Which **side of the road** should you drive on?

	Left	Right
Left	1	0
Right	0	1

Play this game with someone near you. Then find a new partner and play again. Play five times in total.

General Games: Battle of the Sexes

The most interesting games combine elements of *cooperation and competition*.

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

General Games: Battle of the Sexes

The most interesting games combine elements of *cooperation and competition*.

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

Play this game with someone near you. Then find a new partner and play again. Play five times in total.

From Optimality to Equilibrium

Lecture 4

Lecture Overview

- 1 Recap
- 2 Pareto Optimality
- 3 Best Response and Nash Equilibrium
- 4 Mixed Strategies

Non-Cooperative Game Theory

- What is it?
 - mathematical study of interaction between **rational**, **self-interested** agents
- Why is it called non-cooperative?
 - while it's most interested in situations where agents' interests conflict, it's not restricted to these settings
 - the key is that the individual is the basic modeling unit, and that individuals pursue their own interests
 - cooperative/coalitional game theory has teams as the central unit, rather than agents

Defining Games

- Finite, n -person game: $\langle N, A, u \rangle$:
 - N is a finite set of n **players**, indexed by i
 - $A = A_1 \times \dots \times A_n$, where A_i is the **action set** for player i
 - $a \in A$ is an **action profile**, and so A is the space of action profiles
 - $u = \langle u_1, \dots, u_n \rangle$, a **utility function** for each player, where $u_i : A \mapsto \mathbb{R}$
- Writing a 2-player game as a **matrix**:
 - row player is player 1, column player is player 2
 - rows are actions $a \in A_1$, columns are $a' \in A_2$
 - cells are outcomes, written as a tuple of utility values for each player

Prisoner's dilemma

Prisoner's dilemma is any game

	<i>C</i>	<i>D</i>
<i>C</i>	a, a	b, c
<i>D</i>	c, b	d, d

with $c > a > d > b$.

Games of Pure Competition

Players have **exactly opposed** interests

- There must be precisely two players (otherwise they can't have exactly opposed interests)
- For all action profiles $a \in A$, $u_1(a) + u_2(a) = c$ for some constant c
 - Special case: zero sum

	Heads	Tails
Heads	1	-1
Tails	-1	1

Games of Cooperation

Players have **exactly the same** interests.

- no conflict: all players want the same things
- $\forall a \in A, \forall i, j, u_i(a) = u_j(a)$

	Left	Right
Left	1	0
Right	0	1

General Games: Battle of the Sexes

The most interesting games combine elements of *cooperation and competition*.

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

Lecture Overview

- 1 Recap
- 2 Pareto Optimality
- 3 Best Response and Nash Equilibrium
- 4 Mixed Strategies

Analyzing Games

- We've defined some canonical games, and thought about how to play them. Now let's examine the games from the **outside**
- From the point of view of an outside observer, can some outcomes of a game be said to be **better** than others?

Analyzing Games

- We've defined some canonical games, and thought about how to play them. Now let's examine the games from the **outside**
- From the point of view of an outside observer, can some outcomes of a game be said to be **better** than others?
 - we have no way of saying that one agent's interests are more important than another's
 - intuition: imagine trying to find the revenue-maximizing outcome when you don't know what currency has been used to express each agent's payoff
- Are there situations where we can still prefer one outcome to another?

Pareto Optimality

- **Idea:** 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'
 - in this case, it seems reasonable to say that o is better than o'
 - we say that o **Pareto-dominates** o' .

Pareto Optimality

- **Idea:** 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'
 - in this case, it seems 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.

Pareto Optimality

- **Idea:** 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'
 - in this case, it seems 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.
 - can a game have more than one Pareto-optimal outcome?

Pareto Optimality

- **Idea:** 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'
 - in this case, it seems 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.
 - can a game have more than one Pareto-optimal outcome?
 - does every game have at least one Pareto-optimal outcome?

Pareto Optimal Outcomes in Example Games

	<i>C</i>	<i>D</i>
<i>C</i>	-1, -1	-4, 0
<i>D</i>	0, -4	-3, -3

Pareto Optimal Outcomes in Example Games

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

	Left	Right
Left	1	0
Right	0	1

Pareto Optimal Outcomes in Example Games

	<i>C</i>	<i>D</i>
<i>C</i>	-1, -1	-4, 0
<i>D</i>	0, -4	-3, -3

	Left	Right
Left	1	0
Right	0	1

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

Pareto Optimal Outcomes in Example Games

	<i>C</i>	<i>D</i>
<i>C</i>	-1, -1	-4, 0
<i>D</i>	0, -4	-3, -3

	Left	Right
Left	1	0
Right	0	1

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

	Heads	Tails
Heads	1	-1
Tails	-1	1

Lecture Overview

- 1 Recap
- 2 Pareto Optimality
- 3 Best Response and Nash Equilibrium**
- 4 Mixed Strategies

Best Response

- If you knew what everyone else was going to do, it would be easy to pick your own action

Best Response

- If you knew what everyone else was going to do, it would be easy to pick your own action
- Let $a_{-i} = \langle a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n \rangle$.
 - now $a = (a_{-i}, a_i)$
- **Best response:** $a_i^* \in BR(a_{-i})$ iff
$$\forall a_i \in A_i, u_i(a_i^*, a_{-i}) \geq u_i(a_i, a_{-i})$$

Nash Equilibrium

- Now let's return to the setting where no agent knows anything about what the others will do
- What can we say about which actions will occur?

Nash Equilibrium

- Now let's return to the setting where no agent knows anything about what the others will do
- What can we say about which actions will occur?

- Idea: look for **stable** action profiles.
- $a = \langle a_1, \dots, a_n \rangle$ is a (“pure strategy”) **Nash equilibrium** iff $\forall i, a_i \in BR(a_{-i})$.

Nash Equilibria of Example Games

	<i>C</i>	<i>D</i>
<i>C</i>	-1, -1	-4, 0
<i>D</i>	0, -4	-3, -3

Nash Equilibria of Example Games

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

	Left	Right
Left	1	0
Right	0	1

Nash Equilibria of Example Games

	<i>C</i>	<i>D</i>
<i>C</i>	-1, -1	-4, 0
<i>D</i>	0, -4	-3, -3

	Left	Right
Left	1	0
Right	0	1

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

Nash Equilibria of Example Games

	<i>C</i>	<i>D</i>
<i>C</i>	-1, -1	-4, 0
<i>D</i>	0, -4	-3, -3

	Left	Right
Left	1	0
Right	0	1

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

	Heads	Tails
Heads	1	-1
Tails	-1	1

Nash Equilibria of Example Games

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

	Left	Right
Left	1	0
Right	0	1

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

	Heads	Tails
Heads	1	-1
Tails	-1	1

The paradox of *Prisoner's dilemma*: the Nash equilibrium is the only non-Pareto-optimal outcome!

Lecture Overview

- 1 Recap
- 2 Pareto Optimality
- 3 Best Response and Nash Equilibrium
- 4 Mixed Strategies**

Strategies in game theory

- In game theory, a strategy refers to one of the options that a player can choose. That is, every player in a non-cooperative game has a set of possible strategies, and must choose one of the choices.
- A strategy must specify what action will happen in each contingent state of the game ? e.g. if the opponent does a , then take action b , whereas if the opponent does c , take action d .
- Strategies in game theory may be random (mixed) or deterministic (pure). That is, in some games, players choose mixed strategies. Pure strategies can be thought of as a special case of mixed strategies, in which only probabilities 0 or 1 are assigned to actions.

Mixed Strategies

- It would be a pretty bad idea to play any deterministic strategy in matching pennies
- Idea: confuse the opponent by playing **randomly**
- Define a **strategy** s_i for agent i as any probability distribution over the actions A_i .
 - **pure strategy**: only one action is played with positive probability
 - **mixed strategy**: more than one action is played with positive probability
 - these actions are called the **support** of the mixed strategy
- Let the set of **all strategies** for i be S_i
- Let the set of **all strategy profiles** be $S = S_1 \times \dots \times S_n$.

Utility under Mixed Strategies

- What is your **payoff** if all the players follow mixed strategy profile $s \in S$?
 - We can't just read this number from the game matrix anymore: we won't always end up in the same cell

Utility under Mixed Strategies

- What is your **payoff** if all the players follow mixed strategy profile $s \in S$?
 - We can't just read this number from the game matrix anymore: we won't always end up in the same cell
- Instead, use the idea of **expected utility** from decision theory:

$$u_i(s) = \sum_{a \in A} u_i(a) Pr(a|s)$$

$$Pr(a|s) = \prod_{j \in N} s_j(a_j)$$

Best Response and Nash Equilibrium

Our definitions of best response and Nash equilibrium generalize from actions to strategies.

- **Best response:**

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

Best Response and Nash Equilibrium

Our definitions of best response and Nash equilibrium generalize from actions to strategies.

- **Best response:**

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

- **Nash equilibrium:**

- $s = \langle s_1, \dots, s_n \rangle$ is a Nash equilibrium iff $\forall i, s_i \in BR(s_{-i})$

Best Response and Nash Equilibrium

Our definitions of best response and Nash equilibrium generalize from actions to strategies.

- **Best response:**

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

- **Nash equilibrium:**

- $s = \langle s_1, \dots, s_n \rangle$ is a Nash equilibrium iff $\forall i, s_i \in BR(s_{-i})$

- **Every finite game has a Nash equilibrium!** [Nash, 1950]

- e.g., matching pennies: both players play heads/tails 50%/50%

Computing Mixed Nash Equilibria: Battle of the Sexes

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

- It's hard in general to compute Nash equilibria, but it's easy when you can guess the **support**
- For BoS, let's look for an equilibrium where all actions are part of the support

Computing Mixed Nash Equilibria: Battle of the Sexes

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

- Let player 2 play B with p , F with $1 - p$.
- If player 1 best-responds with a mixed strategy, player 2 must make him indifferent between F and B (why?)

Computing Mixed Nash Equilibria: Battle of the Sexes

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

- Let player 2 play B with p , F with $1 - p$.
- If player 1 best-responds with a mixed strategy, player 2 must make him indifferent between F and B (why?)

$$u_1(B) = u_1(F)$$

$$2p + 0(1 - p) = 0p + 1(1 - p)$$

$$p = \frac{1}{3}$$

Computing Mixed Nash Equilibria: Battle of the Sexes

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

- Likewise, player 1 must randomize to make player 2 indifferent.
 - Why is player 1 willing to randomize?

Computing Mixed Nash Equilibria: Battle of the Sexes

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

- Likewise, player 1 must randomize to make player 2 indifferent.
 - Why is player 1 willing to randomize?
- Let player 1 play B with q , F with $1 - q$.

$$u_2(B) = u_2(F)$$

$$q + 0(1 - q) = 0q + 2(1 - q)$$

$$q = \frac{2}{3}$$

- Thus the mixed strategies $(\frac{2}{3}, \frac{1}{3})$, $(\frac{1}{3}, \frac{2}{3})$ are a Nash equilibrium.

Interpreting Mixed Strategy Equilibria

What does it mean to play a mixed strategy? Different interpretations:

- Randomize to **confuse** your opponent
 - consider the matching pennies example
- Players randomize when they are **uncertain** about the other's action
 - consider battle of the sexes
- Mixed strategies are a concise description of what might happen in **repeated play**: count of pure strategies in the limit
- Mixed strategies describe **population dynamics**: 2 agents chosen from a population, all having deterministic strategies. MS is the probability of getting an agent who will play one PS or another.

Strategic dominance

One of the (simple) metrics that allows comparison of two strategies

- Strategic dominance (commonly called simply dominance) occurs when one strategy is better than another strategy for one player, no matter how that player's opponents may play. Many simple games can be solved using dominance.
- Intransitivity (just the opposite) occurs in games where one strategy may be better or worse than another strategy for one player, depending on how the player's opponents may play.

Strategic dominance

B dominates A: choosing B always gives at least as good an outcome as choosing A. There are 2 possibilities:

- B strictly dominates A: choosing B always gives a better outcome than choosing A, no matter what the other player(s) do.

$$\forall s_{-i} \in S_{-i} : u_i(A, s_{-i}) > u_i(B, s_{-i})$$

- B weakly dominates A: There is at least one set of opponents' action for which B is superior, and all other sets of opponents' actions give B at least the same payoff as A.

$$\forall s_{-i} \in S_{-i} : u_i(A, s_{-i}) \geq u_i(B, s_{-i})$$

with at least one strong inequality