Competition and Cooperation in Multi-Agent Systems

Lecture 1 - Models of Selfish Agents

Vincenzo Bonifaci[1]

Institute for Systems Analysis and Informatics (IASI) – CNR, Italy

October 29, 2012

[1]Slides are from prof. K. Leyton-Brown's MAS class

Scope

The mathematical and computational foundations of modern multiagent systems, with a focus on game theoretic analysis of systems in which agents cannot be guaranteed to behave cooperatively.

References

Y. Shoham, K. Leyton-Brown,
Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations,
Cambridge University Press, 2009
`www.masfoundations.org`

N. Nisan et al. (editors),
Algorithmic Game Theory,
Cambridge University Press, 2007

Both officially available as (non-printable) PDFs online

Lecture Overview

1. Self-interested agents
2. What is Game Theory?
3. Example Matrix Games

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
- We capture this by saying that each agent has a utility function: a mapping from states of the world to real numbers, indicating level of happiness with that state of the world
  - quantifies degree of preference across alternatives
  - allows us to understand the impact of uncertainty on these preferences
  - Decision-theoretic rationality: take actions to maximize expected utility.

Why Utility?

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

Why Utility?

- Why would anyone argue with the idea that an agent's preferences could be described using a utility function?
  - 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.
- There's a famous theorem (von Neumann & Morgenstern, 1944) that derives the existence of a utility function from a more basic preference ordering and axioms on such orderings.
  - see Theorem 3.1.18 in the book, which includes a proof.

Lecture Overview

1. Self-interested agents
2. What is Game Theory?
3. 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

TCP Backoff Game

TCP Backoff Game

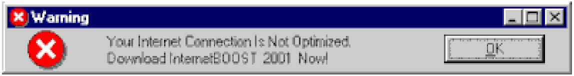

Should you send your packets using correctly-implemented TCP (which has a “backoff” mechanism) or using a defective implementation (which doesn’t)?

- 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.

Defining Games

- Finite, n-person game: $\langle N, A, u \rangle$:
  - N is a finite set of n players, indexed by i
  - $A = A_1 \times \ldots \times A_n$, 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
  - $u = \langle u_1, \ldots, u_n \rangle$, 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").

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

Lecture Overview

1. Self-interested agents
2. What is Game Theory?
3. Example Matrix Games

More General Form

Prisoner's dilemma is any game

	C	D
C	a, a	b, c
D	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

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

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

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

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

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

	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	$0, 0$	$1, 2$

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

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

	Heads	Tails
Heads	1	-1
Tails	-1	1

Lecture Overview

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 $a_{-i} = \langle a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n \rangle$.
    - 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.
- $a = \langle a_1, \ldots, a_n \rangle$ is a ("pure strategy") Nash equilibrium iff $\forall i,\ a_i \in BR(a_{-i})$.

Nash Equilibria of Example Games

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

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

	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	$0, 0$	$1, 2$

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

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

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 $S = S_1 \times \ldots \times S_n$.

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:
  - $s = \langle s_1, \ldots, s_n \rangle$ 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:
  - $s = \langle s_1, \ldots, s_n \rangle$ 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%

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.

A multiplayer game: 2/3 of the mean

- Consider the following game:
  - Everybody bets 1 euro
  - Everyone secretly writes a number between 0 and 100 on a sheet of paper
  - The mean of the numbers is computed and the player(s) closest to 2/3 of the mean split the money
- Model this as a normal form game
- Which are the Nash equilibria, and why?

Lecture Overview

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.

Linear Programming

A linear program is defined by:

- a set of real-valued variables
- a linear objective function
  - a weighted sum of the variables
- a set of linear constraints
  - the requirement that a weighted sum of the variables must be greater than or equal to some constant

Support Enumeration

To compute a Mixed Nash Equilibrium in a 2-player game:

- Enumerate all possible pairs of supports $Z_1 \subseteq A_1$, $Z_2 \subseteq A_2$
- For each pair, check feasibility of this Linear Program:

$$\sum_{a_2 \in Z_2} p_2(a_2) u_1(a_1, a_2) \geq \sum_{a_2 \in Z_2} p_2(a_2) u_1(a', a_2) \qquad \forall a_1 \in Z_1, \forall a' \in A_1$$

$$\sum_{a_1 \in Z_1} p_1(a_1) u_2(a_1, a_2) \geq \sum_{a_1 \in Z_1} p_1(a_1) u_2(a_1, a') \qquad \forall a_2 \in Z_2, \forall a' \in A_2$$

$$\sum_{a_1 \in Z_1} p_1(a_1) = 1, \quad p_1(a_1) \geq 0 \qquad \forall a_1 \in Z_1$$

$$\sum_{a_2 \in Z_2} p_2(a_2) = 1, \quad p_2(a_2) \geq 0 \qquad \forall a_2 \in Z_2.$$

Running time is exponential in $|A_1| + |A_2|$.

Lecture Overview

Maxmin Strategies

- Player i's maxmin strategy is a strategy that maximizes i's worst-case payoff, in the situation where all the other players (whom we denote $-i$) happen to play the strategies which cause the greatest harm to i.
- The maxmin value (or safety level) of the game for player i is that minimum amount of payoff guaranteed by a maxmin strategy.

Definition (Maxmin)

The maxmin strategy for player i is $\arg\max_{s_i} \min_{s_{-i}} u_i(s_1, s_2)$, and the maxmin value for player i is $\max_{s_i} \min_{s_{-i}} u_i(s_1, s_2)$.

- Why would i want to play a maxmin strategy?

Maxmin Strategies

- Player i's maxmin strategy is a strategy that maximizes i's worst-case payoff, in the situation where all the other players (whom we denote $-i$) happen to play the strategies which cause the greatest harm to i.
- The maxmin value (or safety level) of the game for player i is that minimum amount of payoff guaranteed by a maxmin strategy.

Definition (Maxmin)

The maxmin strategy for player i is $\arg\max_{s_i} \min_{s_{-i}} u_i(s_1, s_2)$, and the maxmin value for player i is $\max_{s_i} \min_{s_{-i}} u_i(s_1, s_2)$.

- Why would i want to play a maxmin strategy?
  - a conservative agent maximizing worst-case payoff
  - a paranoid agent who believes everyone is out to get him

Minmax Strategies

- Player i's minmax strategy against player $-i$ in a 2-player game is a strategy that minimizes $-i$'s best-case payoff, and the minmax value for i against $-i$ is payoff.
- Why would i want to play a minmax strategy?

Definition (Minmax, 2-player)

In a two-player game, the minmax strategy for player i against player $-i$ is $\arg\min_{s_i} \max_{s_{-i}} u_{-i}(s_i, s_{-i})$, and player $-i$'s minmax value is $\min_{s_i} \max_{s_{-i}} u_{-i}(s_i, s_{-i})$.

Minmax Strategies

- Player i's minmax strategy against player $-i$ in a 2-player game is a strategy that minimizes $-i$'s best-case payoff, and the minmax value for i against $-i$ is payoff.
- Why would i want to play a minmax strategy?
  - to punish the other agent as much as possible

Definition (Minmax, 2-player)

In a two-player game, the minmax strategy for player i against player $-i$ is $\arg\min_{s_i} \max_{s_{-i}} u_{-i}(s_i, s_{-i})$, and player $-i$'s minmax value is $\min_{s_i} \max_{s_{-i}} u_{-i}(s_i, s_{-i})$.

Minmax Theorem

Theorem (Minimax theorem (von Neumann, 1928))

In any finite, two-player, zero-sum game, in any Nash equilibrium each player receives a payoff that is equal to both his maxmin value and his minmax value.

Minmax Theorem

Theorem (Minimax theorem (von Neumann, 1928))

In any finite, two-player, zero-sum game, in any Nash equilibrium each player receives a payoff that is equal to both his maxmin value and his minmax value.

1. Each player's maxmin value is equal to his minmax value. By convention, the maxmin value for player 1 is called the value of the game.

Minmax Theorem

Theorem (Minimax theorem (von Neumann, 1928))

In any finite, two-player, zero-sum game, in any Nash equilibrium each player receives a payoff that is equal to both his maxmin value and his minmax value.

1. Each player's maxmin value is equal to his minmax value. By convention, the maxmin value for player 1 is called the value of the game.
2. For both players, the set of maxmin strategies coincides with the set of minmax strategies.

Minmax Theorem

Theorem (Minimax theorem (von Neumann, 1928))

In any finite, two-player, zero-sum game, in any Nash equilibrium each player receives a payoff that is equal to both his maxmin value and his minmax value.

1. Each player's maxmin value is equal to his minmax value. By convention, the maxmin value for player 1 is called the value of the game.
2. For both players, the set of maxmin strategies coincides with the set of minmax strategies.
3. Any maxmin strategy profile (or, equivalently, minmax strategy profile) is a Nash equilibrium. Furthermore, these are all the Nash equilibria. Consequently, all Nash equilibria have the same payoff vector (namely, those in which player 1 gets the value of the game).

Saddle Point: Matching Pennies

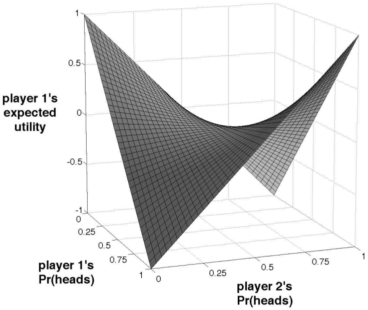

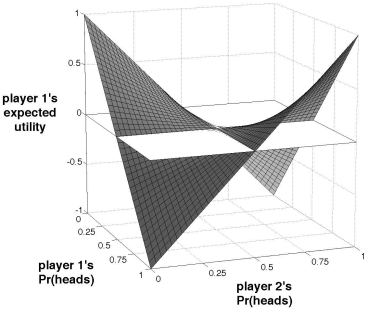

Computing equilibria of zero-sum games

$$\begin{aligned}
\text{minimize} \quad & U_1^* \\
\text{subject to} \quad & \sum_{a_2 \in A_2} u_1(a_1, a_2) \cdot s_2^{a_2} \leq U_1^* && \forall a_1 \in A_1 \\
& \sum_{a_2 \in A_2} s_2^{a_2} = 1 \\
& s_2^{a_2} \geq 0 && \forall a_2 \in A_2
\end{aligned}$$

- First, identify the variables:
  - U_1^* is the expected utility for player 1
  - $s_2^{a_2}$ is player 2's probability of playing action a_2 under his mixed strategy
- each $u_1(a_1, a_2)$ is a constant.

Computing equilibria of zero-sum games

Now let's interpret the LP:

$$\begin{aligned}
\text{minimize} \quad & U_1^* \\
\text{subject to} \quad & \sum_{a_2 \in A_2} u_1(a_1, a_2) \cdot s_2^{a_2} \leq U_1^* & \forall a_1 \in A_1 \\
& \sum_{a_2 \in A_2} s_2^{a_2} = 1 \\
& s_2^{a_2} \geq 0 & \forall a_2 \in A_2
\end{aligned}$$

- s_2 is a valid probability distribution.

Computing equilibria of zero-sum games

Now let's interpret the LP:

$$\begin{array}{lll} \text{minimize} & U_1^* & \\ \text{subject to} & \sum_{a_2 \in A_2} u_1(a_1, a_2) \cdot s_2^{a_2} \leq U_1^* & \forall a_1 \in A_1 \\ & \sum_{a_2 \in A_2} s_2^{a_2} = 1 & \\ & s_2^{a_2} \geq 0 & \forall a_2 \in A_2 \end{array}$$

- U_1^* is as small as possible.

Computing equilibria of zero-sum games

Now let's interpret the LP:

$$\begin{aligned} \text{minimize} \quad & U_1^* \\ \text{subject to} \quad & \sum_{a_2 \in A_2} u_1(a_1, a_2) \cdot s_2^{a_2} \leq U_1^* && \forall a_1 \in A_1 \\ & \sum_{a_2 \in A_2} s_2^{a_2} = 1 \\ & s_2^{a_2} \geq 0 && \forall a_2 \in A_2 \end{aligned}$$

- Player 1's expected utility for playing each of his actions under player 2's mixed strategy is no more than U_1^*.
  - Because U_1^* is minimized, this constraint will be tight for some actions: the support of player 1's mixed strategy.

Computing equilibria of zero-sum games

$$
\begin{array}{lll}
\text{minimize} & U_1^* & \\
\text{subject to} & \displaystyle\sum_{a_2 \in A_2} u_1(a_1, a_2) \cdot s_2^{a_2} \leq U_1^* & \forall a_1 \in A_1 \\
& \displaystyle\sum_{a_2 \in A_2} s_2^{a_2} = 1 & \\
& s_2^{a_2} \geq 0 & \forall a_2 \in A_2
\end{array}
$$

- This formulation gives us the minmax strategy for player 2.
- To get the minmax strategy for player 1, we need to solve a second (analogous) LP.

Lecture Overview

Domination

- Let s_i and s'_i be two strategies for player i, and let S_{-i} be is the set of all possible strategy profiles for the other players

Definition

s_i strictly dominates s'_i if $\forall s_{-i} \in S_{-i}$, $u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i})$

Definition

s_i weakly dominates s'_i if $\forall s_{-i} \in S_{-i}$, $u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$ and $\exists s_{-i} \in S_{-i}$, $u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i})$

Definition

s_i very weakly dominates s'_i if $\forall s_{-i} \in S_{-i}$, $u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$

Equilibria and dominance

- If one strategy dominates all others, we say it is dominant.
- A strategy profile consisting of dominant strategies for every player must be a Nash equilibrium.
  - An equilibrium in strictly dominant strategies must be unique.

Equilibria and dominance

- If one strategy dominates all others, we say it is dominant.
- A strategy profile consisting of dominant strategies for every player must be a Nash equilibrium.
  - An equilibrium in strictly dominant strategies must be unique.
- Consider Prisoner's Dilemma again
  - not only is the only equilibrium the only non-Pareto-optimal outcome, but it's also an equilibrium in strictly dominant strategies!

Lecture Overview

Dominated strategies

- No equilibrium can involve a strictly dominated strategy
  - Thus we can remove it, and end up with a strategically equivalent game
  - This might allow us to remove another strategy that wasn't dominated before
  - Running this process to termination is called iterated removal of dominated strategies.

Iterated Removal of Dominated Strategies: Example

	L	C	R
U	$3, 1$	$0, 1$	$0, 0$
M	$1, 1$	$1, 1$	$5, 0$
D	$0, 1$	$4, 1$	$0, 0$

Iterated Removal of Dominated Strategies: Example

	L	C	R
U	$3, 1$	$0, 1$	$0, 0$
M	$1, 1$	$1, 1$	$5, 0$
D	$0, 1$	$4, 1$	$0, 0$

- R is dominated by L.

Iterated Removal of Dominated Strategies: Example

	L	C
U	$3, 1$	$0, 1$
M	$1, 1$	$1, 1$
D	$0, 1$	$4, 1$

Iterated Removal of Dominated Strategies: Example

	L	C
U	$3,1$	$0,1$
M	$1,1$	$1,1$
D	$0,1$	$4,1$

- M is dominated by the mixed strategy that selects U and D with equal probability.

Iterated Removal of Dominated Strategies: Example

	L	C
U	$3, 1$	$0, 1$
D	$0, 1$	$4, 1$

Iterated Removal of Dominated Strategies: Example

	L	C
U	$3, 1$	$0, 1$
D	$0, 1$	$4, 1$

- No other strategies are dominated.

Iterated Removal of Dominated Strategies

- This process preserves Nash equilibria.
  - strict dominance: all equilibria preserved.
  - weak or very weak dominance: at least one equilibrium preserved.
- Thus, it can be used as a preprocessing step before computing an equilibrium
  - Some games are solvable using this technique.
  - Example: Traveler's Dilemma!
- What about the order of removal when there are multiple dominated strategies?
  - strict dominance: doesn't matter.
  - weak or very weak dominance: can affect which equilibria are preserved.

Lecture Overview

Pithy Quote

If there is intelligent life on other planets, in a majority of them, they would have discovered correlated equilibrium before Nash equilibrium.

– Roger Myerson

Examples

- Consider again Battle of the Sexes.
  - Intuitively, the best outcome seems a 50-50 split between (F, F) and (B, B).
  - But there's no way to achieve this, so either someone loses out (unfair) or both players often miscoordinate
- Another classic example: traffic game

	go	$wait$
go	$-100, -100$	$10, 0$
B	$0, 10$	$-10, -10$

Intuition

- What is the natural solution here?

Intuition

- What is the natural solution here?
  - A traffic light: a fair randomizing device that tells one of the agents to go and the other to wait.
- Benefits:
  - the negative payoff outcomes are completely avoided
  - fairness is achieved
  - the sum of social welfare exceeds that of any Nash equilibrium
- We could use the same idea to achieve the fair outcome in battle of the sexes.
- Our example presumed that everyone perfectly observes the random event; not required.
- More generally, some random variable with a commonly known distribution, and a private signal to each player about the outcome.
  - signal doesn't determine the outcome or others' signals; however, correlated

Formal definition

Definition (Correlated equilibrium)

Given an n-agent game $G = (N, A, u)$, a correlated equilibrium is a probability distribution $(p(a))_{a \in A}$ on the space of strategy profiles such that for each player i and every two strategies a_i, a'_i of i,

$$\sum_{a \in A | a_i \in a} p(a) u_i(a_i, a_{-i}) \geq \sum_{a \in A | a_i \in a} p(a) u_i(a'_i, a_{-i}).$$

Conditioned on the event that a contains a_i, the expected utility of playing a_i is no smaller than that of playing a'_i.

Existence

Theorem

For every mixed Nash equilibrium $(s_1, s_2, \ldots, s_n)$ *there exists a corresponding correlated equilibrium* $(p(a)_{a \in A})$.

- This is easy to show:
  - let $p(a) = \prod_{i \in N} s_i(a_i)$
- Thus, correlated equilibria always exist

Remarks

- Not every correlated equilibrium is equivalent to a Nash equilibrium
  - thus, correlated equilibrium is a weaker notion than Nash
- Any convex combination of the payoffs achievable under correlated equilibria is itself realizable under a correlated equilibrium
  - start with the Nash equilibria (each of which is a CE)
  - introduce a second randomizing device that selects which CE the agents will play
  - regardless of the probabilities, no agent has incentive to deviate
  - the probabilities can be adjusted to achieve any convex combination of the equilibrium payoffs
  - the randomizing devices can be combined

Lecture Overview

Computing CE

$$\sum_{a \in A | a_i \in a} p(a) u_i(a) \geq \sum_{a \in A | a_i \in a} p(a) u_i(a'_i, a_{-i}) \quad \forall i \in N, \forall a_i, a'_i \in A_i$$

$$p(a) \geq 0 \quad \forall a \in A$$

$$\sum_{a \in A} p(a) = 1$$

- variables: $p(a)$; constants: $u_i(a)$

Computing CE

$$\sum_{a \in A | a_i \in a} p(a) u_i(a) \geq \sum_{a \in A | a_i \in a} p(a) u_i(a'_i, a_{-i}) \quad \forall i \in N, \forall a_i, a'_i \in A_i$$

$$p(a) \geq 0 \quad \forall a \in A$$

$$\sum_{a \in A} p(a) = 1$$

- variables: $p(a)$; constants: $u_i(a)$
- we could find the social-welfare maximizing CE by adding an objective function

$$\text{maximize:} \quad \sum_{a \in A} p(a) \sum_{i \in N} u_i(a).$$

Why are CE easier to compute than NE?

$$\sum_{a \in A | a_i \in a} p(a) u_i(a) \geq \sum_{a \in A | a_i' \in a} p(a) u_i(a_i', a_{-i}) \quad \forall i \in N, \forall a_i, a_i' \in A_i$$

$$p(a) \geq 0 \quad \forall a \in A$$

$$\sum_{a \in A} p(a) = 1$$

- intuitively, correlated equilibrium has only a single randomization over outcomes, whereas in NE this is constructed as a product of independent probabilities.
- To change this program so that it finds NE, the first constraint would be

$$\sum_{a \in A} u_i(a) \prod_{j \in N} p_j(a_j) \geq \sum_{a \in A} u_i(a_i', a_{-i}) \prod_{j \in N \setminus \{i\}} p_j(a_j) \quad \forall i \in N, \forall a_i' \in A_i.$$

- This is a nonlinear constraint!