

CS 598:
AI Methods for Market Design

Lecture 2: Intro to Game Theory

Xintong Wang
Spring 2024

Logistics

- Pre-class CQs
 - Due before each lecture
 - Binary grading scheme
 - Two chances to drop
- Paper presentations
 - Bidding on papers
 - Presentation guidelines

Outline

- Simultaneous-move games
 - Normal-form representation
 - Solution concepts
 - Succinct representations
- Sequential-move games
 - Extensive-form representation
 - Solution concepts
 - Repeated games
 - Stackelberg games

Outline

- Simultaneous-move games
 - Normal-form representation
 - Solution concepts
 - Succinct representations
- Sequential-move games
 - Extensive-form representation
 - Solution concepts
 - Repeated games
 - Stackelberg games

Example 1: Prisoners' Dilemma

- Two people are arrested and accused of a crime
- They are questioned in separate rooms (no communication)
- Each prisoner has two choices, Cooperate or Defect, and gets different payoffs depending on the outcome:

		Player 2	
		C	D
Player 1	C	-1, -1	-5, 0
	D	0, -5	-3, -3

Example 1: Prisoners' Dilemma

- Two people are arrested and accused of a crime
- They are questioned in separate rooms (no communication)
- Each prisoner has two choices, Cooperate or Defect, and gets different payoffs depending on the outcome:

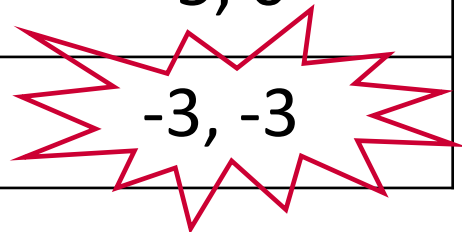
		Player 2	
		C	D
Player 1	C	-1, -1	-5, 0
	D	0, -5	-3, -3

How should each player act?

Example 1: Prisoners' Dilemma

- Two people are arrested and accused of a crime
- They are questioned in separate rooms (no communication)
- Each prisoner has two choices, Cooperate or Defect, and gets different payoffs depending on the outcome:

		Player 2	
		C	D
Player 1	C	-1, -1	-5, 0
	D	0, -5	-3, -3



How should each player act?

Components of a Game

- **Agents / Players**: participants of the game, may be an individual, organization, a machine or algorithm...
- **Strategies**: actions available to each player
- **Outcome**: the profile of player strategies
- **Payoffs**: a function mapping an outcome to a utility / payoff for each player

Simultaneous-Move Game

A **simultaneous-move game** has (N, A, u)

- $N = \{1, 2, \dots, n\}$ **agents**, indexed by i
- $A = A_1 \times \dots \times A_n$, where each agent plays an **action** $a_i \in A_i$ and the **action profile** is $a = (a_1, \dots, a_n) \in A$
 - As a convention, $a_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n)$
- $u = (u_1, \dots, u_n)$, where $u_i: A \rightarrow \mathbb{R}$ is a **utility function (or payoff function)** for agent i , and assigns a utility (or payoff) to every action profile $a \in A$

Simultaneous-Move Game

Some notes on **simultaneous-move game**

- **Simultaneous**: each agent selects an action *without* knowledge about the actions that are selected by others
- **Complete information**: every agent knows the available actions and utility functions of all agents
 - $\{A_i, u_i\}_{i \in [n]}$ are public knowledge

Outline

- Simultaneous-move games
 - Normal-form representation
 - Solution concepts
 - Succinct representations
- Sequential-move games
 - Extensive-form representation
 - Solution concepts
 - Repeated games
 - Stackelberg games

Normal-Form Representation

The **normal-form representation** of a simultaneous-move game (N, A, u) represents the payoffs to agents as a **payoff matrix**

		Player 2	
		C	D
Player 1	C	-1, -1	-5, 0
	D	0, -5	-3, -3

What is the dimension of a payoff matrix with n players, each with m actions?

Example 1: Prisoners' Dilemma

- 2 agents: $N = \{1, 2\}$
- $A_1 = A_2 = \{C, D\}$ and $a \in A = A_1 \times A_2 = \{(C,C), (C,D), (D,C), (D,D)\}$
- $u_1(\cdot)$ and $u_2(\cdot)$ are predefined
 - $u_1(C, C) = -1, u_1(C, D) = -5, u_1(D, C) = 0, u_1(D, D) = -3$
- The whole game is public knowledge
- Agents take actions without knowing others' choice of action

		Player 2	
		C	D
Player 1	C	-1, -1	-5, 0
	D	0, -5	-3, -3

Outline

- Simultaneous-move games
 - Normal-form representation
 - Solution concepts
 - Succinct representations
- Sequential-move games
 - Extensive-form representation
 - Solution concepts
 - Repeated games
 - Stackelberg games

Pareto Optimality

- An action profile $a \in A$ is **Pareto dominated** by another action profile $a' \in A$ if and only if

$$u_i(a') \geq u_i(a) \text{ for all agents } i \in N \text{ and} \\ u_i(a') > u_i(a) \text{ for some agent } i \in N$$

- For example, action profile (D,D) is Pareto dominated by (C,C)

		Player 2	
		C	D
Player 1	C	-1, -1	-5, 0
	D	0, -5	-3, -3

Pareto Optimality

- An action profile $a \in A$ is **Pareto optimal** *if and only if* there is no action profile $a' \in A$ that Pareto dominates a
- For example, action profile (C,C) is Pareto optimal

Any other Pareto optimal action profile?

		Player 2	
		C	D
Player 1	C	-1, -1	-5, 0
	D	0, -5	-3, -3

Dominant Strategy

- An action $a_i \in A_i$ is a **dominant strategy** for player i if a_i is better than any other action $a'_i \in A_i$, *regardless* what actions other players take

$$u_i(a_i, a_{-i}) \geq u_i(a'_i, a_{-i}), \quad \forall a'_i \neq a_i \quad \forall a_{-i}$$

- “Defect” is a dominant strategy for both agents

		Player 2	
		C	D
Player 1	C	-1, -1	-5, 0
	D	0, -5	-3, -3

Dominant Strategy

- Dominant strategies do *not* always exist
 - Consider the following two-player, three-action game

		Player 2		
		L	M	R
Player 1	U	4, 3	5, 1	6, 2
	M	2, 1	8, 4	3, 6
	D	3, 0	9, 6	2, 8

Is there a dominant strategy?

Is there a strictly dominated strategy?

Dominant-Strategy Equilibrium (DSE)

- An action profile $a^* = (a_1^*, \dots, a_n^*) \in A$ is a **dominant-strategy equilibrium (DSE)** if and only if

$$u_i(a_i^*, a_{-i}) \geq u_i(a_i, a_{-i}), \quad \forall i \in n, a_i \in A_i, a_{-i} \in A_{-i}$$

- (D, D) is a dominant-strategy equilibrium
- Predictive power: no need to reason about others' actions!

		Player 2	
		C	D
Player 1	C	-1, -1	-5, 0
	D	0, -5	-3, -3

Dominant-Strategy Equilibrium (DSE)

- Dominant-strategy equilibrium do *not* always exist
 - Consider the following two-player, three-action game

		Player 2		
		L	M	R
Player 1	U	4, 3	5, 1	6, 2
	M	2, 1	8, 4	3, 6
	D	3, 0	9, 6	2, 8

Is there a DSE?

How about other solution concepts?

Pure-Strategy Nash Equilibrium (PSNE)

- An action profile $a^* = (a_1^*, \dots, a_n^*) \in A$ is a **pure-strategy Nash equilibrium** if and only if

$$u_i(a_i^*, a_{-i}^*) \geq u_i(a_i, a_{-i}^*), \quad \forall i \in n, a_i \in A_i$$

- Every agent plays a best response to the actions of others

		Player 2		
		L	M	R
Player 1	U	4, 3	5, 1	6, 2
	M	2, 1	8, 4	3, 6
	D	3, 0	9, 6	2, 8

Which action profile is a PSNE?

Pure-Strategy Nash Equilibrium (PSNE)

- An action profile $a^* = (a_1^*, \dots, a_n^*) \in A$ is a **pure-strategy Nash equilibrium** if and only if

$$u_i(a_i^*, a_{-i}^*) \geq u_i(a_i, a_{-i}^*), \quad \forall i \in n, a_i \in A_i$$

- Every agent plays a best response to the actions of others

		Player 2		
		L	M	R
Player 1	U	4, 3	5, 1	6, 2
	M	2, 1	8, 4	3, 6
	D	3, 0	9, 6	2, 8

Which action profile is a PSNE?

Pure-Strategy Nash Equilibrium (PSNE)

Some notes on Nash equilibrium:

- Require common knowledge of rationality
- Serve a sensible prediction of behavior

		Player 2		
		L	M	R
Player 1	U	4, 3	5, 1	6, 2
	M	2, 1	8, 4	3, 6
	D	3, 0	9, 6	2, 8

Which action profile is a PSNE?

Finding Nash Equilibrium: First Attempt

Iterated elimination of strictly dominated actions

- Step 1: remove action M for player 2

		Player 2		
		L	M	R
Player 1	U	4, 3	5, 1	6, 2
	M	2, 1	8, 4	3, 6
	D	3, 0	9, 6	2, 8

Which action profile is a PSNE?

Finding Nash Equilibrium: First Attempt

Iterated elimination of strictly dominated actions

- Step 1: remove action M for player 2
- Step 2: remove actions M and D for player 1 (M and D are strictly dominated by U as long as player 2 selects L or R)

		Player 2		
		L	M	R
Player 1	U	4, 3		6, 2
	M	2, 1		3, 6
	D	3, 0		2, 8

Which action profile is a PSNE?

Finding Nash Equilibrium: First Attempt

Iterated elimination of strictly dominated actions

- Step 1: remove action M for player 2
- Step 2: remove actions M and D for player 1 (M and D are strictly dominated by U, if player 2 selects L or R)
- Step 3: player 2 plays L, if player 1 chooses U, so (U, L)

		Player 2		
		L	M	R
Player 1	U	4, 3		6, 2
	M			
	D			

Which action profile is a PSNE?

Finding Nash Equilibrium: First Attempt

Iterated elimination of strictly dominated actions

Questions to think about:

- What is the time complexity of iterated elimination of strictly dominated actions by pure actions, for a game of n players each with m actions?
- Will iterated elimination of *weakly* dominated actions work in finding NE?

Pure-Strategy Nash Equilibrium (PSNE)

- Pure-strategy Nash equilibrium does *not* always exist
 - Consider rock-paper-scissor

	Rock	Paper	Scissor
Rock	(0, 0)	(-1, 1)	(1, -1)
Paper	(1, -1)	(0, 0)	(-1, 1)
Scissor	(-1, 1)	(1, -1)	(0, 0)

Pure-Strategy Nash Equilibrium (PSNE)

- Pure-strategy Nash equilibrium does *not* always exist
 - Consider the Matching Pennies game

		Player 2	
		H	T
Player 1	H	1, -1	-1, 1
	T	-1, 1	1, -1

Mixed Strategy

- Pure strategy: take an action deterministically
- Mixed strategy: randomize over actions
 - Described by a distribution s_i where $s_i(a_i) \geq 0$ denotes the probability of taking an action a_i
 - $|A_i|$ -dimensional **simplex** $\Delta(A_i) := \{s_i: \sum_{a_i \in A_i} s_i(a_i) = 1\}$ contains all possible mixed strategies for player i
 - Each agent **independently** draws an action based on its mixed strategy s_i

Mixed Strategy

- A **strategy profile** is then $s = (s_1, \dots, s_n)$
- **The probability of action profile** $a = (a_1, \dots, a_n)$ is then $p(a) = \prod_{i \in [n]} s_i(a_i)$ due to independence
- Given a strategy profile $s = (s_1, \dots, s_n)$, **the expected utility of agent i** is

$$u_i(s) = \sum_{a \in A} u_i(a) \cdot p(a) = \sum_{a \in A} u_i(a) \cdot \prod_{i \in [n]} s_i(a_i)$$

Mixed Strategy

- Exercise: Given strategy $s_1 = (0.4, 0.6)$ for player 1 and strategy $s_2 = (1, 0)$ for player 2, what is the expected utility for player 1?

		Player 2	
		H	T
Player 1	H	1, -1	-1, 1
	T	-1, 1	1, -1

$$u_i(s) = \sum_{a \in A} u_i(a) \cdot p(a) = \sum_{a \in A} u_i(a) \cdot \prod_{i \in [n]} s_i(a_i)$$

Mixed-Strategy Nash Equilibrium (MSNE)

- A **strategy profile** $s^* = (s_1^*, \dots, s_n^*)$ is a **mixed-strategy Nash equilibrium** *if and only if*

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*), \quad \forall i \in n, s_i \in \Delta(A_i)$$

- Every agent plays a best response to the **strategies** of others

Pure-Strategy Nash Equilibrium

- An action profile $a^* = (a_1^*, \dots, a_n^*) \in A$ is a **pure-strategy Nash equilibrium** *if and only if*

$$u_i(a_i^*, a_{-i}^*) \geq u_i(a_i, a_{-i}^*), \quad \forall i \in n, a_i \in A_i$$

- Every agent plays a best response to the **actions** of others

Mixed-Strategy Nash Equilibrium (MSNE)

Some notes on best responses:

- The **support of mixed strategy** s_i is the set of actions played with strictly-positive probability, i.e.,

$$\sigma(s_i) = \{a_i : s_i(a_i) > 0, a_i \in A_i\} \subseteq A_i$$

- A useful property: all actions in the support of s_i^* have the same expected utility
- A **strategy profile** $s^* = (s_1^*, \dots, s_n^*)$ is a **mixed-strategy Nash equilibrium** if and only if

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*), \quad \forall i \in n, s_i \in \Delta(A_i)$$

$$u_i(a_i, s_{-i}^*) \geq u_i(s_i, s_{-i}^*), \quad \forall i \in n, a_i \in \sigma(s_i^*), s_i \in \Delta(A_i)$$

$$u_i(a_i, s_{-i}^*) \geq u_i(a'_i, s_{-i}^*), \quad \forall i \in n, a_i \in \sigma(s_i^*), a'_i \in A_i$$

Mixed-Strategy Nash Equilibrium (MSNE)

What is the mixed-strategy Nash equilibrium?

		Player 2	
		H	T
Player 1	H	1, -1	-1, 1
	T	-1, 1	1, -1

Mixed-Strategy Nash Equilibrium (MSNE)

Theorem (Nash, 1951): **Every finite simultaneous-move game has at least one mixed-strategy Nash equilibrium**

- A fundamental result in game theory
- An equilibrium outcome is *not* necessarily the best for players
 - Describe where the game may stabilize at
 - Understand how self-interested behaviors reduces overall social welfare (Price of Anarchy (PoA))

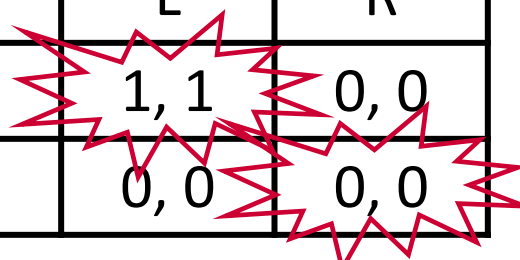
		Player 2		PoA = 3
		C	D	
Player 1	C	-1, -1	-5, 0	
	D	0, -5	-3, -3	

Mixed-Strategy Nash Equilibrium (MSNE)

Theorem (Nash, 1951): **Every finite simultaneous-move game has at least one mixed-strategy Nash equilibrium**

- A game may have many, even infinitely many, NEs
 - When facing multiple equilibria, may need additional assumptions

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



	L	R
L	3, 1	0, 1
R	0, 1	4, 1

Mixed-Strategy Nash Equilibrium (MSNE)

Theorem (Nash, 1951): **Every finite simultaneous-move game has at least one mixed-strategy Nash equilibrium**


Equilibrium is a **prediction** of agent behaviors and a prediction of system outcomes from strategic interactions

- ML: data-driven
- Equilibrium analysis: model-driven & data-driven
 - Learn what game agents are playing (e.g., game parameters)
 - Learn payoff functions
 - Learn how rational agents are (e.g., behavioral economics)
 - ...

Example 2: Traffic Light Game

- Exercise: What are the equilibria of the game?


		Player 2	
		W	G
Player 1	W	0, 0	0, 2
	G	2, 0	-4, -4



Example 2: Traffic Light Game

- Exercise: What are the equilibria of the game?

		Player 2	
		W	G
Player 1	W	0, 0	0, 2
	G	2, 0	-4, -4



Two pure-strategy NE: (G, W) and (W, G)


A mixed-strategy NE: (2/3, 1/3) for both players

Chance of a crash: 1/9, when *each agent draws an action independently*

Example 2: Traffic Light Game

- Introduce an external, shared signal: the traffic light


		Player 2	
		W	G
Player 1	W	0, 0	0, 2
	G	2, 0	-4, -4



Example 2: Traffic Light Game

- Introduce an external, shared signal: the traffic light

		Player 2	
		W	G
Player 1	W	0, 0	0, 2
	G	2, 0	-4, -4




- Each player's strategy selects an action
 - $a_1 = G$ if signal is 0, $a_1 = W$ otherwise
 - $a_2 = W$ if signal is 0, $a_2 = G$ otherwise
 - An equilibrium: each agent plays a best response to the other, **conditioned on the signal**

Correlated Equilibrium (CE)

- A **recommendation policy** π assigns probability $\pi(a)$ for each action profile $a \in A$
- A mediator samples $a \sim \pi$, then recommends a_i to agent i
- A (fair) correlated equilibrium: $\begin{pmatrix} 0 & 1/2 \\ 1/2 & 0 \end{pmatrix}$
- Nash equilibria: $\begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}$ $\begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$ $\begin{pmatrix} 4/9 & 2/9 \\ 2/9 & 1/9 \end{pmatrix}$

		Player 2	
		W	G
Player 1	W	0, 0	0, 2
	G	2, 0	-4, -4



Correlated Equilibrium (CE)

$a_i \in \sigma(\pi_i)$: an action a_i may be suggested to player i

$\pi_{-i}(a_{-i} | a_i)$: the probability of $a_{-i} \in A_{-i}$ suggested for others, conditioned on action a_i being suggested to agent i

- A probability distribution π on action profiles A is a **correlated equilibrium** *if and only if*

$$\sum_{a_{-i} \in A_{-i}} u_i(a_i, a_{-i}) \cdot \pi_{-i}(a_{-i} | a_i) \geq \sum_{a_{-i} \in A_{-i}} u_i(a'_i, a_{-i}) \cdot \pi_{-i}(a_{-i} | a_i),$$

$\forall i \in n, a_i \in \sigma(\pi_i), a'_i \in A_i$

Correlated Equilibrium (CE)

Some notes on correlated equilibrium

- π is public knowledge
- No agent wants to deviate from its suggested action, assuming others also follow their suggested actions
- **When actions are drawn independently**, $\pi_{-i}(a_{-i} | a_i)$ is the product of the marginal probability with other players play their corresponding action in a_{-i}

Equivalent to NE!


Correlated Equilibrium

- Fact: Any Nash equilibrium is also a correlated equilibrium
- Corollary: Every finite, simultaneous-move game has at least one correlated equilibrium
- How to compute correlated equilibrium?

Correlated Equilibrium

In practice, what are the correlation devices?

- Traffic lights
- Google Maps
- A shared history of play



Coarse Correlated Equilibrium (CCE)

- A **weaker** notion of correlated equilibrium
- A probability distribution π on action profiles A is a **coarse correlated equilibrium** *if and only if*

$$\sum_{a \in A} u_i(a) \cdot \pi(a) \geq \sum_{a \in A} u_i(a'_i, a_{-i}) \cdot \pi(a), \quad \forall i \in n, a'_i \in A_i$$

Correlated Equilibrium (CE)

$a_i \in \sigma(\pi_i)$: an action a_i may be suggested to player i

$\pi_{-i}(a_{-i} | a_i)$: the probability of $a_{-i} \in A_{-i}$ suggested for others, **conditioned on action a_i being suggested to agent i**

- A probability distribution π on action profiles A is a **correlated equilibrium** *if and only if*

$$\sum_{a_{-i} \in A_{-i}} u_i(a_i, a_{-i}) \cdot \pi_{-i}(a_{-i} | a_i) \geq \sum_{a_{-i} \in A_{-i}} u_i(a'_i, a_{-i}) \cdot \pi_{-i}(a_{-i} | a_i),$$

$\forall i \in n, a_i \in \sigma(\pi_i), a'_i \in A_i$

Coarse Correlated Equilibrium (CCE)

- A **weaker** notion of correlated equilibrium
- A probability distribution π on action profiles A is a **coarse correlated equilibrium** *if and only if*


$$\sum_{a \in A} u_i(a) \cdot \pi(a) \geq \sum_{a \in A} u_i(a'_i, a_{-i}) \cdot \pi(a), \quad \forall i \in n, a'_i \in A_i$$

- CCE vs. CE: after an action profile is drawn
 - CCE: playing a_i is a best response for player i , **in expectation before** seeing a_i
 - CE: playing a_i is a best response for player i , **conditioned on** seeing a_i

$$\sum_{a_i} \sum_{a_{-i} \in A_{-i}} u_i(a_i, a_{-i}) \cdot \pi_{-i}(a_{-i} | a_i) \geq \sum_{a'_i} \sum_{a_{-i} \in A_{-i}} u_i(a'_i, a_{-i}) \cdot \pi_{-i}(a_{-i} | a_i),$$

$\forall i \in n, a_i \in \sigma(\pi_i), a'_i \in A_i$

Equilibrium Hierarchy for Simultaneous-Move Games



Outline


- Simultaneous-move games
 - Normal-form representation
 - Solution concepts
 - Succinct representations
- Sequential-move games
 - Extensive-form representation
 - Solution concepts
 - Repeated games
 - Stackelberg games

Succinct Representations

- Congestion games
- Agent-graph games
- Action-graph games

Example 3: Network Flow


- There are 2,000 people who commute to work everyday from point “Start” to point “End”
- Every driver has to choose a path, without seeing what others do



How to represent this simultaneous-move game?

Example 3: Network Flow

- There are 2,000 people who commute to work everyday from point “Start” to point “End”
- Every driver has to choose a path, without seeing what others do



How to represent this simultaneous-move game?

3^{2000} possible action profiles

Congestion Games


A **congestion game** (N, A, R, c) has

- $N = \{1, \dots, n\}$ agents, indexed by i ,
- $R = \{1, \dots, q\}$ resources, indexed by r ,
- Action profiles $A = A_1 \times \dots \times A_n$
 - $A_i \subseteq 2^R \setminus \emptyset$: the action set of agent i
 - $a_i \in A_i$: the set of resources used, *i.e.*, $a_i \subseteq R$
- Cost function $c_r(\cdot) \in \mathbb{R}$ can depend on the number of agents that use the resource r
 - $c_i(a) = \sum_{r \in a_i} c_r(x_{r,a})$: the cost to agent i , given action profile a
 - $x_{r,a}$: the number of agents that select resource r , given action profile a

Example 3: Network Flow

- Agents: $n = 2000$
- Resources: $R = \{SA, SB, AB, AE, BE\}$
- Action set for agent i : $A_i = \{\{SA, AE\}, \{SB, BE\}, \{SA, AB, BE\}\}$
- Cost functions:
 - $c_{SA}(x) = c_{BE}(x) = x/2000$
 - $c_{SB}(x) = c_{AE}(x) = 1, c_{AB}(x) = 0$
- Eq. cost: $c_i(a^*) = c_{SA}(x) + c_{AB}(x) + c_{BE}(x) = \frac{2000}{2000} + 0 + \frac{2000}{2000} = 2$

Equilibrium profile



Congestion Games

Some notes on congestion games

- Symmetry: the payoff depends on the number of agents choosing an action, *not* which particular player(s)
- The cost / payoff representation scales linearly in the number of agents
- There always exists a pure-strategy Nash equilibrium of a congestion game
- Expressiveness: Not every game can be represented as a congestion game

Agent-Graph Games


An **agent-graph game** is defined via a graph $G = (V, E)$, which can be either directed or undirected

- V : the set of agents
- $e \in E$: the payoff dependence between connected agents
- Each agent has a set A_i of feasible actions
- u_i : utility as a function of the actions of the neighbors of agent i

Agent-Graph Games

Some notes on agent-graph games:

- Consider a game of n players, each with m actions; its agent-graph representation has a maximum degree of d
 - Each agent's utility can be represented with at most m^d numbers
 - Payoffs: normal-form $O(nm^n)$ vs. agent-graph $O(nm^d)$
 - The representation size is polynomial in the number of agents and actions, when the degree d is bounded by some constant
- Payoff dependence vs. strategic dependence
- Agent-graph games are fully expressive



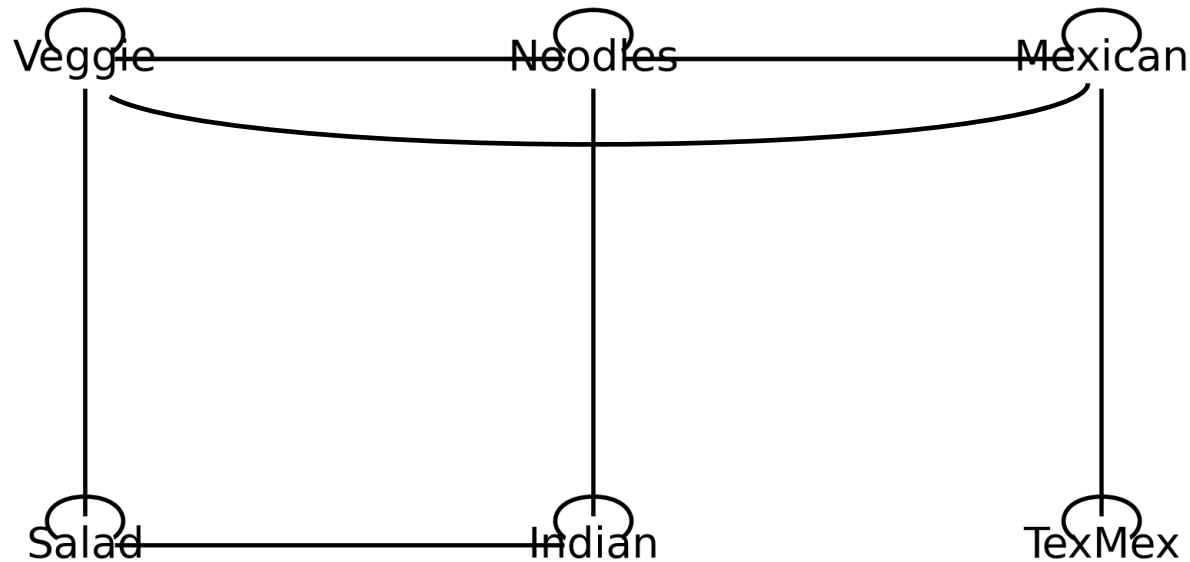
Action-Graph Games

An **action-graph game** is defined via a graph $G = (V, E)$, which can be either directed or undirected:

- $v \in V$: an **action** in the game
 - Each agent has a set $V_i \subseteq V$ of feasible actions ($A_i = V_i$)
- $e \in E$: the payoff dependence between agents who take the corresponding actions
- w_j : utility to any player who takes the action j , dependent on the number of agents who play neighboring actions to j

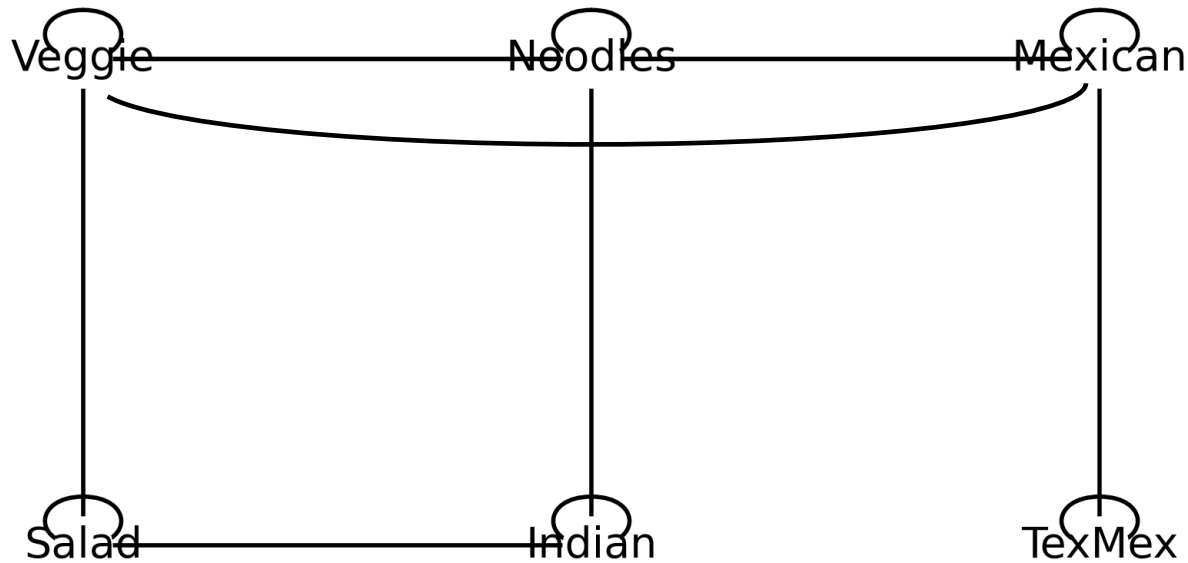
Example 4: Food Truck Games

- $n = 20$ sellers compete for business
- $m = 6$ different actions to choose from
- Symmetry: utility depends on #agents taking certain actions



Example 4: Food Truck Games

- $n = 20$ sellers compete for business
- $m = 6$ different actions to choose from
- Symmetry: utility depends on #agents taking certain actions




Payoffs: normal-form $(20)(6^{20}) \approx 10^{17}$ vs.
action-graph $(3)(20^4) + (2)(20^3) + (1)(20^2) = 496400$

Action-Graph Games

Some notes on action-graph games:

- Consider a game of n players, each with m actions; its action-graph representation has a maximum degree of d
 - Each **action utility** can be represented with at most n^d numbers
 - Payoffs: normal-form $O(nm^n)$ vs. symmetric action-graph $O(mn^d)$
 - Non-symmetric action-graph: mn vertices, payoffs $O(mn^{d+1})$
 - The representation size is polynomial in the number of agents and actions, when the degree d is bounded by some constant
- Payoff dependence vs. strategic dependence
- Action-graph games are fully expressive



Comparing Succinct Representations

Questions to think about:

- Show that the action-graph representation is exponentially more succinct than the agent-graph representation

Outline

- Simultaneous-move games
 - Normal-form representation
 - Solution concepts
 - Succinct representations
- Sequential-move games
 - Extensive-form representation
 - Solution concepts
 - Repeated games
 - Stackelberg games

Example 5: Bargaining Game

- Step 1: Player 1 determines a split of \$4
 - Choose from “me” (3, 1), “even” (2,2), and “you” (1,3)
- Step 2: Player 2 decides to accept or decline
 - Choose “Y” or “N” for each possible proposal

Example 5: Bargaining Game

- Step 1: Player 1 determines a split of \$4
- Step 2: Player 2 decides to accept or decline

Normal-form representation

	$\langle N, N, N \rangle$	$\langle N, N, Y \rangle$	$\langle N, Y, N \rangle$	$\langle N, Y, Y \rangle$	$\langle Y, N, N \rangle$	$\langle Y, N, Y \rangle$	$\langle Y, Y, N \rangle$	$\langle Y, Y, Y \rangle$
<i>me</i>	0, 0	0, 0	0, 0	0, 0	3, 1	3, 1	3, 1	3, 1
<i>even</i>	0, 0	0, 0	2, 2	2, 2	0, 0	0, 0	2, 2	2, 2
<i>you</i>	0, 0	1, 3	0, 0	1, 3	0, 0	1, 3	0, 0	1, 3

What are the Nash equilibria?

Example 5: Bargaining Game

- Step 1: Player 1 determines a split of \$4
- Step 2: Player 2 decides to accept or decline

Normal-form representation

	$\langle N, N, N \rangle$	$\langle N, N, Y \rangle$	$\langle N, Y, N \rangle$	$\langle N, Y, Y \rangle$	$\langle Y, N, N \rangle$	$\langle Y, N, Y \rangle$	$\langle Y, Y, N \rangle$	$\langle Y, Y, Y \rangle$
<i>me</i>	0, 0	0, 0	0, 0	0, 0	3, 1	3, 1	3, 1	3, 1
<i>even</i>	0, 0	0, 0	2, 2	2, 2	0, 0	0, 0	2, 2	2, 2
<i>you</i>	0, 0	1, 3	0, 0	1, 3	0, 0	1, 3	0, 0	1, 3

What are the Nash equilibria?

Example 5: Bargaining Game

- Step 1: Player 1 determines a split of \$4
- Step 2: Player 2 decides to accept or decline


Normal-form representation

	$\langle N, N, N \rangle$	$\langle N, N, Y \rangle$	$\langle N, Y, N \rangle$	$\langle N, Y, Y \rangle$	$\langle Y, N, N \rangle$	$\langle Y, N, Y \rangle$	$\langle Y, Y, N \rangle$	$\langle Y, Y, Y \rangle$
<i>me</i>	0, 0	0, 0	0, 0	0, 0	3, 1	3, 1	3, 1	3, 1
<i>even</i>	0, 0	0, 0	2, 2	2, 2	0, 0	0, 0	2, 2	2, 2
<i>you</i>	0, 0	1, 3	0, 0	1, 3	0, 0	1, 3	0, 0	1, 3

What are the Nash equilibria?

NE only requires that a strategy is a best response for the part of the game that can be “reached” in equilibrium

Sequential-Move Game



Outline

- Simultaneous-move games
 - Normal-form representation
 - Solution concepts
 - Succinct representations
- Sequential-move games
 - Extensive-form representation
 - Solution concepts
 - Repeated games
 - Stackelberg games

Extensive-Form Representation

The **extensive-form representation** of a **sequential-move game** Γ consists of the following:

- A set $N = \{1, \dots, n\}$ of agents or players, indexed by i
- A set of history H
 - **Terminal histories:** $h \in Z \subset H$, each as a leaf with a defined utility $u_i(h) \in \mathbb{R}$
E.g., $Z = \{(me, Y), (me, N), (even, Y), (even, N), (you, Y), (you, N)\}$
 - **Non-terminal histories:** $h \in H \setminus Z$, each as a decision node with a player $P(h) \in N$ and a set of feasible actions $A_i(h)$
E.g., $P(\epsilon) = 1, P(me) = P(you) = P(even) = 2$
 $A_1(\epsilon) = \{me, even, you\}$
 $A_2((me)) = A_2((you)) = A_2((even)) = \{Y, N\}$


Extensive-Form Game

- A **strategy** s_i of player i in an **extensive-form game** defines an action $s_i(h) \in A_i(h)$ for all non-terminal histories h when it is player i 's turn

E.g., $s_1(\epsilon) = \text{you}$; $s_2((me)) = N$, $s_2((you)) = Y$, $s_2((even)) = N$

Extensive-Form Game

- The **subgame** at history h , denoted Γ_h , of an extensive-form game Γ is the extensive-form game rooted at the decision node in Γ that corresponds to history h
- A strategy in the full game Γ defines a strategy in a subgame Γ_h



Outline

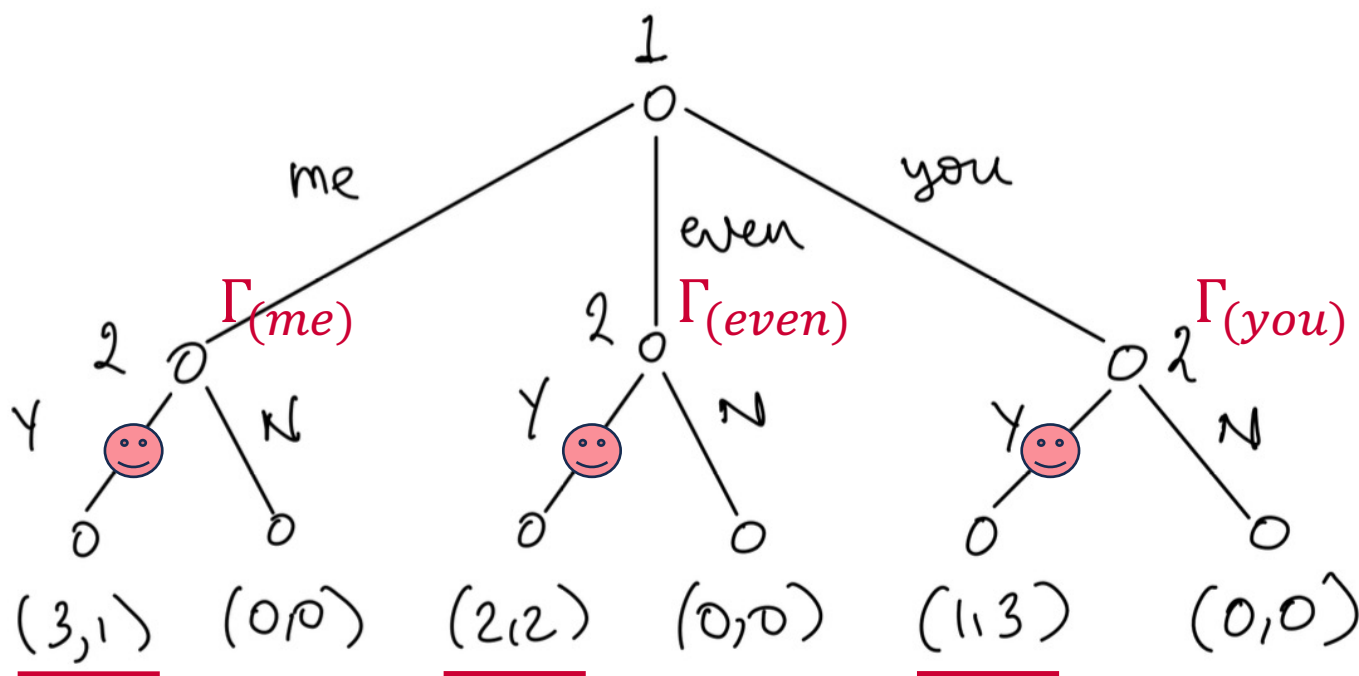
- Simultaneous-move games
 - Normal-form representation
 - Solution concepts
 - Succinct representations
- Sequential-move games
 - Extensive-form representation
 - Solution concepts
 - Repeated games
 - Stackelberg games

Subgame-Perfect Equilibrium (SPE)

- A strategy profile $s^* = (s_1^*, \dots, s_n^*)$ is a **subgame-perfect equilibrium** (SPE) of an extensive-form game Γ , *if* the strategy profile is a Nash equilibrium of game Γ and of subgame Γ_h **for every non-terminal history h**
- Best responses in every subgame, *not just* the subgames that are reached on the equilibrium path

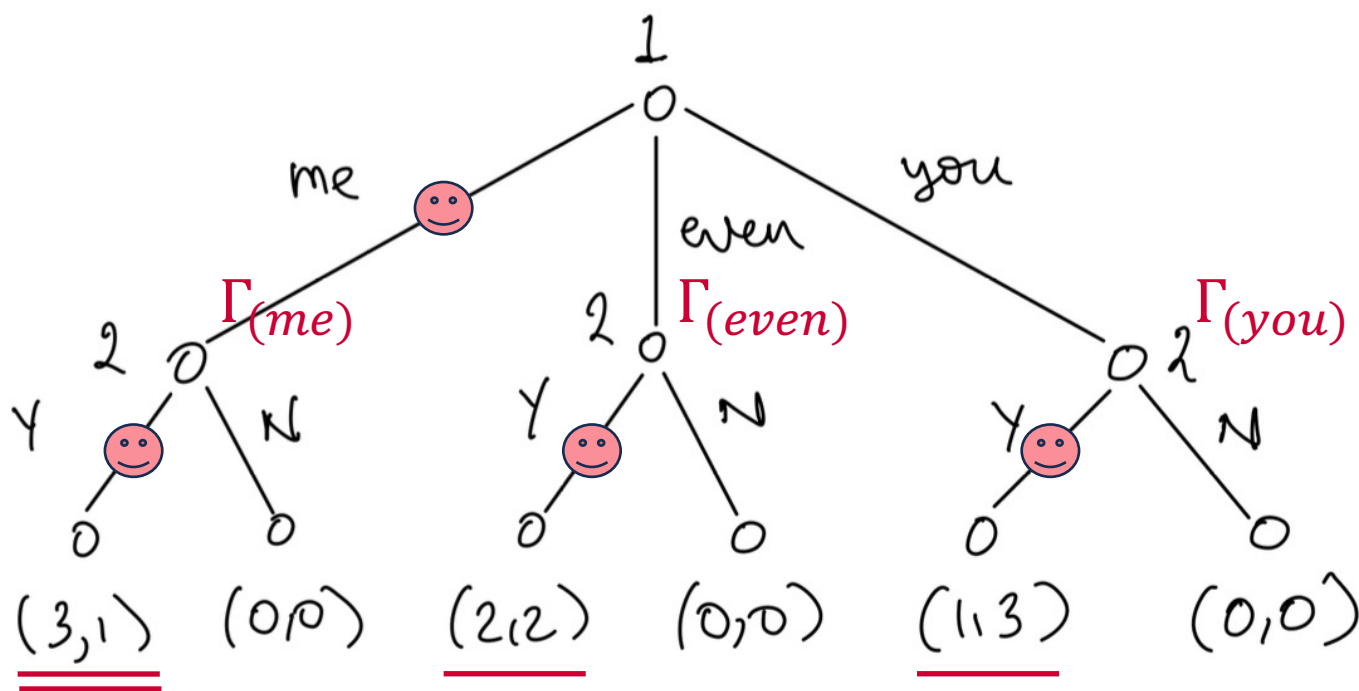
Finding Subgame-Perfect Equilibrium

The backward induction procedure



Finding Subgame-Perfect Equilibrium

The backward induction procedure




SPE: (me, <Y, Y, Y>)

Finding Subgame-Perfect Equilibrium

The backward induction procedure

Take time linear in the number of nodes in tree



SPE: (me, <Y, Y, Y>)

Checking Subgame-Perfect Equilibrium

Single-deviation principle

- A **single deviation** from strategy s_i at history h is a strategy s'_i that differs **only in the action played at history h**
- A single deviation is useful if

$$u_i(s'_i, s_{-i} \mid \Gamma_h) > u_i(s_i, s_{-i} \mid \Gamma_h)$$

Checking Subgame-Perfect Equilibrium

Single-deviation principle

Theorem: A strategy profile s^* is a SPE of a finite extensive-form game *if and only if* there's no useful *single* deviation for any player

Proof:

(If) Basic idea: proof by contradiction. If a more complicated, multi-step deviation is useful, then a simpler deviation will be as well

(Only if) SPE \rightarrow NE in subgames \rightarrow no useful single deviation

Checking Subgame-Perfect Equilibrium

Theorem: Any finite extensive-form game has a SPE

Proof:

- (1) Use backward induction to find the strategy profile(s)
- (2) The found strategy satisfies the single-deviation principle

When do we have a unique SPE?

Outline

- Simultaneous-move games
 - Normal-form representation
 - Solution concepts
 - Succinct representations
- Sequential-move games
 - Extensive-form representation
 - Solution concepts
 - Repeated games
 - Stackelberg games

Repeated Games

- A class of sequential-move games
- In a **finitely-repeated game** G^T , the same simultaneous-move game $G = (N, \tilde{A}, \tilde{u})$ (i.e., the **stage game**) is played by the same players for $T \geq 1$ periods
 - Perfect information about the history of actions
 - G^∞ : **infinitely-repeated games**, the stage game G is repeated forever
- E.g., same players play a Prisoners' Dilemma for 8 times
same players play rock-paper-scissors

Finitely-Repeated Games

- A strategy s_i in a finitely-repeated game defines an action after every history
- Total utility at a terminal history: $u_i(h) = \sum_{k=0}^{T-1} \tilde{u}_i(a^{(k)})$

Finitely-Repeated Games

Single-deviation principle holds for finitely-repeated games

- Theorem: A strategy profile s^* is an SPE of a finitely-repeated game G^T if and only if there is no useful single deviation

Finitely-Repeated Games

- Theorem (Unique SPE): If the stage game G has a *unique* Nash equilibrium, then the **only** SPE s^* of the finitely-repeated game G^T is to play the Nash equilibrium of the stage game after every history

Proof:

(1) SPE: a deviation from NE at any h is not useful

$$u_i(s'_i, s^*_{-i} | h) = w'_i + \sum_{k'=k+1}^{T-1} w_i \leq w_i + \sum_{k'=k+1}^{T-1} w_i = u_i(s^*_i, s^*_{-i} | h)$$

(2) Uniqueness: backward induction + unique NE

- E.g., playing Prisoners' Dilemma or R-P-S multiple times

Infinitely-Repeated Games

- Total discounted utility:


$$u_i(h) = \sum_{k=0}^{\infty} \delta^k \cdot \tilde{u}_i(a^{(k)})$$

- $0 < \delta < 1$ is a discount factor, s.t. $u_i(h)$ is bounded if $\tilde{u}_i(a^{(k)})$ is bounded for all k
- Single-deviation principle holds for infinitely-repeated games *with discounting*

Infinitely-Repeated Games

- An **open-loop strategy** s_i for player i in a repeated game has $s_i(h) = s_i(h')$ for any history h and h' of the same length
- Not dependent on the play in previous periods
- E.g., always “Go”; “Go” or “Wait” with prob=0.5; Cycle through “Go”, “Go”, “Wait”

		Player 2	
		W	G
Player 1	W	0, 0	0, 2
	G	2, 0	-4, -4



Infinitely-Repeated Games

- Theorem: An **open-loop, stage-Nash strategy profile** s^* is a SPE of a repeated game, either finite or infinite

Proof:

A single deviation from stage-NE at any h is not useful

$$\begin{aligned} u_i(s'_i, s^*_{-i} | h) &= w'_i + \delta \cdot u_i(s'_i, s^*_{-i} | h, a') = w'_i + \delta \cdot u_i(s^*_i, s^*_{-i} | h, a') \\ \text{open-loop, independent of previous play} &= w'_i + \delta \cdot u_i(s^*_i, s^*_{-i} | h, a) \\ &\leq w_i + \delta \cdot u_i(s^*_i, s^*_{-i} | h, a) = u_i(s^* | h) \end{aligned}$$

- E.g., the cyclic play (W, G), (G, W), (W, G), (G, W)

Outline

- Simultaneous-move games
 - Normal-form representation
 - Solution concepts
 - Succinct representations
- Sequential-move games
 - Extensive-form representation
 - Solution concepts
 - Repeated games
 - Stackelberg games

Stackelberg Games

- One player (the “leader”) moves first, and the other player (the “follower”) moves after
- Can be generalized to multiple leaders/followers
- Applications
 - Public policy: a policymaker and other participants
 - Security domain: a defender and an attacker
 - Online marketplace: the marketplace and buyers/sellers

Stackelberg Equilibrium

- A two-player game: a leader l and a follower f , with corresponding sets of actions A_l and A_f . $A = A_l \times A_f$
- Strategies: $x \in \Delta(A_l)$ and $y \in \Delta(A_f)$
- Utility for a player $i \in \{l, f\}$:

$$u_i(x, y) = \mathbb{E}_{a_l \sim x, a_f \sim y}[u_i(a_l, a_f)]$$

- The leader knows *ex ante* that the follower observes its action

Stackelberg Equilibrium

- Given any leader strategy x , the **follower** chooses their strategy from the *best-response set* to strategy x

$$BR(x) = \operatorname{argmax}_{y \in \Delta(A_f)} u_f(x, y)$$

- Based on the best response assumption, the **leader** chooses their strategy x

$$\max_{x \in \Delta(A_l)} u_l(x, y) \quad \text{s.t. } y \in BR(x)$$

Stackelberg Equilibrium

- Given any leader strategy x , the **follower** chooses their strategy from the *best-response set* to strategy x

$$BR(x) = \operatorname{argmax}_{y \in \Delta(A_f)} u_f(x, y)$$

- Based on the best response assumption, the **leader** chooses their strategy x

$$\max_{x \in \Delta(A_l)} u_l(x, y) \quad \text{s.t. } y \in BR(x)$$

- *Which $y \in BR(x)$ will the follower choose?*

Stackelberg Equilibrium

- **Strong Stackelberg equilibrium (SSE):** the follower breaks ties in favor of the leader

$$\max_{x \in \Delta(A_l), y \in BR(x)} u_l(x, y)$$

- **Weak Stackelberg equilibrium (WSE):** the follower breaks ties adversarially to the leader

$$\max_{x \in \Delta(A_l)} \min_{y \in BR(x)} u_l(x, y)$$

Stackelberg Equilibrium

- **Strong Stackelberg equilibrium (SSE):** the follower breaks ties in favor of the leader

$$\max_{x \in \Delta(A_l), y \in BR(x)} u_l(x, y)$$

- **Weak Stackelberg equilibrium (WSE):** the follower breaks ties adversarially to the leader

$$\max_{x \in \Delta(A_l)} \min_{y \in BR(x)} u_l(x, y)$$

- Comparing to playing NE, will the leader benefit from firstly committing to a strategy?

Stackelberg Equilibrium

- Commit to pure actions $a_l \in A_l$?
- Commit to any $x \in \Delta(A_l)$?
- Theorem: In a general-sum game, the leader achieves weakly more utility in SSE than in *any* Nash equilibrium

Proof: Consider the NE (x, y) that yields the highest utility for the leader

Stackelberg Equilibrium

- Commit to pure actions $a_l \in A_l$?
- Commit to any $x \in \Delta(A_l)$?
- Theorem: In a general-sum game, the leader achieves weakly more utility in SSE than in *any* Nash equilibrium
Proof: Consider the NE (x, y) that yields the highest utility for the leader
- Theorem: In a general-sum game, the WSE provides the leader a utility at least as good as *some* Nash equilibrium

Logistics (Reminder)

- Pre-class CQs
 - Due before each lecture
 - Binary grading scheme
 - Two chances to drop
- Paper presentations
 - Bidding on papers
 - Presentation guidelines
- Class survey for newly registered
- Office hour
 - After class till 2pm today
 - 2-3pm for future weeks