CS 598: AI Methods for Market Design

Lecture 3: Computing Equilibrium

Xintong Wang

Spring 2024

Announcements

- Paper presentation assignment is out
  - Paper reading and presenting guidelines
  - Peer evaluation and grading scheme
- Office hour for today: 1:20pm—2pm
- CS Colloquium: 2pm–3pm at CoRE 301
  - “Eliciting Information without Verification from Humans and Machines” by Yuqing Kong

Recap

- Simultaneous-move games
    - Normal-form representation
    - Solution concepts
    - Succinct representations

- Sequential-move games
    - Extensive-form representation
    - Solution concepts
    - Repeated games
    - Stackelberg games

Recap

- Simultaneous-move games
  - Normal-form representation
  - Solution concepts
  - Succinct representations

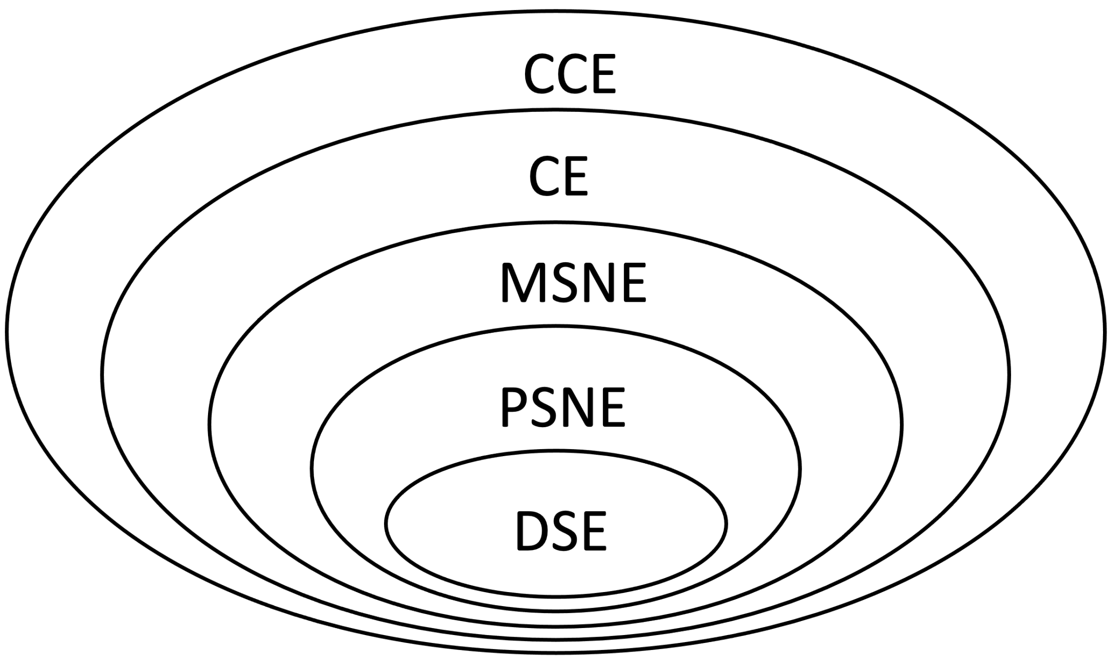

Recap

- Sequential-move games
  - Extensive-form representation
  - Solution concepts
  - Repeated games
  - Stackelberg games

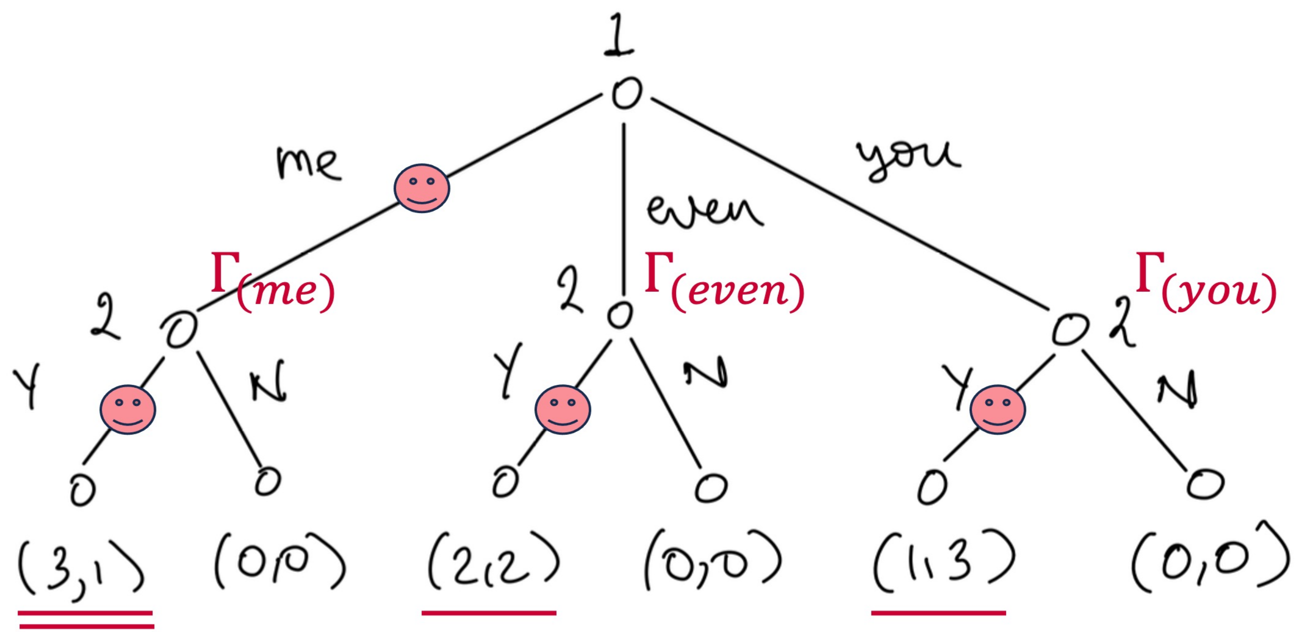

Outline

- Simultaneous-move games
  - Normal-form representation
  - Solution concepts
  - Succinct representations

- Sequential-move games
  - Extensive-form representation
  - Solution concepts
  - Repeated games
  - Stackelberg games

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: $u_i(h) = \sum_{k=0}^{T-1} \widetilde{u}_i(a^{(k)})$

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

$$u_i(s_i', s_{-i}^* \mid h) = w_i' + \sum_{k'=k+1}^{T-1} w_i \leq w_i + \sum_{k'=k+1}^{T-1} w_i = u_i(s_i^*, s_{-i}^* \mid h)$$

(2) Uniqueness: backward induction + unique NE

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

Infinitely-Repeated Games

- Total discounted utility:

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

- $0 < \delta < 1$ is a discount factor, s.t. $u_i(h)$ is bounded if $\widetilde{u_i}\left(a^{(k)}\right)$ 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

$$u_i(s_i', s_{-i}^* \mid h) = w_i' + \delta \cdot u_i(s_i', s_{-i}^* \mid h, a') = w_i' + \delta \cdot u_i(s_i^*, s_{-i}^* \mid h, a')$$

open-loop, independent of previous play

$$= w_i' + \delta \cdot u_i(s_i^*, s_{-i}^* \mid h, a)$$

$$\leq w_i + \delta \cdot u_i(s_i^*, s_{-i}^* \mid h, a) = u_i(s^* \mid h)$$

- 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. $\mathrm{A} = A_l \times \mathrm{A}_f$
- Strategies: $x \in \Delta(A_l)$ and $\mathrm{y} \in \Delta(A_f)$
- Utility for a player $i \in \{l, f\}$:

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

- 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

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

- 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

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

- 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

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

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

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

- 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

Today: Equilibrium Computation

- NE in a two-player, zero-sum game
- PSNE in a general-sum game
- MSNE in a two-player, general-sum game
- MSNE in a general-sum game
- CE in a general-sum game
- SSE in Stackelberg game (one leader, one follower)

Today: Equilibrium Computation

- NE in a two-player, zero-sum game
- PSNE in a general-sum game
- MSNE in a two-player, general-sum game
- MSNE in a general-sum game
- CE in a general-sum game
- SSE in Stackelberg game (one leader, one follower)

Recap: Two-Player, Zero-Sum Game

- Matching Pennies game

		Player 2: H	Player 2: T
Player 1	H	$1, -1$	$-1, 1$
	T	$-1, 1$	$1, -1$

Recap: Two-Player, Zero-Sum Game

- 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)

Example 1: Odd-or-Even Game

- Each player chooses to play \$1 or \$2

Player 2

		1D	*2D*
Player 1	*1D*	$-2,\ 2$	$3,\ -3$
	2D	$3,\ -3$	$-4,\ 4$

What is the MSNE?

Maximin Strategy

- Player 1 selects a strategy to maximize its expected utility, given that player 2 knows the goal and selects an action to minimize player 1's expected utility

- A maximin strategy for player 1 in a two-player game

$$\bar{x} \in \text{argmax}_{x \in \Delta(A_1)} [\min_{a_2 \in A_2} u_1(x, a_2)]$$

- Maximin value for player 1

$$\overline{v_1} = \min_{a_2 \in A_2} u_1(\bar{x}, a_2)$$

Minimax Strategy

- Player 1 selects a strategy to minimize player 2's expected utility, given that player 2 knows the goal and selects an action to maximize its expected utility

- A minimax strategy for player 1 in a two-player game

$$\underline{x} \in \mathrm{argmin}_{x \in \Delta(A_1)} [\max_{a_2 \in A_2} u_2(x, a_2)]$$

- Minimax value for player 2

$$\underline{v_2} = \max_{a_2 \in A_2} u_2(\underline{x}, a_2)$$

Example 1: Odd-or-Even Game

		Player 2	
		1D	*2D*
Player 1	*1D*	$-2,\ 2$	$3,\ -3$
	2D	$3,\ -3$	$-4,\ 4$

What is the maximin strategy for player 1?

Example 1: Odd-or-Even Game

		Player 2: $1D$	Player 2: $2D$
Player 1	$1D$	$-2,\ 2$	$3,\ -3$
	$2D$	$3,\ -3$	$-4,\ 4$

What is the maximin strategy for player 1?

x: player 1's probability of choosing 1D

Player 1 will choose the x that maximizes

$$\min(u_1(x, 1D), u_1(x, 2D)) = \min(-2x + 3(1-x), 3x - 4(1-x))$$

Example 1: Odd-or-Even Game

		Player 2	
		1D	*2D*
Player 1	*1D*	$-2,\ 2$	$3,\ -3$
	2D	$3,\ -3$	$-4,\ 4$

Exercise: What is the maximin strategy for player 2?

The Minimax Theorem

Theorem 3.4 (Minimax). In any two-player, zero-sum game,

(1) For each player, the set of maximin strategies is equal to the set of minimax strategies

Proof:

The Minimax Theorem

Theorem 3.4 (Minimax). In any two-player, zero-sum game,

(1) For each player, the set of maximin strategies is equal to the set of minimax strategies

(2) Each player’s maximin value is equal to its minimax value, and equal to its expected utility in any Nash equilibrium

Proof:

The Minimax Theorem

Theorem 3.4 (Minimax). In any two-player, zero-sum game,

(1) For each player, the set of maximin strategies is equal to the set of minimax strategies

(2) Each player's maximin value is equal to its minimax value, and equal to its expected utility in any Nash equilibrium

(3) Any maximin or minimax strategy for player 1 and any maximin or minimax strategy for player 2 form a Nash equilibrium, and these correspond to all Nash equilibria

Proof:

Solving for the Maximin Strategy

For player 1

$$\max_{v_1,\ x} \quad v_1$$

$$\text{s.t.} \quad \sum_{j \in A_1} u_1(j,k) \cdot x_j \geq v_1, \qquad \forall k \in A_2$$

$$\sum_{j \in A_1} x_j = 1$$

$$x_j \geq 0, \qquad \forall j \in A_1$$

the prob. of playing action *j*

Solving for the Maximin Strategy

Odd-or-Even Game

		Player 2	
		1D	*2D*
Player 1	*1D*	$-2,\ 2$	$3,\ -3$
	2D	$3,\ -3$	$-4,\ 4$

What is the maximin strategy for player 1?

Linear Programming

Given n variables and m constraints, variables x and constants w a and b:

$$\begin{aligned} \text{maximize} \quad & \sum_{i=1}^{n} w_i x_i & \\ \text{subject to} \quad & \sum_{i=1}^{n} a_{ij} x_i \leq b_j & \forall j = 1 \ldots m \\ & x_i \in \{0, 1\} & \forall i = 1 \ldots n \end{aligned}$$

LPs can be solved in polynomial time using interior point methods. In practice, the (worst-case exponential) simplex method is often faster.

Solving for the Maximin Strategy

Theorem. *FindNash* in a two-player, zero-sum, normal-form game can be solved in worst-case polynomial time in the size of the payoff matrix

Proof:

- The LP for finding a maximin strategy has $1 + |A_1|$ variables, and $|A_2| + 1 + |A_1|$ constraints
- By the Minimax theorem, the maximin strategies for each player provide a Nash equilibrium

Stackelberg Equilibrium

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

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

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

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

In two-player, zero-sum game, Nash equilibrium and Stackelberg equilibrium are equivalent!

Today: Equilibrium Computation

- NE in a two-player, zero-sum game
- PSNE in a general-sum game
- MSNE in a two-player, general-sum game
- MSNE in a general-sum game
- CE in a general-sum game
- SSE in Stackelberg game (one leader, one follower)

Computing PSNE in a General-Sum Game

Subroutines:

- $\text{IsNash}((a_1, \dots, a_n), G)$: check whether an action profile is a Nash equilibrium — $n(m-1)$ single deviations
- $\text{next}(a, G)$: return the next action profile

```
begin
  NashFound := false
  initialize (a_1, ..., a_n) ∈ A_1 × ... × A_n          m^n action profiles
  while ¬NashFound do
    (NashFound, NashEq) := IsNash((a_1, ..., a_n), G)
    (a_1, ..., a_n) := next(a_1, ..., a_n)
  Output: (NashFound, NashEq)
```

The size of payoff matrix: $O(nm^n)$; The runtime of computing PSNE: $O(nm^{n+1})$

Pre-Processing: Iterated Elimination

An action $a_i \in A_i$ is strictly dominated if there exists a mixed strategy x that places no probability on a_i s.t.

$$u_i(x, a_{-i}) > u_i(a_i, a_{-i}). \quad \forall a_{-i} \in A_{-i}$$

Question: Can we use an LP to determine whether a_i is strictly dominated by some mixed strategy?

Pre-Processing: Iterated Elimination

An action $a_i \in A_i$ is strictly dominated if there exists a mixed strategy x that places no probability on a_i s.t.

$$u_i(x, a_{-i}) > u_i(a_i, a_{-i}). \quad \forall a_{-i} \in A_{-i}$$

minimize $\sum x_j$ (the prob. of playing action *j*)

subject to $\sum_{j \in A_i} u_i(j, a_{-i})\, x_j \geq u_i(a_i, a_{-i}) \quad \forall a_{-i} \in A_{-i}$

$x_j \geq 0 \quad \forall j \in A_i$

Check $\sum x_j < 1$ to see whether a_i is strictly dominated

Pre-Processing: Iterated Elimination

Runtime analysis

- For each game, there can be at most $n(m-1)$ stages
- In each stage, it needs to check at most m actions per agent, i.e., we run mn LPs
- Note: solving a polynomial number of LPs is still in $\mathcal{P}$

Today: Equilibrium Computation

- NE in a two-player, zero-sum game
- PSNE in a general-sum game
- MSNE in a two-player, general-sum game
- MSNE in a general-sum game
- CE in a general-sum game
- SSE in Stackelberg game (one leader, one follower)

Computing MSNE in a Two-Player, General-Sum Game

Subroutines:

- $\mathrm{next}(X, Y)$: perform the support enumeration by returning the next pair of action sets, $(X, Y) \subseteq A_1 \times A_2$

  E.g., ({R},{R}), ({R},{P}), …,({R},{P,S}),…,({R,P},{P,S})…,({R,P,S},{R,P,S})

  $O(2^{2m})$ action set pairs!

- $\mathrm{CheckNash}((X, Y), G)$: look for a Nash equilibrium (x, y) that have support $\sigma(x) \subseteq X, \sigma(y) \subseteq Y$ and satisfy

  (P1) Player 1 is indifferent across every action in X, given strategy y and weakly prefers any action in X to any other action

  (P2) Player 2 is indifferent across every action in Y, given strategy x and weakly prefers any action in Y to any other action

  Polynomial time

Computing MSNE in a Two-Player, General-Sum Game

Subroutines:

- $\text{next}(X, Y)$: perform the support enumeration by returning the next pair of action sets, $(X, Y) \subseteq A_1 \times A_2$
- $\text{CheckNash}((X, Y), G)$: look for a Nash equilibrium (x, y) that have support $\sigma(x) \subseteq X, \sigma(y) \subseteq Y$

```
begin
    NashFound := false
    initialize (X, Y) ⊆ A_1 × A_2
    while ¬NashFound do
        (NashFound, NashEq) := CheckNash((X, Y), G)
        (X, Y) := next(X, Y)
    Output: (NashFound, NashEq)
```

Traffic Light Game

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

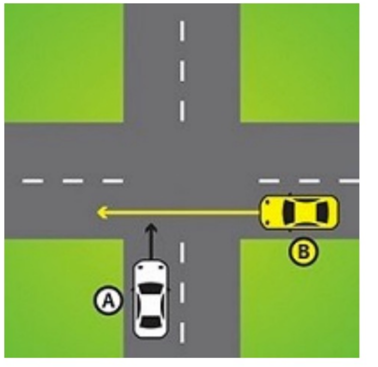

Two pure-strategy NE: (G, W) and (W, G)

Support of 1

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

Support of 2

Traffic Light Game (Variation)

		Player 2 W	Player 2 G	Player 2 C
Player 1	W	0, 0	0, 2	0, 1/4
	G	2, 0	−4, −4	−4, −1
	C	1/4, 0	−1, −4	1/2, 1/2

X = {W, G}, Y = {W, G}

CheckNash returns MSNE (2/3, 1/3, 0) for both players

W and G are better than C: $u_1(C, y) = u_2(x, C) = -1/6$

Computing MSNE in a Two-Player, General-Sum Game

Theorem. (1) CheckNash is guaranteed to return a NE when the input, (X, Y), corresponds to the support of a NE

(2)The CheckNash problem can be solved in polynomial time

Proof for (1):

(P1) and (P2) guarantees a Nash equilibrium.

when (X, Y) corresponds to the support of a Nash equilibrium, there is at least one strategy profile that satisfies (P1) and (P2)

Computing MSNE in a Two-Player, General-Sum Game

Subroutines:

- $\mathrm{next}(X, Y)$: perform the support enumeration by returning the next pair of action sets, $(X, Y) \subseteq A_1 \times A_2$

  E.g., ({R},{R}), ({R},{P}), …,({R},{P,S}),…,({R,P},{P,S})…,({R,P,S},{R,P,S})

  $O(2^{2m})$ action set pairs!

- $\mathrm{CheckNash}((X, Y), G)$: look for a Nash equilibrium (x, y) that have support $\sigma(x) \subseteq X, \sigma(y) \subseteq Y$ and satisfy

  (P1) Player 1 is indifferent across every action in X, given strategy y and weakly prefers any action in X to any other action

  (P2) Player 2 is indifferent across every action in Y, given strategy x and weakly prefers any action in Y to any other action

  Polynomial time

Computing MSNE in a Two-Player, General-Sum Game

Theorem. (1) CheckNash is guaranteed to return a NE when the input, (X, Y), corresponds to the support of a NE

(2)The CheckNash problem can be solved in polynomial time

Proof for (2): a linear feasibility program following (P1) and (P2)

$$\sum_{j \in A_1} u_2(j,k) \cdot x_j = v_2, \; \forall k \in Y, \quad \sum_{j \in A_1} u_2(j,k) \cdot x_j \leq v_2, \; \forall k \in A_2 \setminus Y$$

$$\sum_{j \in A_1} x_j = 1, \; x_j \geq 0, \; \forall j \in X, \quad x_j = 0, \; \forall j \in A_1 \setminus X$$

$$\sum_{k \in A_2} u_1(j,k) \cdot y_k = v_1, \; \forall j \in X, \quad \sum_{k \in A_2} u_1(j,k) \cdot y_k \leq v_1, \; \forall j \in A_1 \setminus X$$

$$\sum_{k \in A_2} y_k = 1, \; y_k \geq 0, \; \forall k \in Y, \quad y_k = 0, \; \forall k \in A_2 \setminus Y$$

$|A_1| + |A_2| + 2$ variables; $|A_2| + (1 + |A_1|) + |A_1| + (1 + |A_2|)$ constraints

Computing MSNE in a Two-Player, General-Sum Game

Exercise. For X={M, D}, Y={L, R}, solve CheckNash for the game

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

Today: Equilibrium Computation

- NE in a two-player, zero-sum game
- PSNE in a general-sum game
- MSNE in a two-player, general-sum game
- MSNE in a general-sum game
- CE in a general-sum game
- SSE in Stackelberg game (one leader, one follower)

Computing MSNE in a Multi-Player, General-Sum Game

- Extend support enumeration method to multi-player, general-sum game
- The number of support tuples: $(2^m - 1)^n$
- $\text{CheckNash}((X, Y, Z), G)$ for three-player game
  - Player 1 is indifferent across all actions in X

$$\sum_{k \in A_2} \sum_{l \in A_3} u_1(j, k, l) \cdot y_k z_l = v_1 \;\; \forall j \in X$$

Nonlinear feasible problem!

- Suitable for solving multi-player games that have Nash equilibria with small supports

Today: Equilibrium Computation

- NE in a two-player, zero-sum game
- PSNE in a general-sum game
- MSNE in a two-player, general-sum game
- MSNE in a general-sum game
- CE in a general-sum game
- SSE in Stackelberg game (one leader, one follower)

Recap: Correlated Equilibrium (CE)

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

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

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

$$\sum_{a_{-i} \in A_{-i}} u_i(j, a_{-i}) \cdot p^*_{-i}(a_{-i} \mid j) \geq \sum_{a_{-i} \in A_{-i}} u_i(j', a_{-i}) \cdot p^*_{-i}(a_{-i} \mid j),$$

$$\forall i \in n, j \in \sigma(p^*_i), j' \in \mathrm{A}_i$$

Recap: Correlated Equilibrium (CE)

A probability distribution p^* on action profiles A is a correlated equilibrium *if and only if*

$$\sum_{a_{-i} \in A_{-i}} u_i(j, a_{-i}) \cdot p^*_{-i}(a_{-i} \mid j) \cdot p^*_i(j) \geq \sum_{a_{-i} \in A_{-i}} u_i(j', a_{-i}) \cdot p^*_{-i}(a_{-i} \mid j) \cdot p^*_i(j),$$

$$\forall i \in n, j \in \mathrm{A}_i, j' \in \mathrm{A}_i$$

$$\sum_{a_{-i} \in A_{-i}} u_i(j, a_{-i}) \cdot p^*(j, a_{-i}) \geq \sum_{a_{-i} \in A_{-i}} u_i(j', a_{-i}) \cdot p^*(j, a_{-i}),$$

$$\forall i \in n, j \in \mathrm{A}_i, j' \in \mathrm{A}_i$$

Computing CE in a Multi-Player, General-Sum Game

- A linear feasibility program to find CE

$$\sum_{a_{-i}\in A_{-i}} u_i(j,a_{-i})\cdot p(j,a_{-i}) \geq \sum_{a_{-i}\in A_{-i}} u_i(j',a_{-i})\cdot p(j,a_{-i}), \quad \forall i, \quad \forall j,j'\in A_i$$

$$\sum_{a\in A} p(a) = 1, \quad p(a)\geq 0, \quad \forall a\in A$$

Computing CE in a Two-Player, General-Sum Game

- A linear feasibility program to find CE

$|A_1|(|A_1|-1)$ constraints

$$\sum_{k \in A_2} u_1(j,k) \cdot p(j,k) \geq \sum_{k \in A_2} u_1(j',k) \cdot p(j,k), \quad \forall j \in A_1,\ \forall j' \in A_1,$$

$|A_2|(|A_2|-1)$ constraints

$$\sum_{j \in A_1} u_2(j,k) \cdot p(j,k) \geq \sum_{j \in A_1} u_2(j,k') \cdot p(j,k), \quad \forall k \in A_2,\ \forall k' \in A_2,$$

$$\sum_{j \in A_1, k \in A_2} p(j,k) = 1, \quad p(j,k) \geq 0, \quad \forall j \in A_1,\ \forall k \in A_2.$$

$|A_1||A_2|+1$ constraints

$|A_1||A_2|$ variables

Computing CE in a Multi-Player, General-Sum Game

Theorem. A correlated equilibrium of a multi-player, general-sum, normal-form game can be computed in polynomial time in the size of the payoff matrix.

$O(nm^n)$ entries

Computing CE in a Multi-Player, General-Sum Game

Theorem. A correlated equilibrium of a multi-player, general-sum, normal-form game can be computed in polynomial time in the size of the payoff matrix.

$O(nm^n)$ entries

Proof:

For a game with m actions per agent and n agents, there are m^n variables.

For each agent to follow recommended action, there are $m(m-1)$ constraints; for n agents, there are $m(m-1)n$ constraints.

There are $1 + m^n$ constraints to guarantee a valid probability distribution.

Traffic Light Game

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

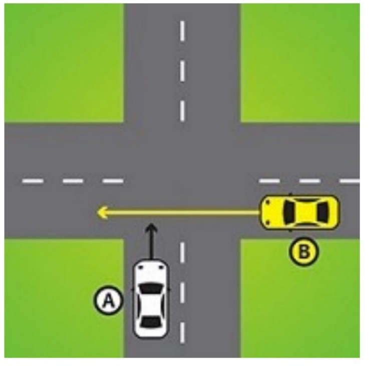

Exercise: What are the correlated equilibrium?

Equilibrium Hierarchy for Simultaneous-Move Games

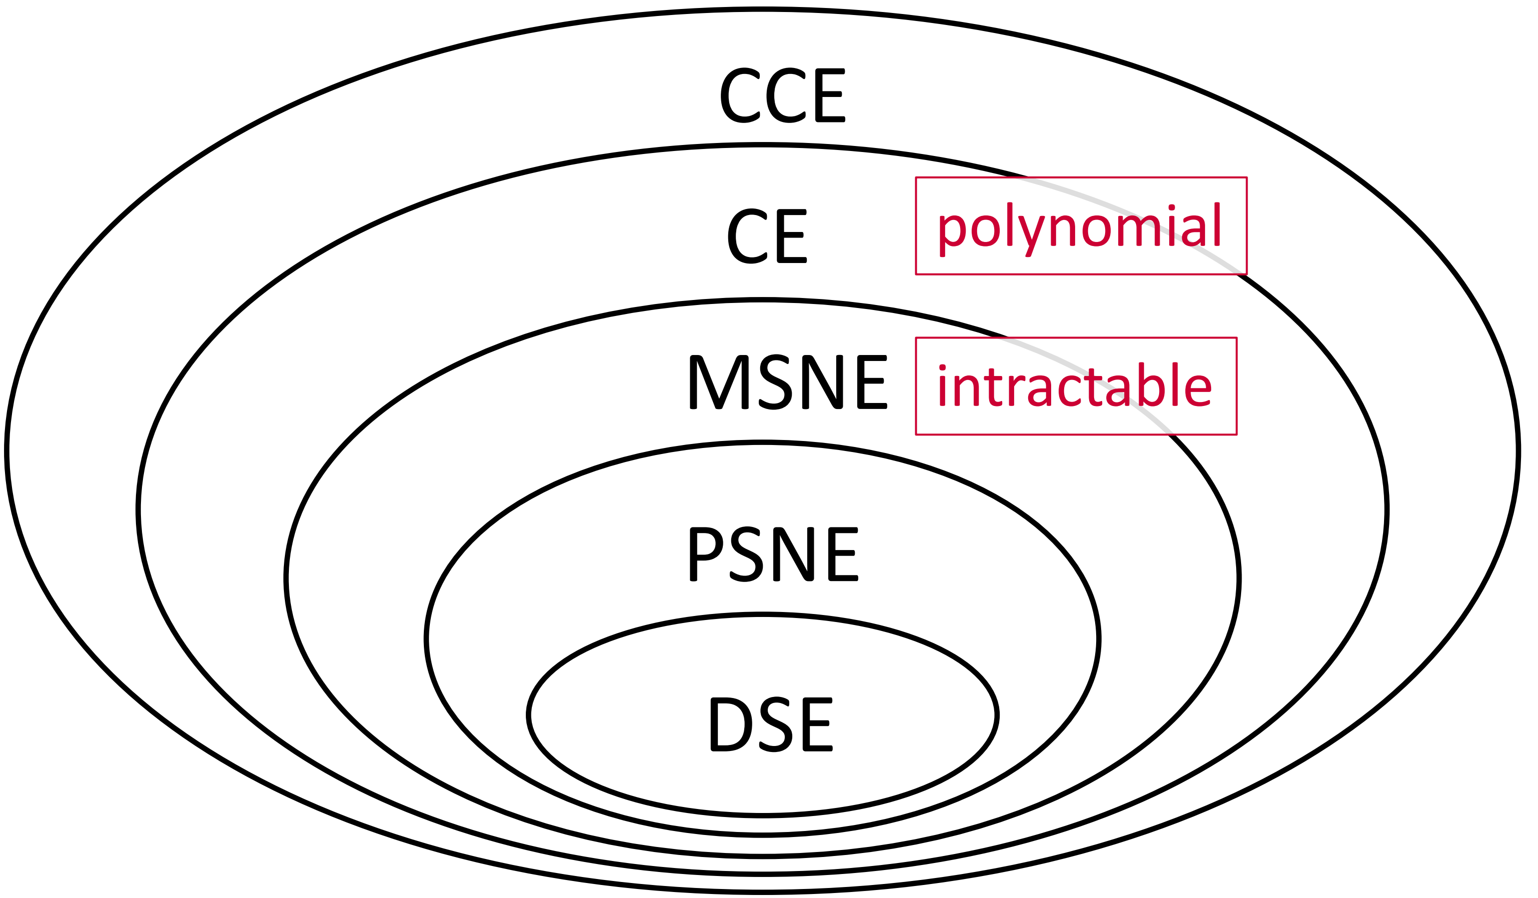

Computing CE, NE under Succinct Representations

- Normal-form representation is exponential in #players

- Theorem. A correlated equilibrium can be computed in polynomial time in the size of the congestion-game, agent-graph, and action-graph representations of simultaneous-move games

- Succinct game representations also grant faster computation of the expected utility of a mixed strategy, thus also of NE

Similar Approaches to FindNash?

For a Nash equilibrium, we have

$$\sum_{a_{-i} \in A_{-i}} u_i(j, a_{-i}) \cdot p^*_{-i}(a_{-i}) \geq \sum_{a_{-i} \in A_{-i}} u_i(j', a_{-i}) \cdot p^*_{-i}(a_{-i})$$

$$\forall i \in [n], j \in \sigma(s_i^*), j' \in A_i$$

$$\sum_{a_{-i} \in A_{-i}} u_i(j, a_{-i}) \cdot p^*_{-i}(a_{-i}) \cdot s^*_{ij} \geq \sum_{a_{-i} \in A_{-i}} u_i(j', a_{-i}) \cdot p^*_{-i}(a_{-i}) \cdot s^*_{ij}$$

$$\forall i \in [n], j, j' \in A_i$$

s^*_{ij}: the probability of agent i playing strategy j in s^*_i

Similar Approaches to FindNash?

For a two-player game, we have

$$\sum_{k\in A_2} u_1(j,k)\cdot \boxed{x_j\cdot y_k} \geq \sum_{k\in A_2} u_1(j',k)\cdot x_j\cdot y_k, \quad \forall j\in A_1,\ \forall j'\in A_1$$

$$\sum_{j\in A_1} u_2(j,k)\cdot x_j\cdot y_k \geq \sum_{j\in A_1} u_2(j,k')\cdot x_j\cdot y_k, \quad \forall k\in A_2,\ \forall k'\in A_2$$

$$\sum_{j\in A_1} x_j = 1,\ \sum_{k\in A_2} y_k = 1,\quad x_j, y_k \geq 0,\ \ \forall j\in A_1,\ \forall k\in A_2.$$

A nonlinear feasibility program to find NE

Computing CE in a Two-Player, General-Sum Game

A linear feasibility program to find CE

$$\sum_{k \in A_2} u_1(j,k) \cdot \boxed{p(j,k)} \geq \sum_{k \in A_2} u_1(j',k) \cdot p(j,k), \quad \forall j \in A_1,\ \forall j' \in A_1,$$

$$\sum_{j \in A_1} u_2(j,k) \cdot p(j,k) \geq \sum_{j \in A_1} u_2(j,k') \cdot p(j,k), \quad \forall k \in A_2,\ \forall k' \in A_2,$$

$$\sum_{j \in A_1, k \in A_2} p(j,k) = 1, \quad p(j,k) \geq 0, \quad \forall j \in A_1,\ \forall k \in A_2.$$

Today: Equilibrium Computation

- NE in a two-player, zero-sum game
- PSNE in a general-sum game
- MSNE in a two-player, general-sum game
- MSNE in a general-sum game
- CE in a general-sum game
- SSE in Stackelberg game (one leader, one follower)

Example: Inspection Game

An inspector chooses whether to inspect or not;
The inspectee chooses whether to cheat or not.

	Cheat	No Cheat
Inspect	-6, -9	-1, 0
No Inspection	-10, 1	0, 0

Exercise: What is the NE of the game? What are the expected utilities for each player?

Example: Inspection Game

An inspector chooses whether to inspect or not;
The inspectee chooses whether to cheat or not.

	Cheat	No Cheat
Inspect	-6, -9	-1, 0
No Inspection	-10, 1	0, 0

Exercise: What is the NE of the game? What are the expected utilities for each player?

(1/10, 1/5) with expected utility (-2, 0)

Example: Inspection Game

An inspector chooses whether to inspect or not;
The inspectee chooses whether to cheat or not.

	Cheat	No Cheat
Inspect	-6, -9	-1, 0
No Inspection	-10, 1	0, 0

What is the SSE of the game? What are the expected utilities for each player?

What is the WSE of the game? What are the expected utilities for each player?

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

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

Example: Inspection Game

An inspector chooses whether to inspect or not;
The inspectee chooses whether to cheat or not.

	Cheat	No Cheat
Inspect	-6, -9	-1, 0
No Inspection	-10, 1	0, 0

SSE: (1/10, No Cheat) with utility (-1/10, 0)

WSE: does not exist ($p > 1/10$, No Cheat)

Compute SSE in a Stackelberg Game

- Iterate over $a_f \in A_f$

$$\max_{x \in \Delta^\ell} \sum_{a \in A_\ell} x_a u_\ell(a, a_f)$$ $|A_l|$ variables

$$s.t. \sum_{a \in A_\ell} x_a u_f(a, a_f) \geq \sum_{a \in A_\ell} x_a u_f(a, a'_f), \ \forall a'_f \in A_f$$

$|A_f| - 1$ constraints

- Find the x^* associated to the LP with the highest objective
- Find the a_f^* that best respond to x^*