Hot Topics in Algorithmic Game Theory

Nicole Immorlica

Northwestern University and CWI

Chicago, USA, and Amsterdam, Holland

Course Outline

- Day One (today)
  - Introduction to normal form games
  - Equilibria notions and computability
  - Repeated games
- Day Two – Auction Theory
- Day Three – Social Networks

Let's Play a Game

The Median Game

- Guess an integer between 1 and 100
- Write your name and number on the card and pass it to the front
- The winner is person whose number is closest to 2/3rds of the median

- P R I Z E : this box of chocolate

The Median Game

Example: If the numbers are

Median is 45, and Ali wins because his guess is closest to 2/3 of the median, or 30.

The Median Game

Outline

- Definitions
  - Normal form games
  - Bi-matrix games
- Equilibria notions
  - Dominant strategy equilibria
  - Pure Nash equilibria
  - Mixed Nash equilibria
  - Approximate Nash equilibria
- Repeated games

Normal Form Games

- A game consists of a set of players {1, ..., n}, each with a set of strategies $S_1, \ldots, S_n$
- The strategy space S of the game is the set of vectors or strategy profiles $S_1 \times \ldots \times S_n$
- For any profile of strategies $s \in S$ and any player i, there is a payoff $\rho_i(s)$

Normal Form Games

Example: The Median Game

- Players: you
- Strategies: integers between 1 and 100
- Payoff: payoff to player i given profile s is a box of chocolates if s_i is closest to two-thirds of the median of the numbers in s and zero otherwise.

Bi-Matrix Games

- Two players, Row and Column
- Row has m strategies
- Column has n strategies
- Payoffs represented an (m x n) matrix A whose entries are pairs of numbers (x, y)
- Entry A_{ij} = (x, y) means that when Row plays i and Column plays j, the payoff to Row is x and the payoff to Column is y

Bi-Matrix Games

Example: Prisoners' Dilemma

	Deny	Confess
Deny	(-2, -2)	(-5, -1)
Confess	(-1, -5)	(-4, -4)

Outline

- Definitions
  - Normal form games
  - Bi-matrix games
- Equilibria notions
  - Dominant strategy equilibria
  - Pure Nash equilibria
  - Mixed Nash equilibria
  - Approximate Nash equilibria
- Repeated games

Game Theory

Given a game, can we predict which strategies the players will play?

Predicting Game Play

- Example: Prisoners' Dilemma

	Deny	Confess
Deny	(-2, -2)	(-5, -1)
Confess	(-1, -5)	(-4, -4)

Predicting Game Play

- In Prisoner's Dilemma, the best strategy of a player is to confess no matter what the other player does
- This is called a dominant strategy equilibrium
- Dominant strategy equilibria are very predictive, but often don't exist

Predicting Game Play

Example: The Median Game

- No dominant strategy equilibrium
- But,
  - consider a player i whose number $s_i \neq (2/3) \times \text{median}(s)$
  - then i should change s_i to equal 2/3 of the median
  - hence the only stable strategy profile is when everyone guesses 2/3 of the median
  - this can only happen when everyone guesses zero, and so the vector of all-zeros is the only stable strategy profile

Pure Nash Equilibria

This is called a pure Nash equilibrium.

- A profile is a pure Nash equilibrium if each player's strategy is his or her best choice given the other players' strategies:

$$\rho_i((s_1,\ldots, s_i, \ldots, s_n)) \geq \rho_i((s_1, \ldots, s'_i, \ldots, s_n))$$

for all i and s'_i.

Equilibria

- Equilibria attempt to determine which strategy profiles will be played
- A good equilibrium notion should
  - Always exist
  - Be natural
  - Be computable
- Is dominant strategy a good equilibrium notion? Is pure Nash equilibria a good notion?

Equilibria Notions

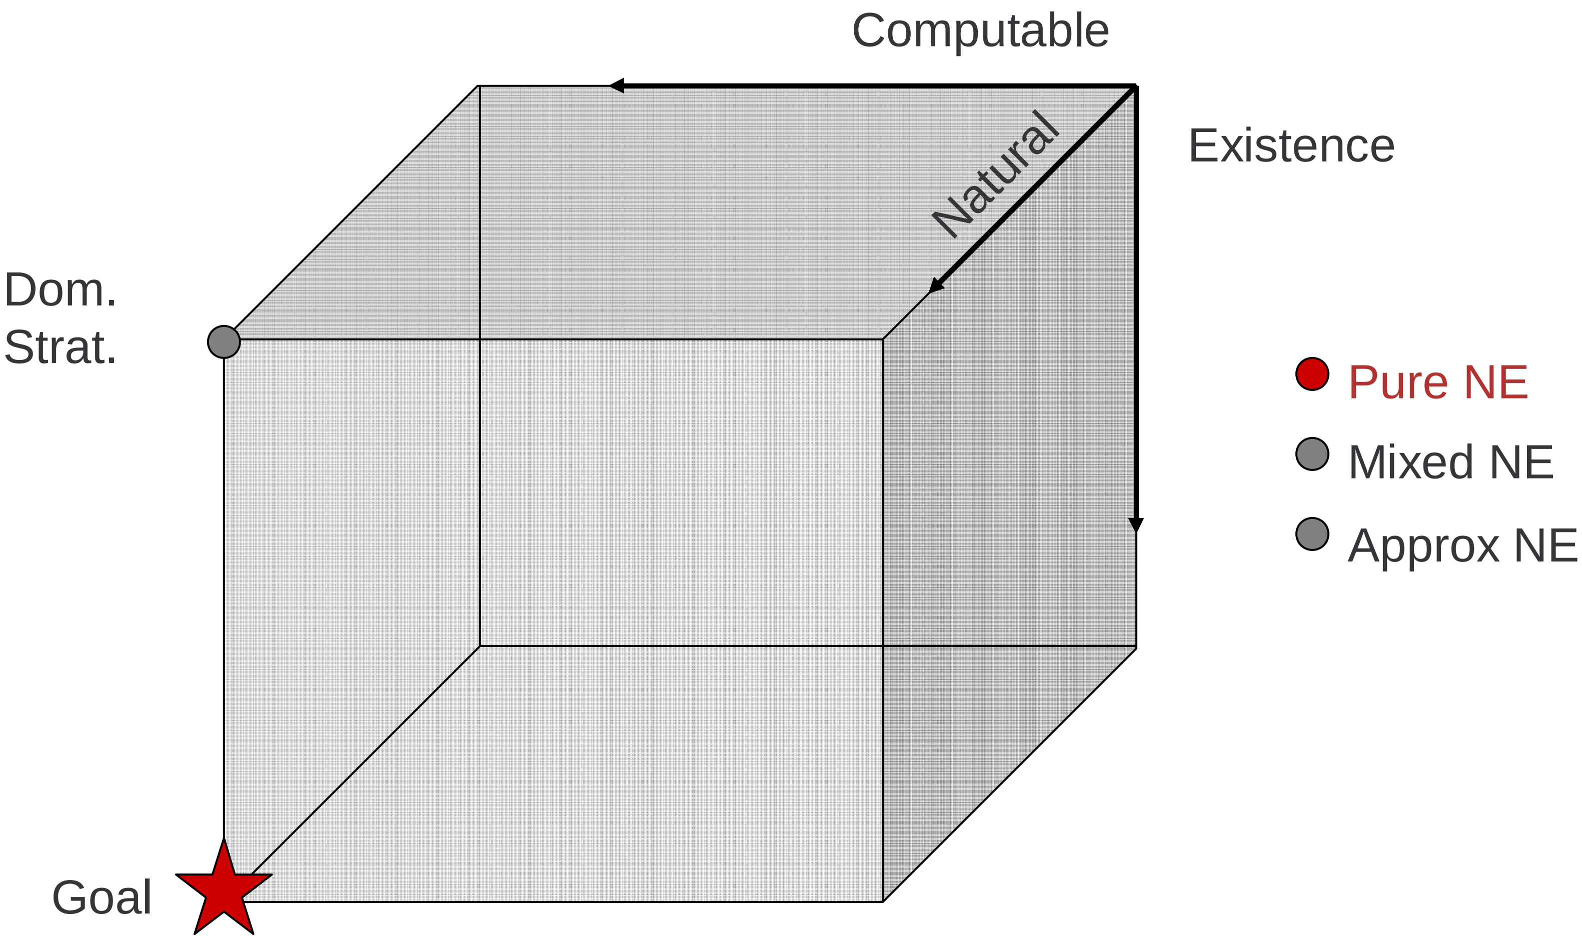

Pure Nash Equilibria

- Exist? Not necessarily.
- Example: Matching pennies game

	Heads	Tails
Heads	(1, -1)	(-1, 1)
Tails	(-1, 1)	(1, -1)

Pure Nash Equilibria

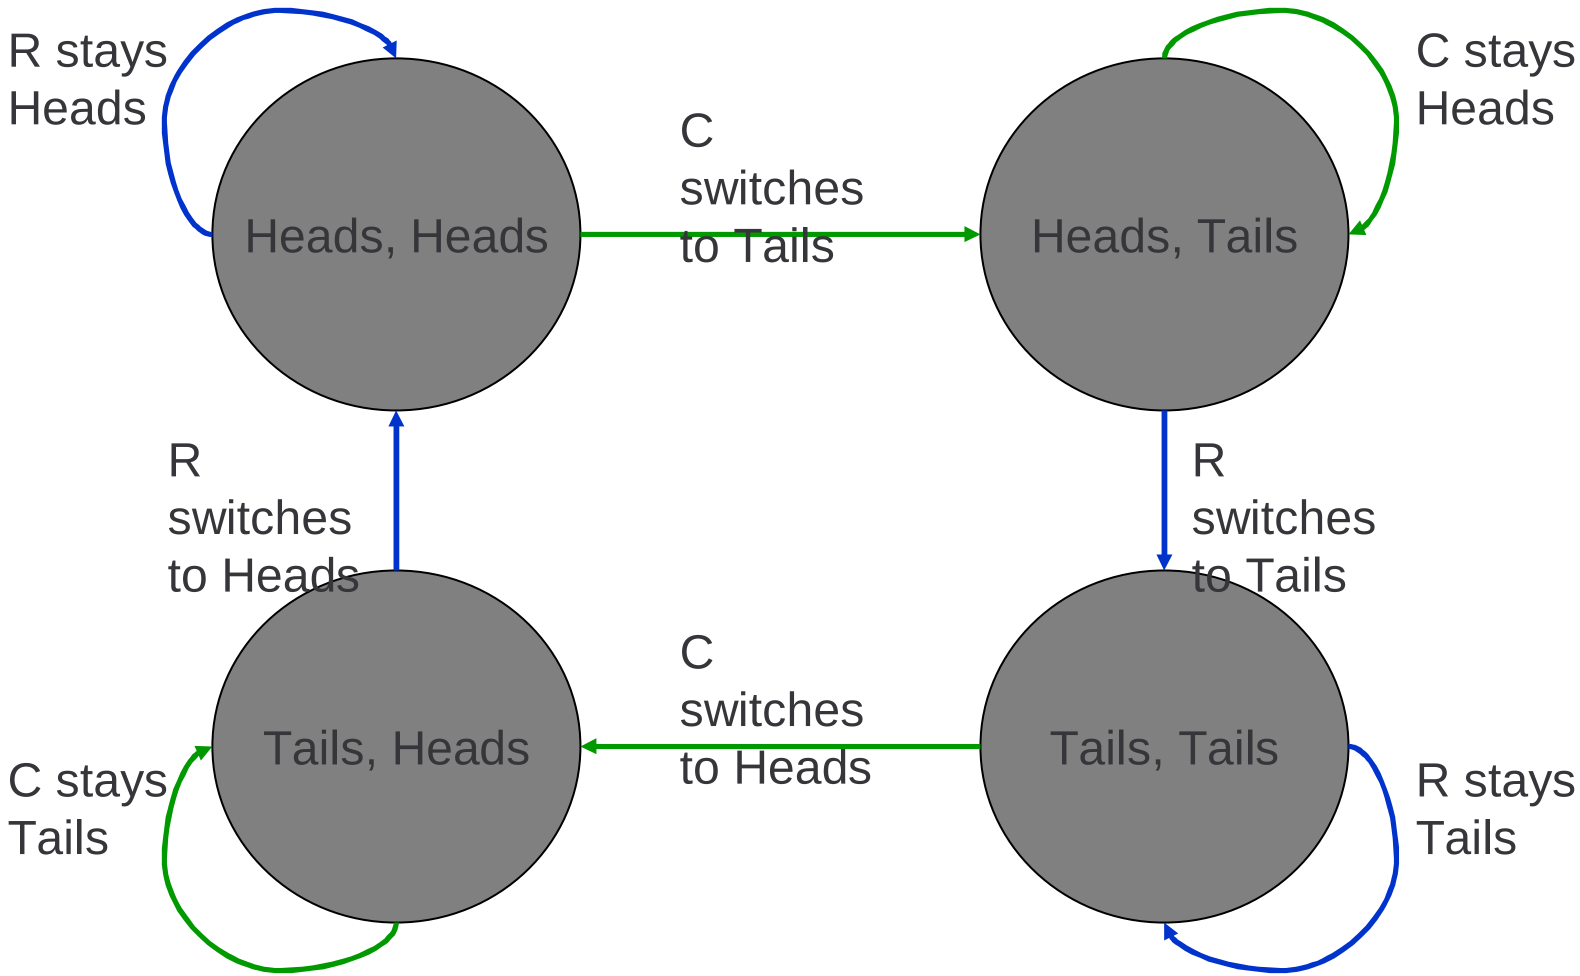

Pure Nash Equilibria

- Natural? Yes, but multiple equilibria.
- Example: Coordination game

	Theater	Football
Theater	(5, 4)	(2, 1)
Football	(1, 2)	(4, 5)

Pure Nash Equilibria

- Computable?
- Depends on game representation.
  - Bi-matrix games, can just check all entries.
  - In general, can be NP-hard to decide if one exists.
  - If one exists, finding it can be PLS-complete [FPT '03]

Equilibria Notions

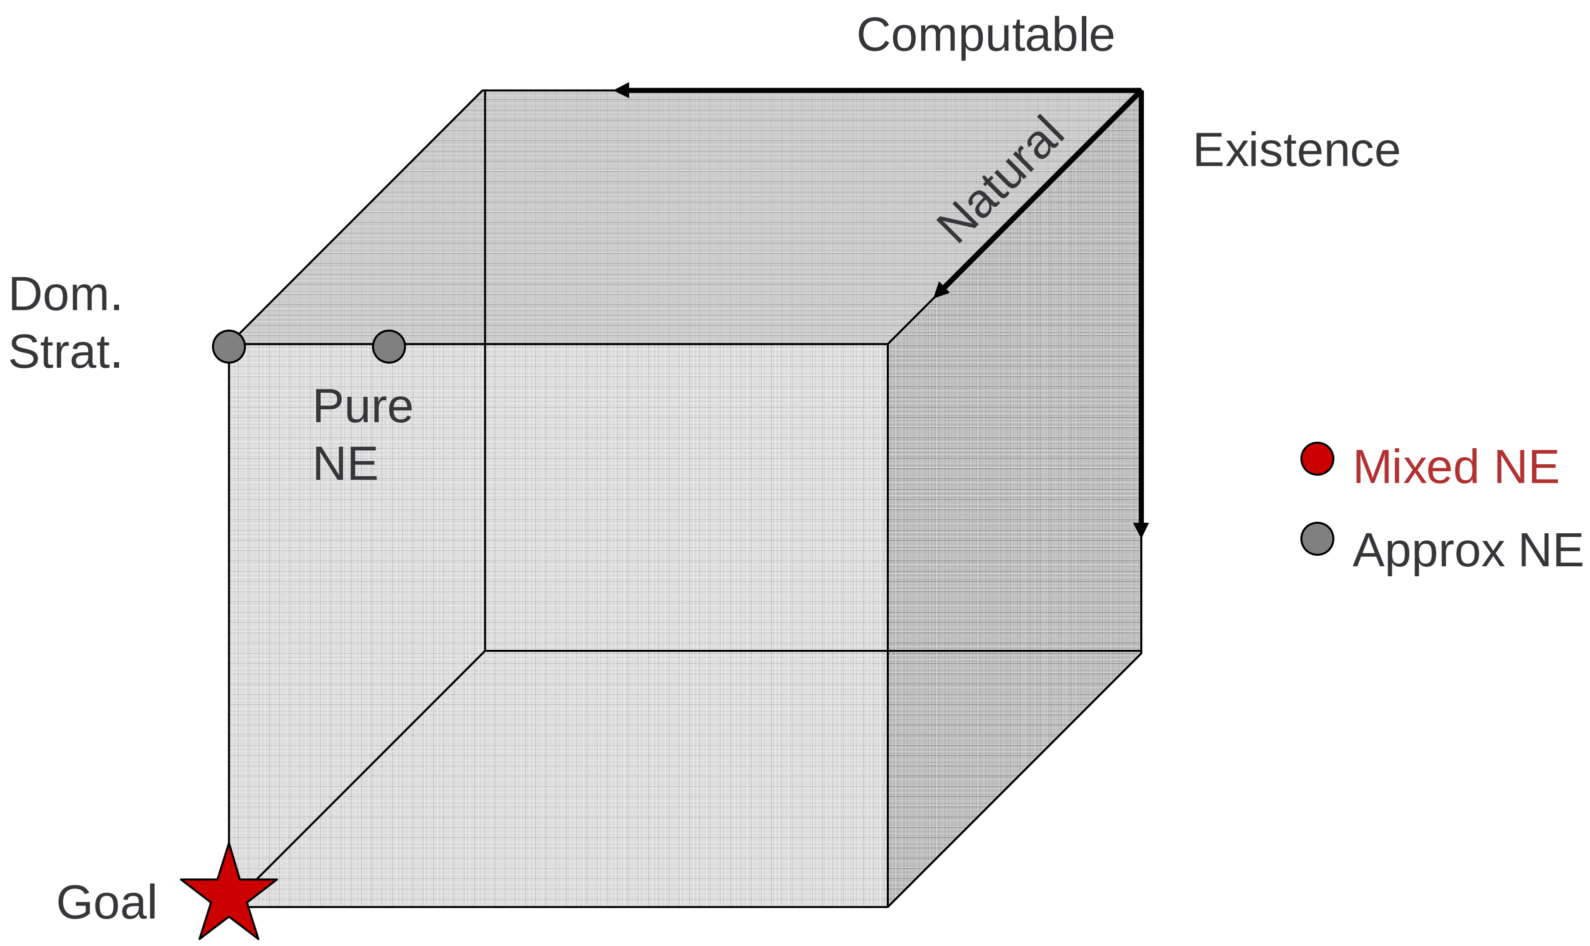

Mixed Nash Equilibria

- Recall Matching Pennies game
- Let players chose strategies probabilisitically

	Heads	Tails
Heads	(1, -1)	(-1, 1)
Tails	(-1, 1)	(1, -1)

Mixed Nash Equilibria

	1/2 Heads	1/2 Tails
1/2 Heads	(1, -1)	(-1, 1)
1/2 Tails	(-1, 1)	(1, -1)

Expected Payoff: (1/4) (1 + -1 + -1 + 1) = 0

Mixed Nash Equilibria

- This is the maximum payoff Row can acheive fixing the strategy of Column

	1/2	1/2
p	(1, -1)	(-1, 1)
1-p	(-1, 1)	(1, -1)

$$E[\rho_{Row}] = (1/2)p – (1/2)(1\text{-}p) – (1/2)(p) + (1/2)(1\text{-}p) = 0$$

Mixed Nash Equilibria

This is called a mixed Nash equilibrium.

- We enhance the strategy set S of a player to include probability distributions over pure strategies
- A mixed strategy σ_i is thus

$$\sigma_i : S_i \rightarrow [0,1] \text{ such that } \sum_{s \in S_i} \sigma_i(s) = 1$$

- And a mixed strategy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$

Mixed Nash Equilibria

This is called a mixed Nash equilibrium.

- A profile (of mixed strategies) is a mixed Nash equilibrium if:

$$E[\rho_i((s_1,\ldots, s_i, \ldots, s_n))] \geq E[\rho_i((s_1, \ldots, s'_i, \ldots, s_n))]$$

for all i and (mixed) σ'_i, where s_i are random vars independently distributed according to σ_i.

Mixed Nash Equilibria

- Exist? Yes, every game has a mixed NE.
- John Nash proved that NE exist in all finite games with finite strategy stets in 1950.
- Von Neumann remarked ``That’s trivial, you know. That’s just a fixed point theorem.’’

Mixed Nash Equilibria

- Proof:
  - Consider simplex S of mixed strategy profiles
  - Let F : S $\rightarrow$ S be the map which shifts profile in direction of best-response:

    $F(\sigma_1,\ldots,\sigma_n) = (\sigma_1+\delta_1, \ldots,\sigma_n + \delta_n)$ where δ_i is small shift in direction of best response of i to σ
  - By Brouwer's fixed point theorem, there is a σ^* such that $\sigma^* = F(\sigma^*)$
  - σ^* is a mixed NE since no one wants to deviate

Mixed Nash Equilibria

- Fixed point theorem – Sperner's Lemma
  (interpretation: color of node = gradient of F)

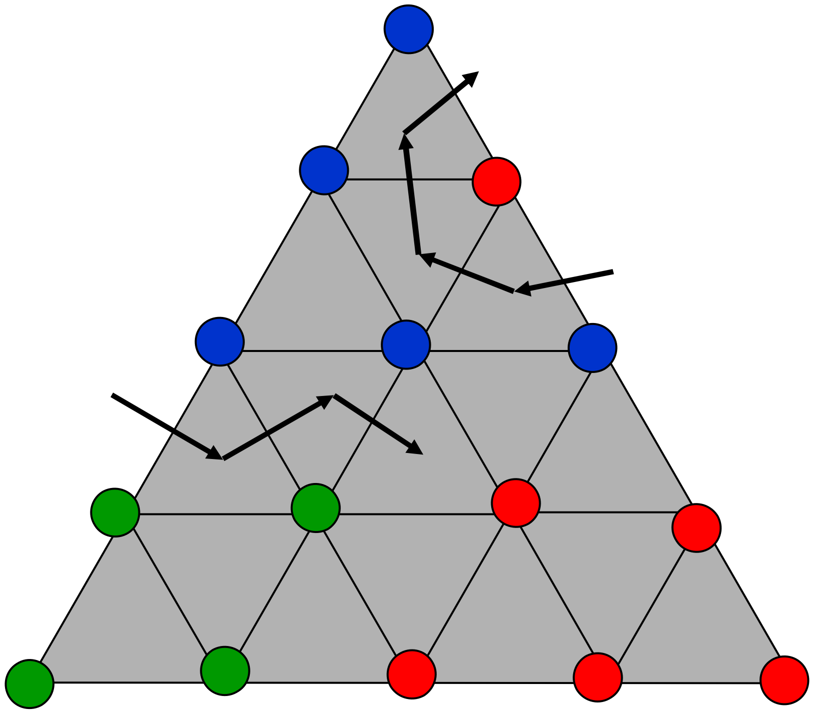

Mixed Nash Equilibria

Computable?

- Exponential-time algorithm for bi-matrix games
- PPAD-hard in general

Mixed Nash Equilibria

Bi-matrix games

- Let T_1 be the support of player 1's mixed strategy, p_i for $i \in T_1$ be probability 1 assigns to i
- Let T_2 be the support of player 2's mixed strategy, q_i for $i \in T_2$ be probability 2 assigns to i
- Let u be payoff of player 1 and v be payoff of player 2

Mixed Nash Equilibria

- Lemma: Each pure strategy in the support of a mixed NE must be a best-response to opponent's strategy.
- Proof: Suppose strategy i is not a best-response. Then shifting probability mass away from i improves payoff, contradiction assumption that profile was mixed NE.

Mixed Nash Equilibria

We have the following conditions

- p, q are valid probability distributions:

$$\sum_{i \in T1} p_i = 1,\ p_i > 0$$

$$\sum_{i \in T2} q_i = 1,\ q_i > 0$$

Mixed Nash Equilibria

We have the following conditions

- T_1 and T_2 are the support

$$p_i = 0 \text{ for } i \text{ not in } T_1 \text{ and } q_i = 0 \text{ for } i \text{ not in } T_2$$

Mixed Nash Equilibria

We have the following conditions

- Each $i \in T_1$ is a best-response to q (and vice-versa)

$$\sum_{j \in T2} q_j A_{ij}(1) = u \text{ for all } i \in T_1$$

$$\sum_{j \in T1} p_j A_{ij}(2) = v \text{ for all } i \in T_2$$

Mixed Nash Equilibria

Summarizing,

$$\sum_{i \in T1} p_i = 1,\ p_i > 0$$

$$\sum_{i \in T2} q_i = 1,\ q_i > 0$$

$$p_i = 0 \text{ for i not in } T_1$$

$$q_i = 0 \text{ for i not in } T_2$$

$$\sum_{j \in T2} q_j A_{ij}(1) = u \text{ for all } i \in T_1$$

$$\sum_{j \in T1} p_j A_{ij}(2) = v \text{ for all } i \in T_2$$

$2(n+1)+|T_1|+|T_2|$ equations, $2(n+1)$ unknowns.

Mixed Nash Equilibria

Computability

- Bi-matrix games: for all $2^n \times 2^n$ possible supports T_1 and T_2, check whether the system of linear equations has a solution such that u and v are best-response payoffs
- General games: Lemke-Howson, pivot-based algorithm similar to simplex algorithm for linear programming, also exponentional

Mixed Nash Equilibria

Computability

- Polytime algorithm?
- No! Nash equilibria are PPAD-complete, even for bi-matrix games [DGP'06, CDT'06]

Equilibria Notions

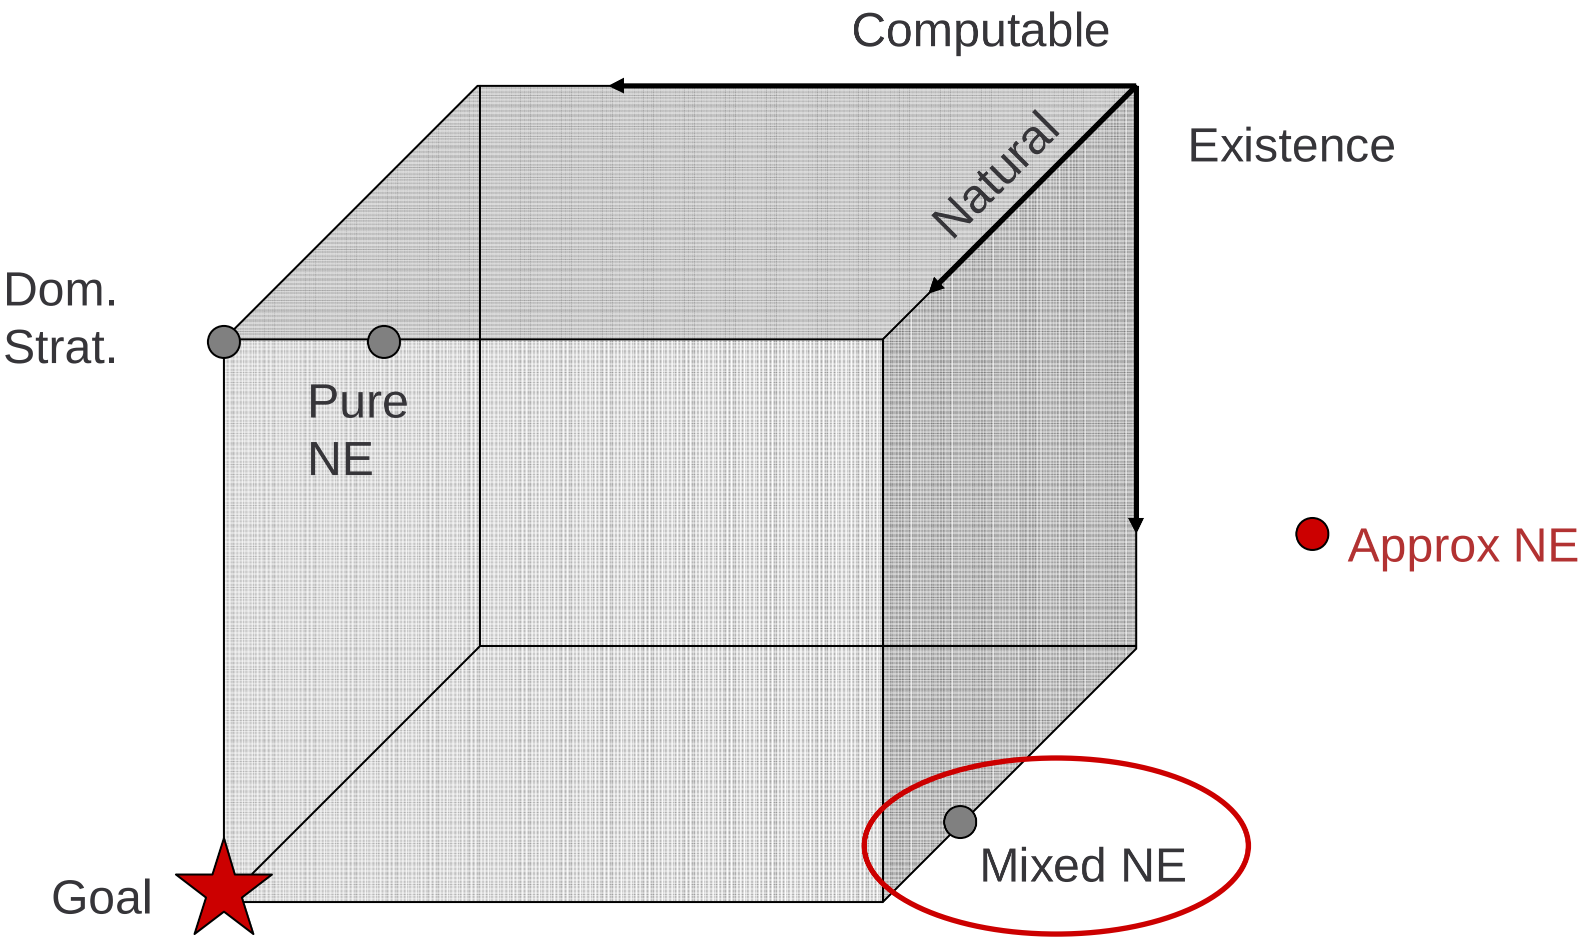

Approximate Nash Equilibria

- Mixed NE is a set of probability distributions s.t.

$$E[\rho_i((s_1,\ldots, s_i, \ldots, s_n))] \geq E[\rho_i((s_1, \ldots, s'_i, \ldots, s_n))]$$

- Approximate NE (ϵ-NE) allows additive error

$$E[\rho_i((s_1,\ldots, s_i, \ldots, s_n))] \geq E[\rho_i((s_1, \ldots, s'_i, \ldots, s_n))] - \epsilon$$

(normalization: payoffs are all in [0,1])

Approximate Nash Equilibria

- Computability in bi-matrix games

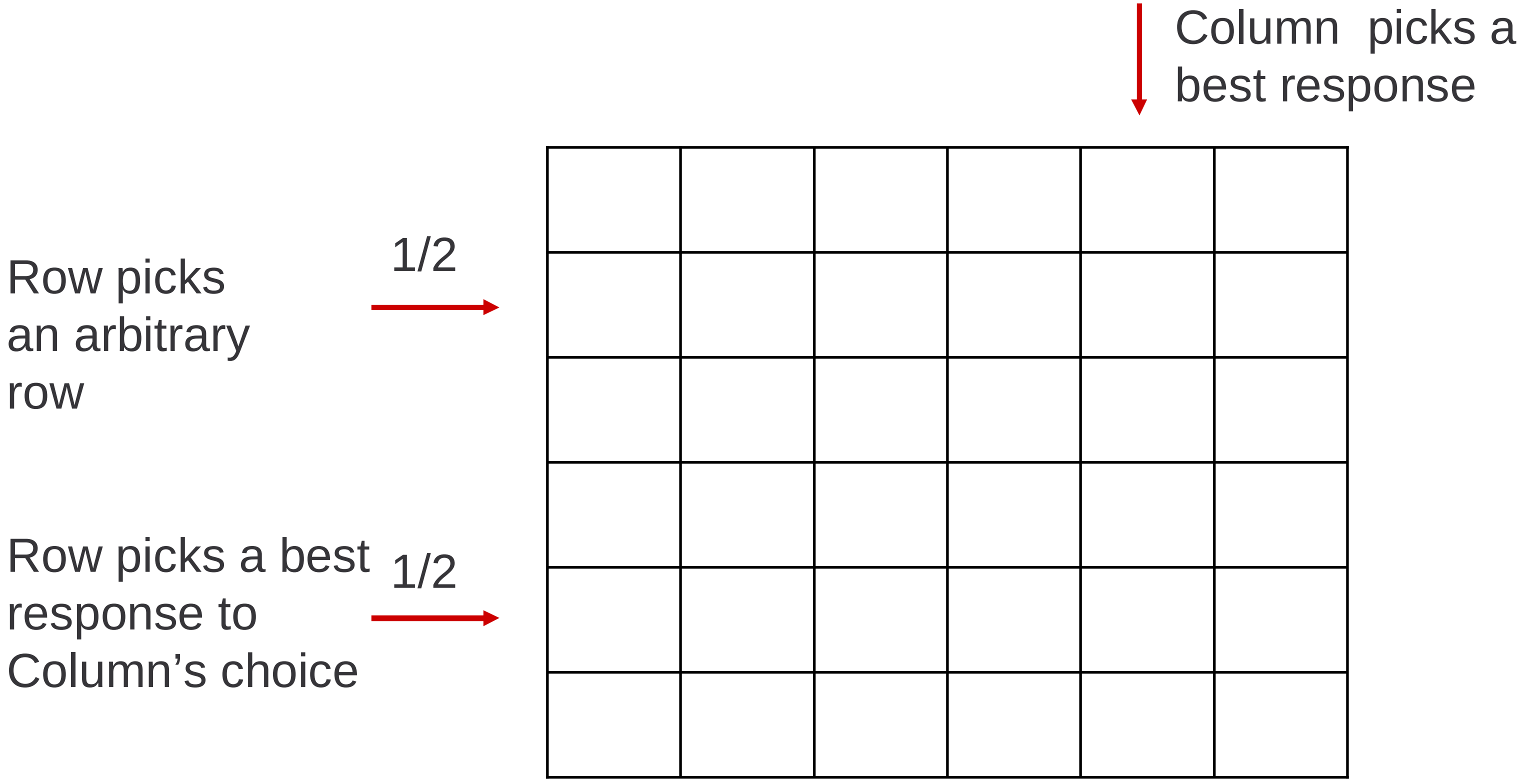

Approximate Nash Equilibria

- Theorem: The above algorithm yields a (1/2)-approximate NE.
- Proof: Each player is playing a best response to opponent’s strategy realization with probability (1/2). Hence deviations only improve payoff half the time, and since maximum payoff is 1, the result follows.

Approximate Nash Equilibria

- Theorem [FNS'07]: For ϵ<1/2, need strategies of support Ω(log n).
- Theorem [DMP'07, BBM'07, ST'07]: There exist algorithms for $\epsilon \approx$ 1/3.
- Critique: Approximate NE is not natural for large epsilon. Really, we want a PTAS.

Approximate Nash Equilibria

A sub-exponential time algorithm for all $\epsilon > 0$

Lipton, Markakis, Mehta 2003

- Definition: A k-uniform strategy is a mixed strategy where all probabilities are integer multiples of (1/k)

  e.g., σ_1 = (1/k, 10/k, 3/k, ..., 5/k)

Approximate Nash Equilibria

A sub-exponential time algorithm for all $\epsilon > 0$

- Key Lemma: For any $\epsilon \in (0,1]$ and for every k > O(log n / ϵ^2), there exists a pair of k-uniform strategies that form an ϵ-NE.

- Result: An algorithm for ϵ-NE that runs in time n^{2k}, sub-exponentional in n.
  - Simply check all n^{2k} possibilities for the k-uniform strategies and verify whether they are ϵ-NE

Approximate Nash Equilibria

- Proof Sketch (of Key Lemma): Based on probabilistic method
  - Let p^*, q^* be a NE
  - Sample k times from strategies of Row according to p^* to get k-uniform strategy p:

    If strategy i sampled m times, assign it probability m/k.

  - Same for Column to get k-uniform strategy q

Approximate Nash Equilibria

- Proof Sketch (of Key Lemma): Based on probabilistic method
  - Show that Pr[p,q are an ϵ-NE] > 0 (by Chernoff-Hoeffding bounds)

Approximate Nash Equilibria

Open Question:

Polynomial-time approximation scheme

Equilibria Notions

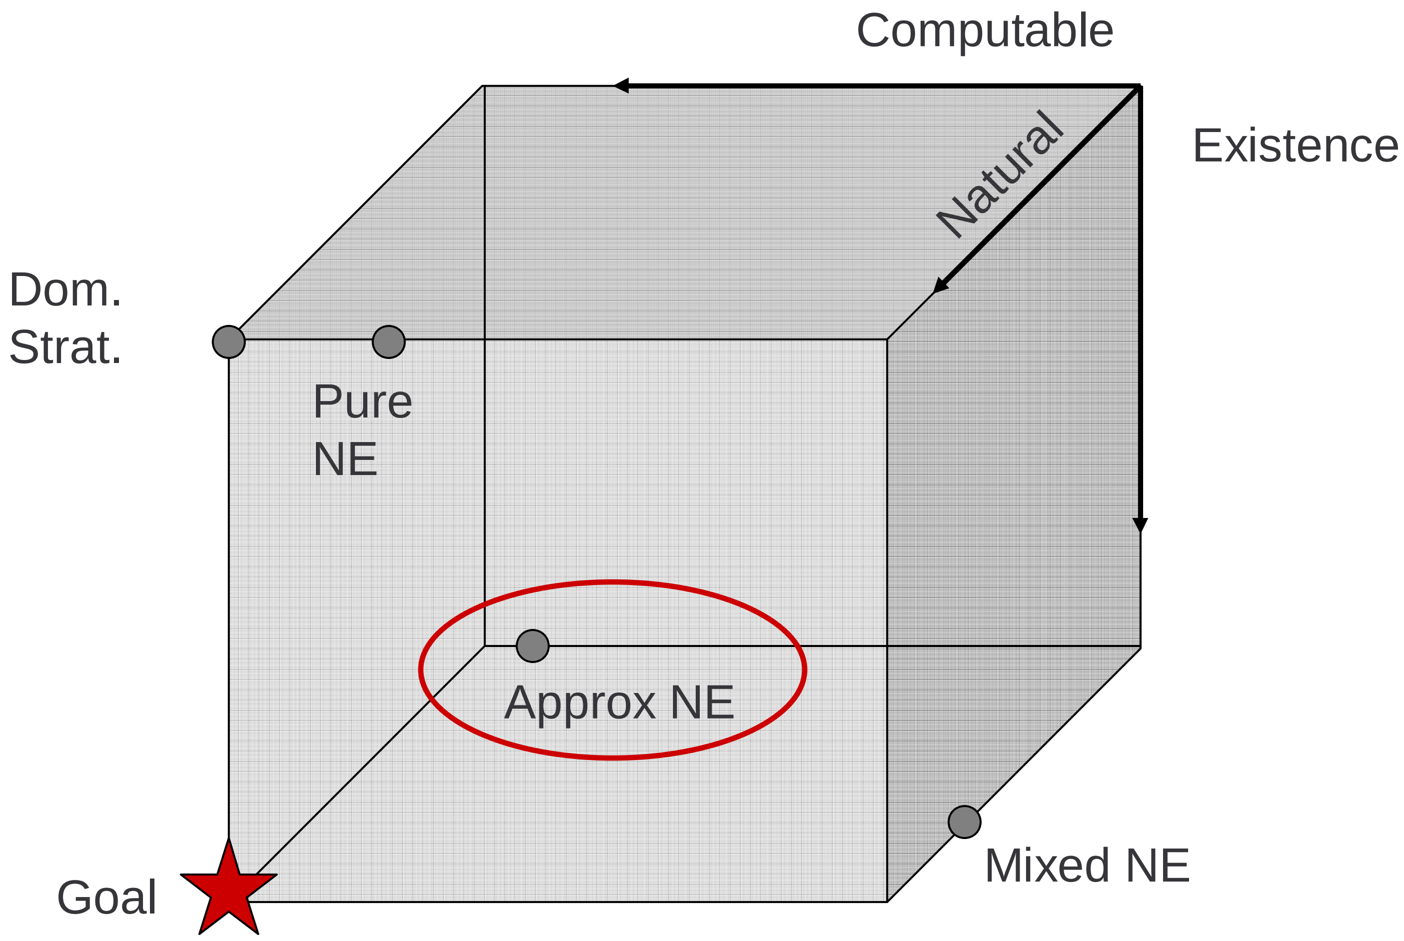

Break

Please return in 15 minutes.

Median Game

- Guesses:

1, 1, 1, 2, 13, 15, 20, 20, 22, 27,
30, 30, 32, 34, 34, 37, 46, 49

Median = 24.5

(2/3) of Median = 16.xx

Winner = Eaman

Outline

- Definitions
  - Normal form games
  - Bi-matrix games
- Equilibria notions
  - Dominant strategy equilibria
  - Pure Nash equilibria
  - Mixed Nash equilibria
  - Approximate Nash equilibria
- Repeated games

Repeated Games

- Recall Prisoner's Dilemma
- Unique NE, both Confess

	Deny	Confess
Deny	(-2, -2)	(-5, -1)
Confess	(-1, -5)	(-4, -4)

Repeated Games

- Suppose Prisoner’s Dilemma is played every day
- Row and Column can agree to both Deny
- If either ever plays Confess, opponent will punish by playing Confess for rest of time
- Threat of lower payoff in future encourages optimal play

Repeated Games

- There is an underlying game with actions A_i
- In each stage t, each player plays an action a^t_i from A_i and gets payoff $u_i(a^t)$
- The history $h^t = (a^1, ..., a^t) \in (A)^t$ describes the game play in the first t stages
- Let h denote the infinite game play

Repeated Games

- A pure strategy is now a function mapping any history of play to the next action

$$s_i : A^* \rightarrow A_i$$

  where $A^* = \cup_t (A)^t$ is all possible histories

- A mixed strategy σ_i is a function mapping any history into a probability distribution over actions

Repeated Games

- Discount factor δ describes tradeoff between present and future payoff
- Mixed strategies induce probability distributions over histories
- The total payoff to a player is expected discounted stage payoffs

$$E[\rho_i(h)] = \delta \sum_t (1-\delta)^t E[u_i(a^t)]$$

where the expectation is over histories

Repeated Games

- A Nash equilibrium is a mixed strategy profile such that no player can deviate and improve his or her payoff

Repeated Games

- Example: Prisoner's Dilemma
- Consider strategies $\sigma_1=\sigma_2$

$$\sigma_1(\text{Deny}^*) = \text{Deny}$$

$$\sigma_1(h) = \text{Confess for any other } h$$

- Then the realized history will be Deny*
- Payoff will be -2 for each player
- Consider deviation in which Row plays Confess at stage T+1

Repeated Games

Then the payoff to Row is

$$\delta \left[\sum^{T}_{t=1}(-2)(1-\delta)^{t} + (-1)(1-\delta)^{T+1} + \sum_{t>T+1}(-4)(1-\delta)^{t}\right]$$

$$= \delta \sum_{t}(-2)(1-\delta)^{t} + \delta\ (1-\delta)^{T+1} + \delta \sum_{t>T+1}(-2)(1-\delta)^{t}$$

$$= -2 + \delta\ (1-\delta)^{T+1} - 2(1-\delta)^{T+2}$$

$$< -2$$

whenever $\delta\ (1-\delta)^{T+1} - 2(1-\delta)^{T+2} < 0$ or $\delta < 2/3$

Repeated Games

- Hence, strategies $\sigma_1 = \sigma_2$

$$\sigma_1(\text{Deny}^*) = \text{Deny}$$
$$\sigma_1(h) = \text{Confess for any other } h$$

are a Nash equilibrium for small enough δ

Repeated Games

- Folk Theorem: Consider the achievable payoffs of the stage game.

(3, 3)	(0, 4)
(4, 0)	(1, 1)

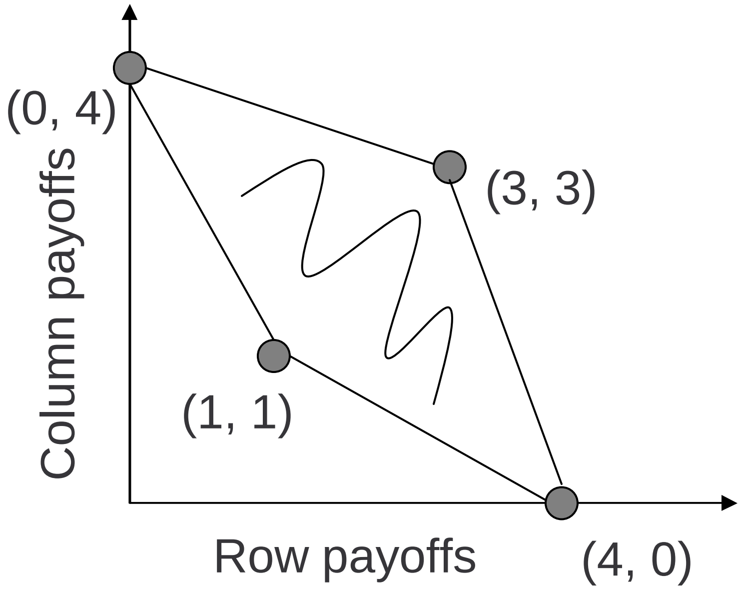

Repeated Games

- Folk Theorem: Compute min-max payoffs for each player (the best payoff a player can guarantee him or herself) – the Threat Point.

(3, 3)	(0, 4)
(4, 0)	(1, 1)

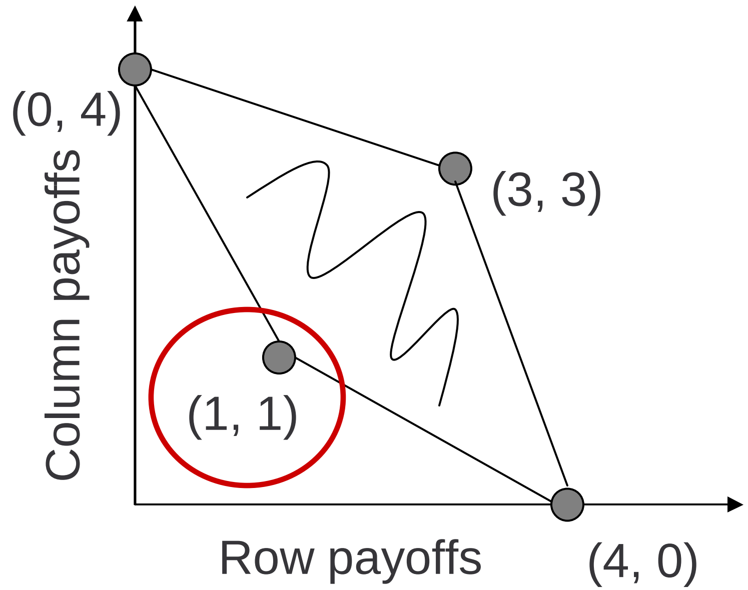

Repeated Games

- Folk Theorem: Individually rational region is set of payoffs above the Threat Point.

(3, 3)	(0, 4)
(4, 0)	(1, 1)

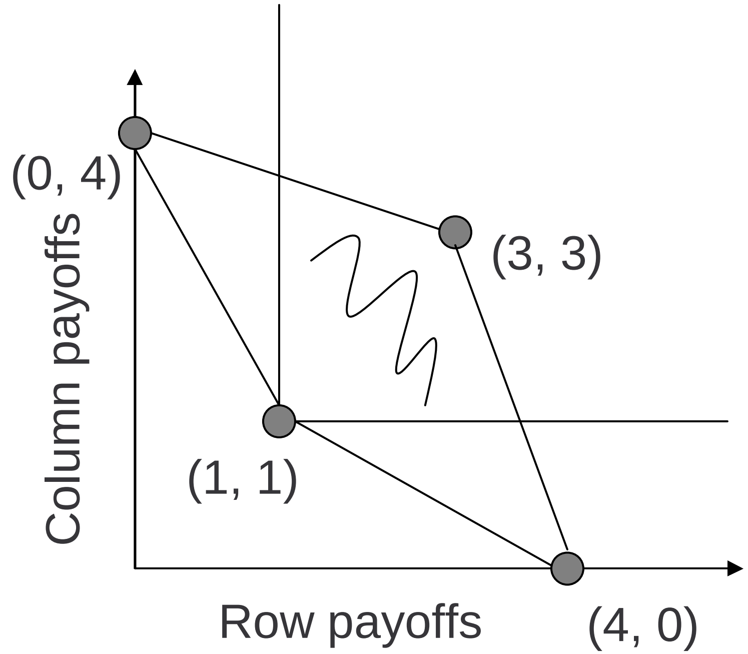

Repeated Games

- Folk Theorem: Any payoff in the individually rational region is acheivable in a Nash equilibrium.

(3, 3)	(0, 4)
(4, 0)	(1, 1)

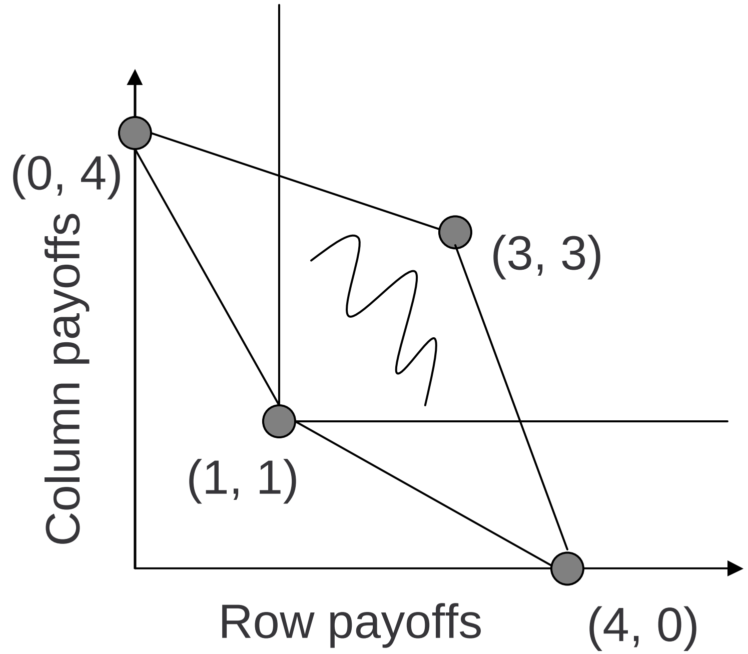

Repeated Games

- Folk Theorem: Pick program of play that acheives chosen payoffs. If a player deviates, switch to Threat Point. This is a NE.

(3, 3)	(0, 4)
(4, 0)	(1, 1)

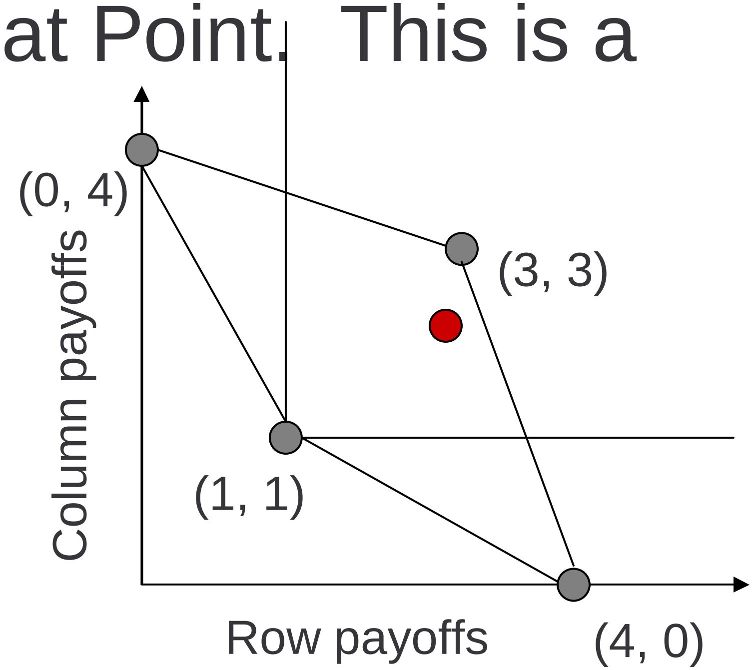

Repeated Games

- Problem! Threat Point may be NP-hard to compute.
- We reduce the problem of computing the Threat Point to 3-coloring.

Repeated Games

- Let G=(V,E) be graph we wish to 3-color
- There are three players 1, 2, and 3
- Show computing minmax payoff of 3 is hard

Repeated Games

- Stragies of 1 and 2: pick a node and a color

$$S_1 = S_2 = V \times \{Red, Green, Blue\}$$

- Strategy of 3: pick a player and a node

$$S_3 = \{1, 2\} \times V$$

Repeated Games

- Let (v_1, c_1), (v_2, c_2), (i, v_3) be strategy profile

- Payoff to 3 is
  - 1 if $v_1 = v_2$ and $c_1 \neq c_2$
  - 1 if $(v_1, v_2) \in E$ and $c_1 = c_2$
  - 1 if $v_3 = v_i$
  - 0 otherwise

1 and 2 are exposed

3 found i

Repeated Games

If G is 3-colorable, then min-max payoff of 3 is 1/n

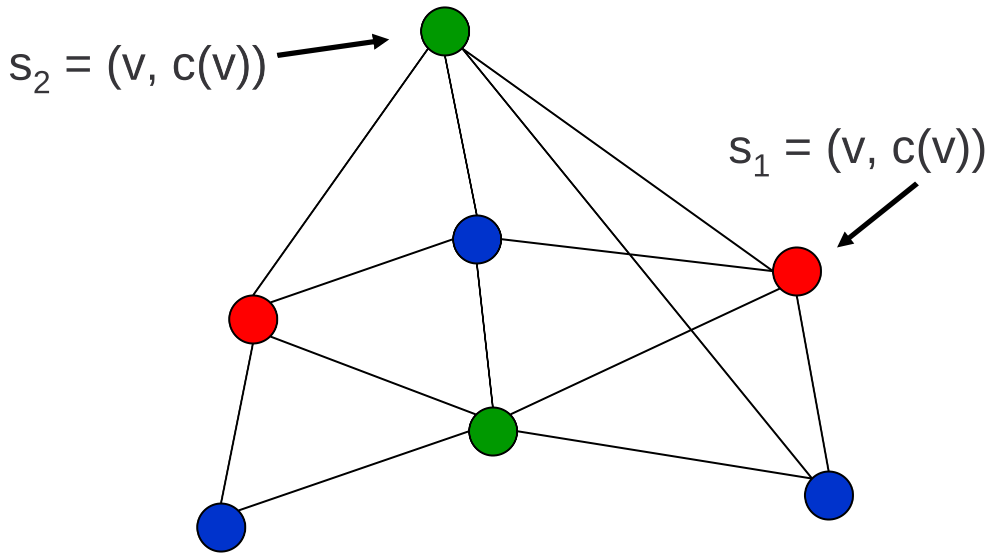

Players 1 and 2 pick uniformly random strategies consistent with 3-coloring c

Repeated Games

If G is 3-colorable, min-max payoff of 3 is 1/n:

- Player 3 gets payoff of 1 only by guessing the node of a player. Happens with probability 1/n.
- This is min that 1 and 2 can force upon 3 since there must be some node selected with probability at least 1/n

Repeated Games

If G is not 3-colorable, min-max payoff of 3 is at least $1/n + 1/(3n^2)$

- Consider any mixed strategies σ_1, σ_2
- Case 1:
  - For some $i \in \{1, 2\}$ and $v \in V$, i chooses v with probability at least $1/n + 1/(3n^2)$
  - Then 3 plays (i, v_i)

Repeated Games

- Case 2:
  - For all $i \in \{1, 2\}$ and $v \in V$, i chooses v with probability at most $1/n + 1/(3n^2)$
  - Then 3 plays (i, v_3) for a random i and v_3
  - With probability $1/n$, $v_3 = v_i$ and 3 gets 1
  - We show 1 and 2 are exposed with prob. $> 4/(9n^2)$
  - This implies payoff of 3 is at least

$$(1/n) + (1-1/n)(4/9n^2) \geq 1/n + 1/(3n^2)$$

Repeated Games

- Case 2:
  - For all $i \in \{1, 2\}$ and $v \in V$, i chooses v with probability at most $1/n + 1/(3n^2)$

    $\rightarrow$ i chooses node v with probability at least $2/(3n)$

Repeated Games

- Case 2:
  - Let E be event 1 and 2 are exposed

$$\Pr[\,E\,] = \Sigma_{x,y \in V} \Pr[v_1=x] \Pr[v_2=y] \Pr[\,E \mid v_1=x, v_2=y]$$
$$\geq 4/(9n^2)\, \Sigma_{x,y \in V} \Pr[\,E \mid v_1=x, v_2=y]$$

  - But sum is expected number of inconsistencies in random colorings induced by mixed strategies.
  - This is at least 1 by probabilistic method.

Repeated Games

Theorem [BCIKMP'08]: Given a 3-player n x n x n game with payoffs in {0,1} it is NP-hard to approximate the min-max payoff for each of the players to within $1/(3n^2)$.

Repeated Games

- What about other algorithms for ϵ-NE?
- Theorem [BCIKMP'08]: Finding an ϵ-NE of a k-player repeated game is as hard as finding an ϵ-NE of a (k-1)-player single-shot game.
- Proof Sketch:
  - Take (k-1)-player game G
  - Construct repeated Kibitzer version (next slide)
  - Extract equilibrium from Kibitzer version

Repeated Games

Kibitzer version of G

Player 2 ↓

Kibitzer →

Player 1 →

				5,5	
				8,7	
1,2	5,5	4,7	2,2	1,0	0,0
				3,8	

Payoffs: $\rho_1 = -7$, $\rho_2 = 0$, $\rho_{kibitizer} = 7$

Repeated Games

- Consider an ϵ-NE of the repeated Kibitzer version of a game G.
- If players are not playing an ϵ-NE of G, then Kibitzer has a response that improves payoff by more than ϵ.

Open Question:

Complexity in specific game classes

Next Time

Auction Theory