Equilibrium Computation

Ruta Mehta

ILLINOIS
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

AGT Mentoring Workshop
18^{th} June, 2018

Q: What outcome to expect?

Multiple self-interested agents interacting in the same environment

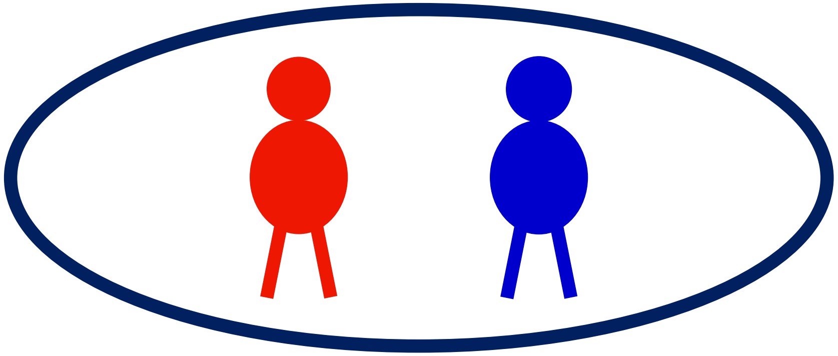

Deciding what to do.

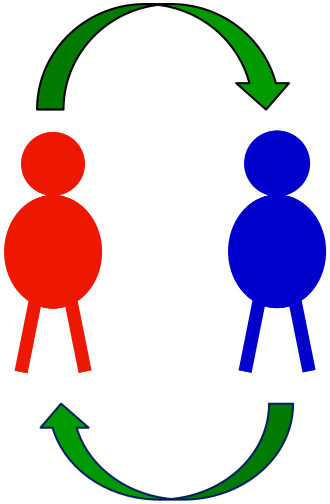

Q: What to expect?

Probably a “stable outcome” = equilibrium

Fig courtesy Vince Contizer

100+ Years of Extensive Work

Walras (1874)

von Neumann (1928)

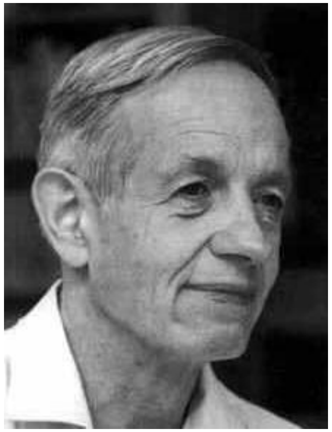

Nash (1950)

Arrow-Debreu (1954)

Gale-Shapley (1962)

This Talk

Games, Nash equilibrium, Algorithms, Complexity

Potential Games

- Network-flow, congestion

Extensive form games.

Commitment: Stackleberg equilibrium

- Application: Security games

Repeated games

(sessions 3B and 7B)

This Talk

- Games, Nash equilibrium, Algorithms, Complexity
- Potential Games
  - ☐ Network-flow, congestion
- Extensive form games
- Commitment: Stackleberg equilibrium
  - ☐ Application: Security games
- Repeated games

(sessions 3B and 7B)

Games

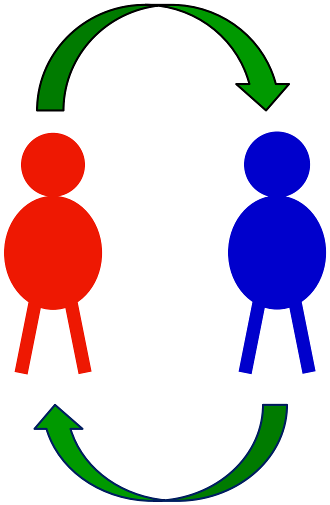

Players

Payoffs

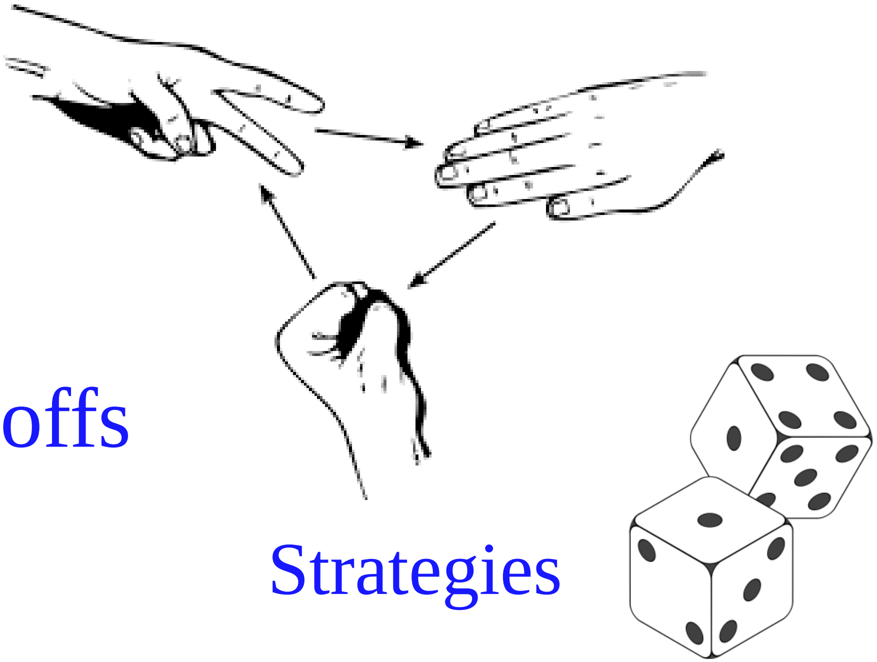

Strategies

Randomize!

Games (normal-form)

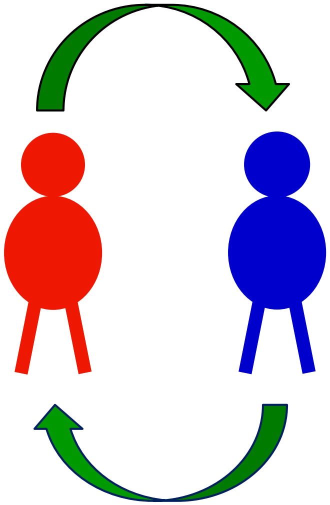

Players

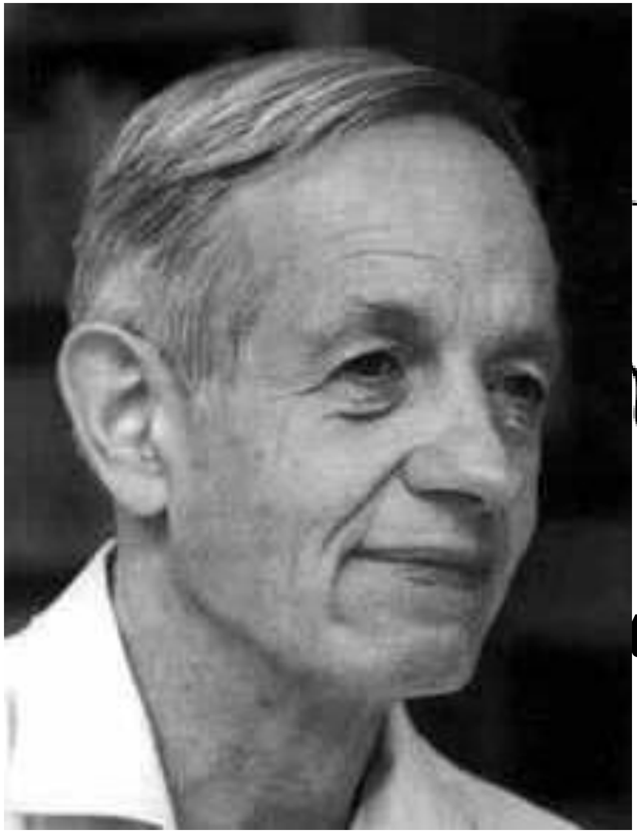

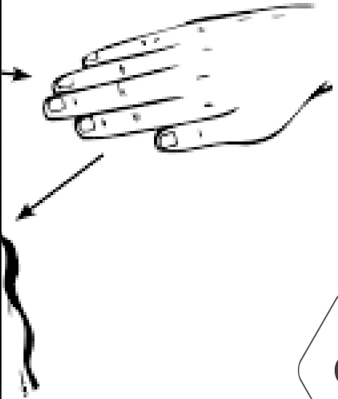

tegies

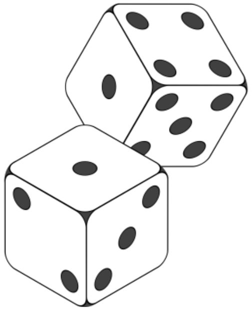

Randomize!

Nash (1950):

There exists a (stable) state where no player gains by unilateral deviation.

Nash equilibrium (NE)

Computation?

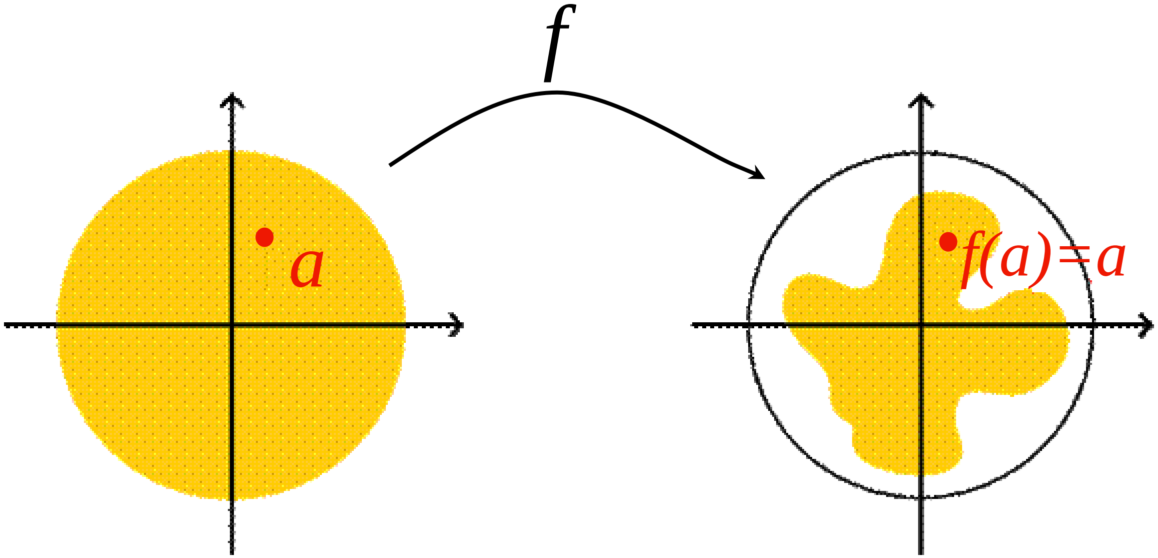

NE existence via fixed-point theorem.

Computation? (in Econ)

- Special cases: Dantzig’51, Lemke-Howson’64, Elzen-Talman’88, Govindan-Wilson’03, …
- Scarf’67: Approximate fixed-point.
  - Numerical instability
  - Not efficient!

Most are path following (complementary pivot) algorithms

Visualizing Fixed Point

Given $f\colon [0,1]^2 \rightarrow [0,1]^2$, direction vectors of $(f(x) - x)$

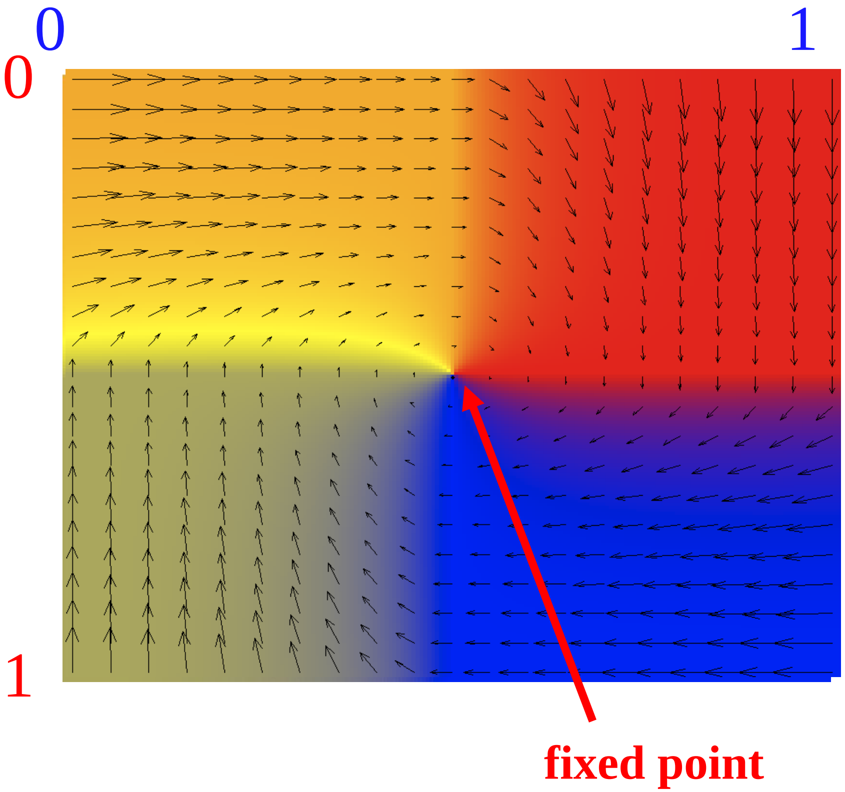

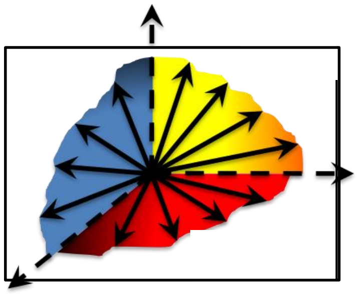

Next 5 slides are curtesy Costis Daskalakis

Visualizing Discrete Fixed Point

Given $f: [0,1]^2 \to [0,1]^2$, direction vectors of $(f(x) - x)$

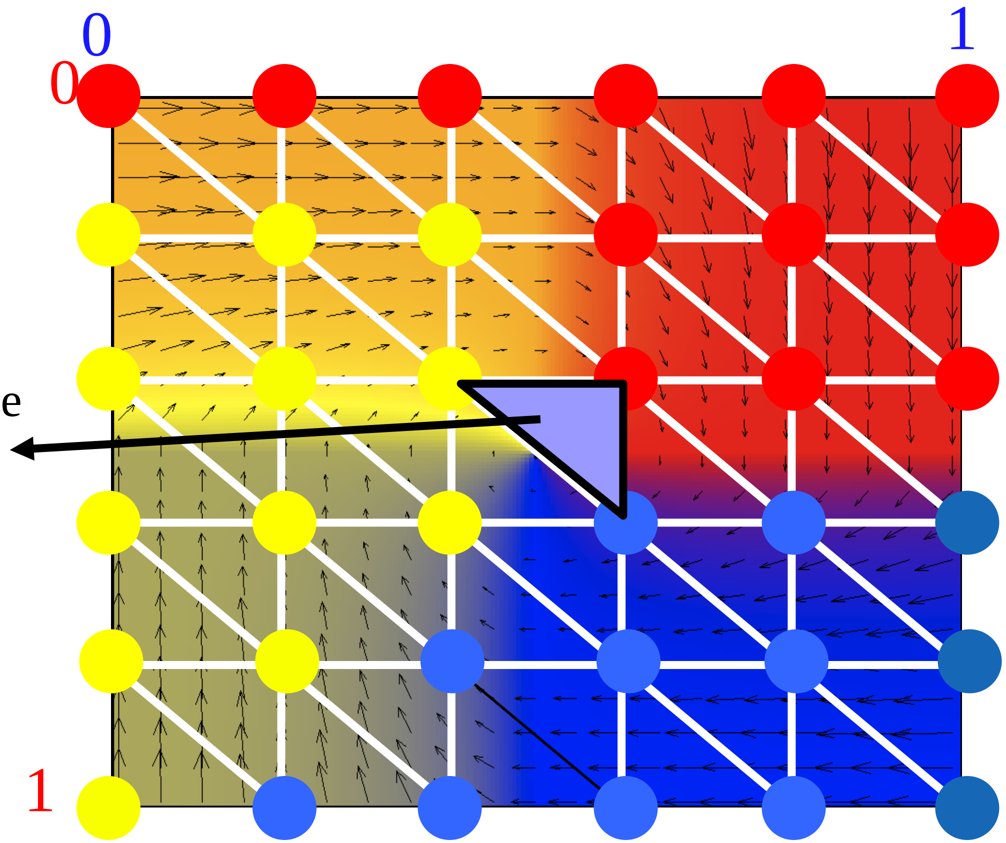

Trichromatic triangle
= fixed point

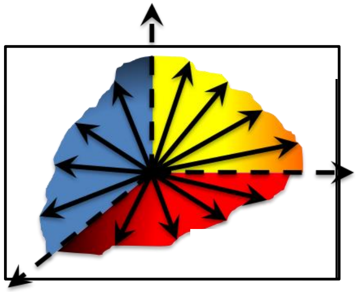

Fixed Point → Sperner's Lemma

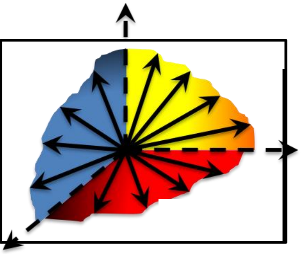

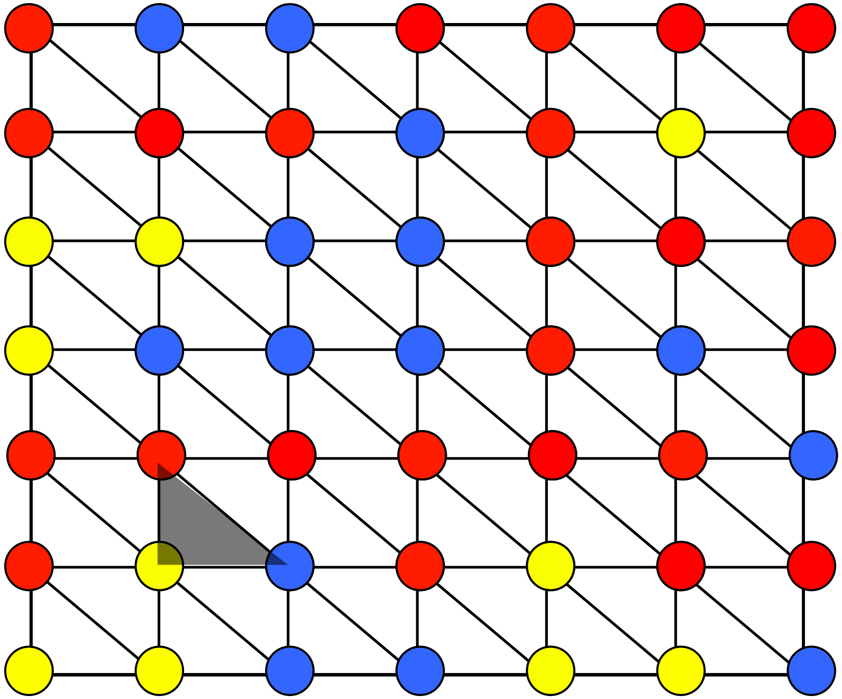

[Sperner 1928]: Color the boundary using three colors in a "legal way". No matter how the internal nodes are colored, there exists a tri-chromatic triangle. In fact an odd number of those.

Sperner's Lemma

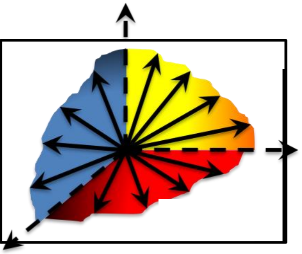

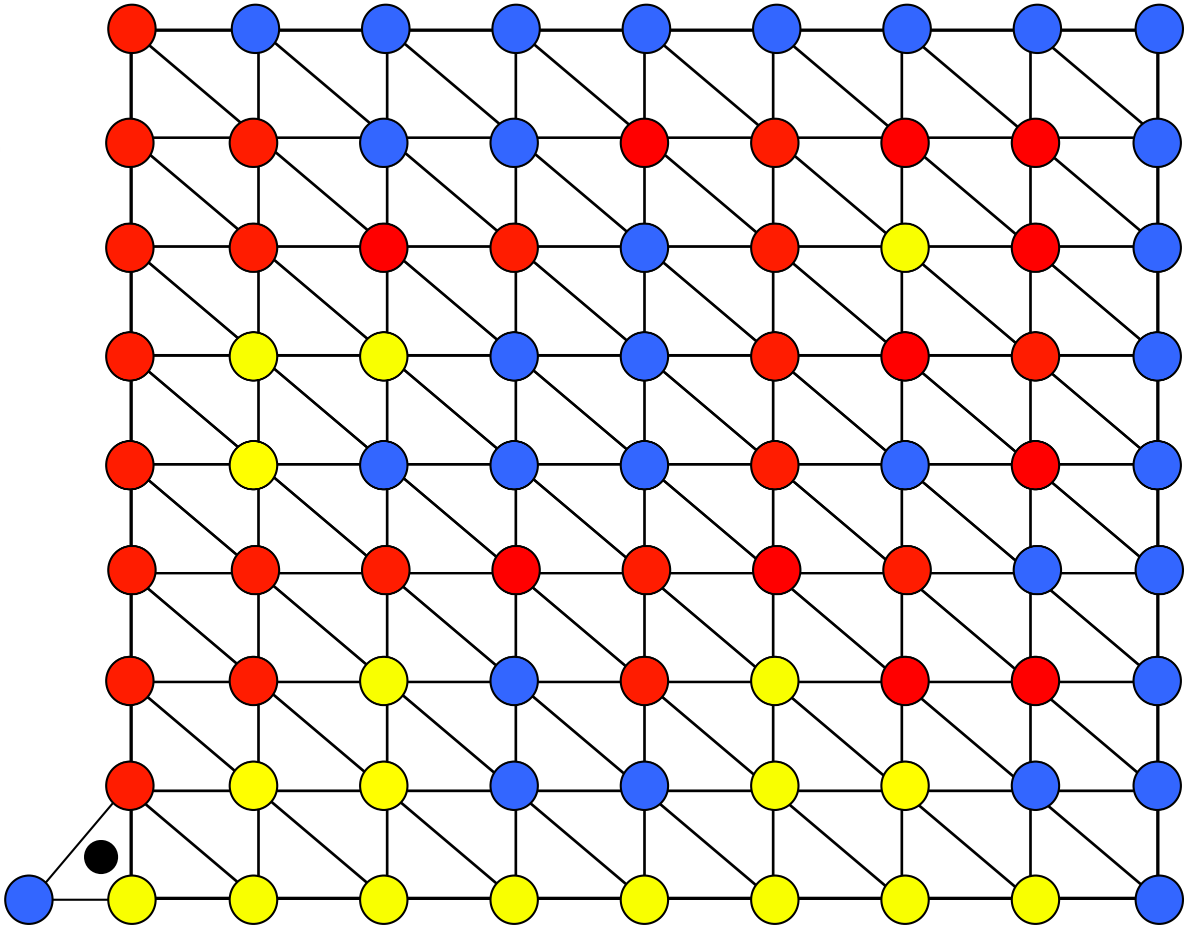

For convenience we introduce an outer boundary, that does not create new tri-chromatic triangles.

Also introduce an artificial tri-chromatic triangle.

Define a directed walk starting from the artificial tri-chromatic triangle.

[Sperner 1928]: Color the boundary using three colors in a "legal way". No matter how the internal nodes are colored, there exists a tri-chromatic triangle. In fact an odd number of those.

Sperner's Lemma: Directed walk

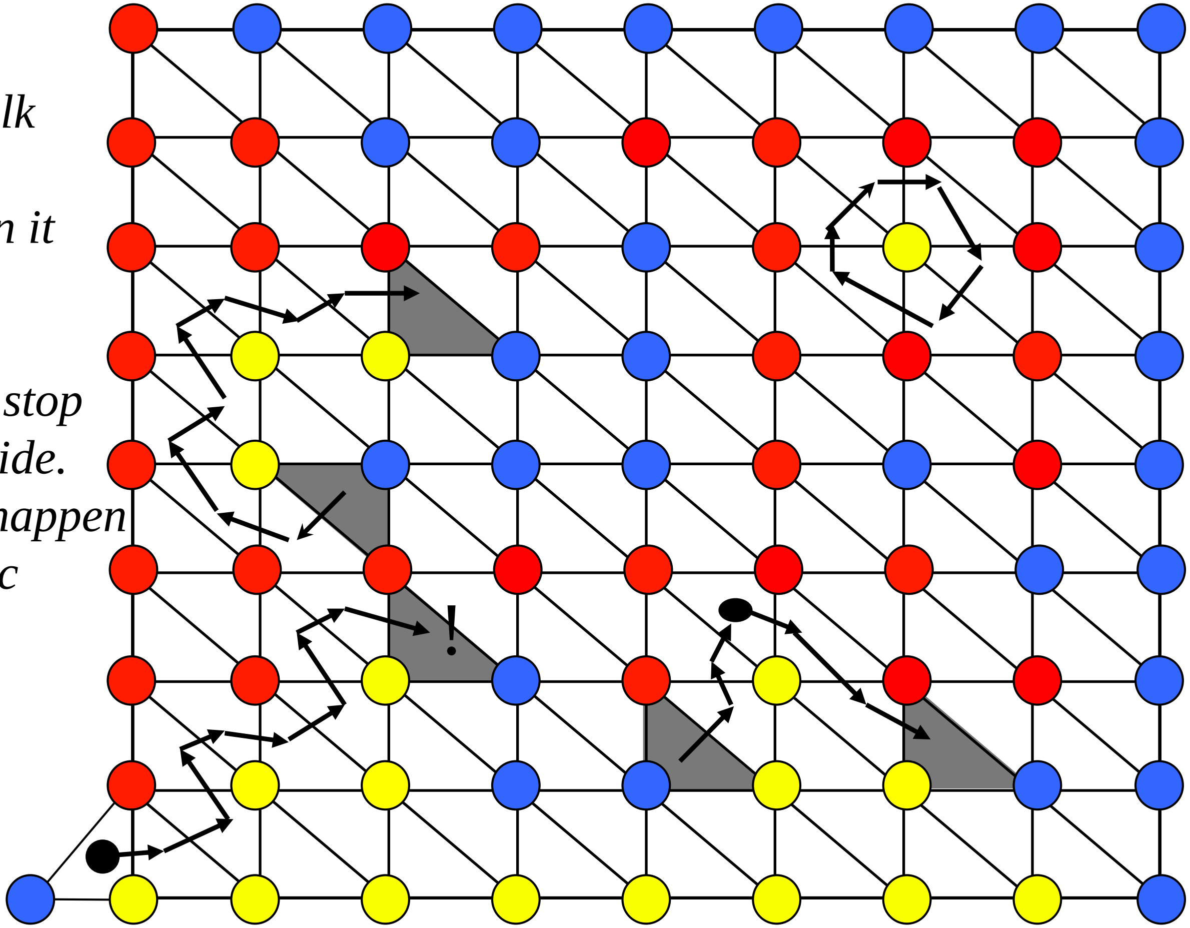

Claim**:** *The walk cannot exit the square, nor can it loop into itself.*

Hence, it must stop somewhere inside. This can only happen at tri-chromatic triangle…

For convenience introduce an outer boundary, that does not create new tri-chromatic triangles.

Also introduce an artificial tri-chromatic triangle.

Next we define a directed walk.

[Sperner 1928]: Color the boundary using three colors in a legal way. No matter how the internal nodes are colored, there exists a tri-chromatic triangle. In fact an odd number of those.

Computation? (in CS)

Not easy!

∃ solution?

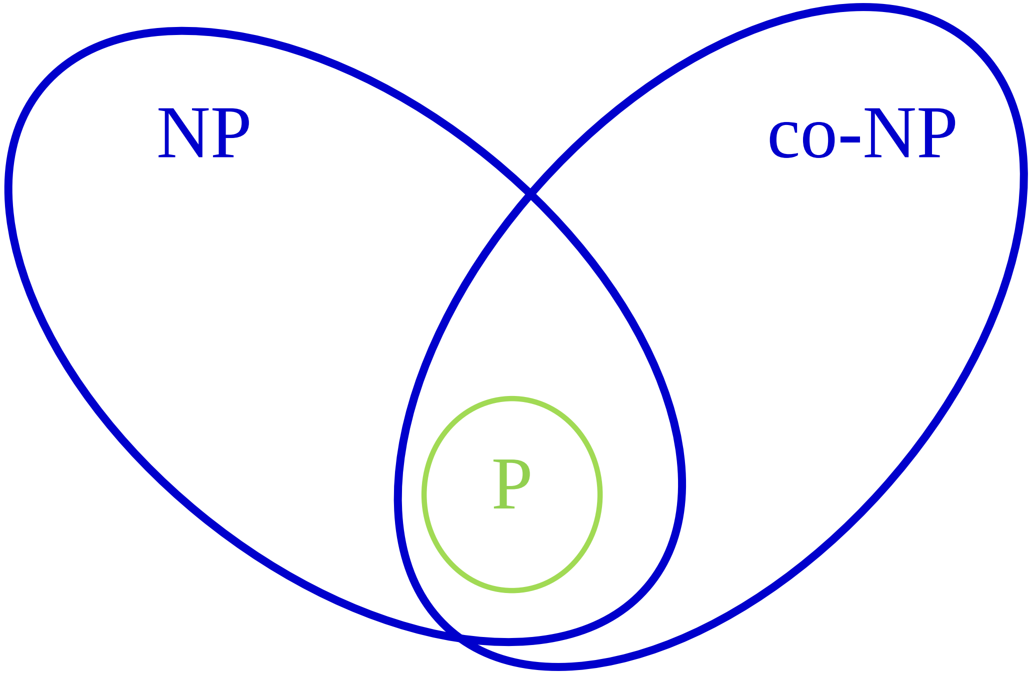

What if solution always exists? Like Nash Eq.?

Computation? (in CS)

Megiddo and Papadimitriou'91 :

Nash is NP-hard ⇒ NP=Co-NP

NP-hardness is ruled out!

Papadimitriou'94

PPAD Polynomial Parity Argument for Directed graph

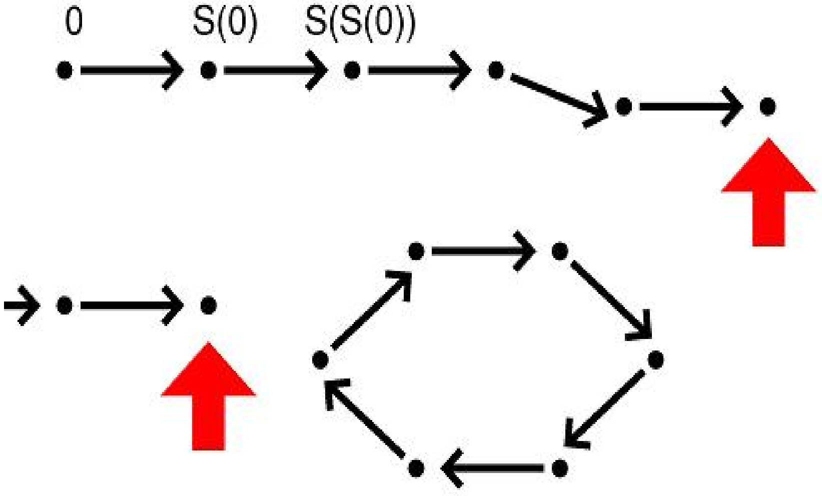

Find an end

Approximate fixed-point is
PPAD-complete. $|f(x) - x| < \epsilon$

$f(x) = x$

Papadimitriou’94

PPAD

Etessami & Yannakakis’07

FIXP

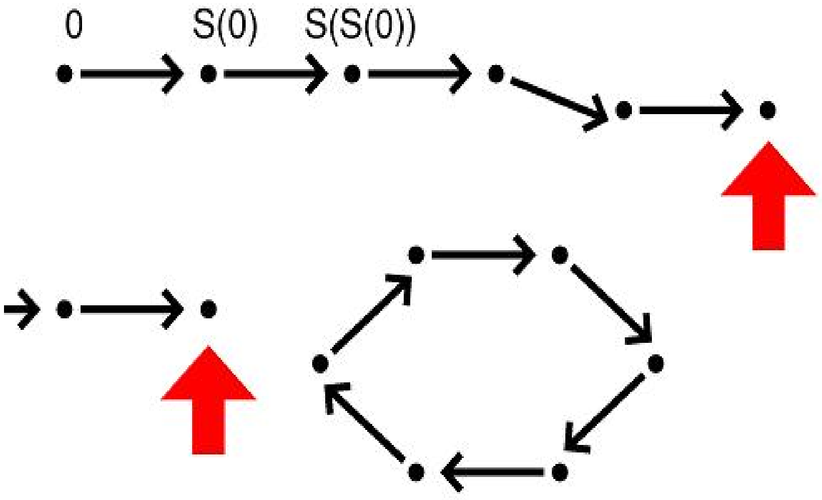

Find an end

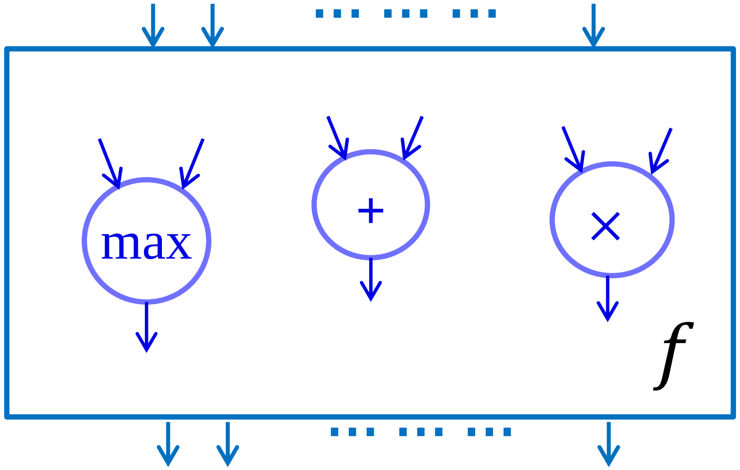

Approximate fixed-point is PPAD-complete. $|f(x) - x| < \epsilon$

Find a fixed-point

$f(x) = x$

$x \longleftrightarrow x$

away

Rational

Irrational but algebraic

Complexity Classes

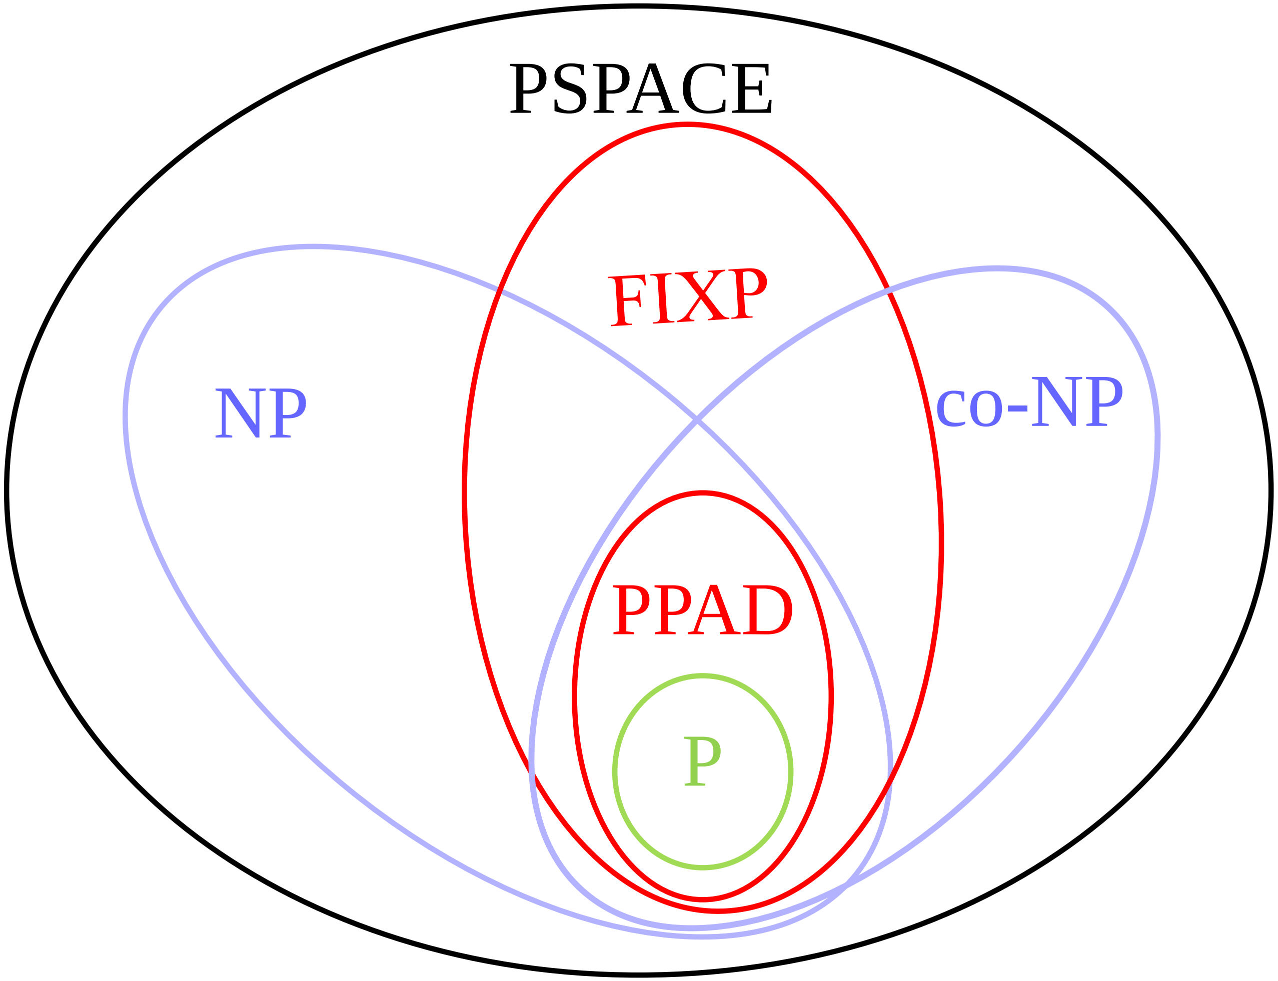

NE in 2-player game

	2-Nash	**k-Nash, $k > 2$**
Nature of solution	Rational	Algebraic; Irrational e.g.: Nash’51
Complexity	PPAD-complete [DaskalakisGoldbergPapadimitriou’06, ChenDeng’06]	FIXP-complete [EtessamiYannakakis’07]
Practical algorithm	Lemke-Howson’64 algorithm	

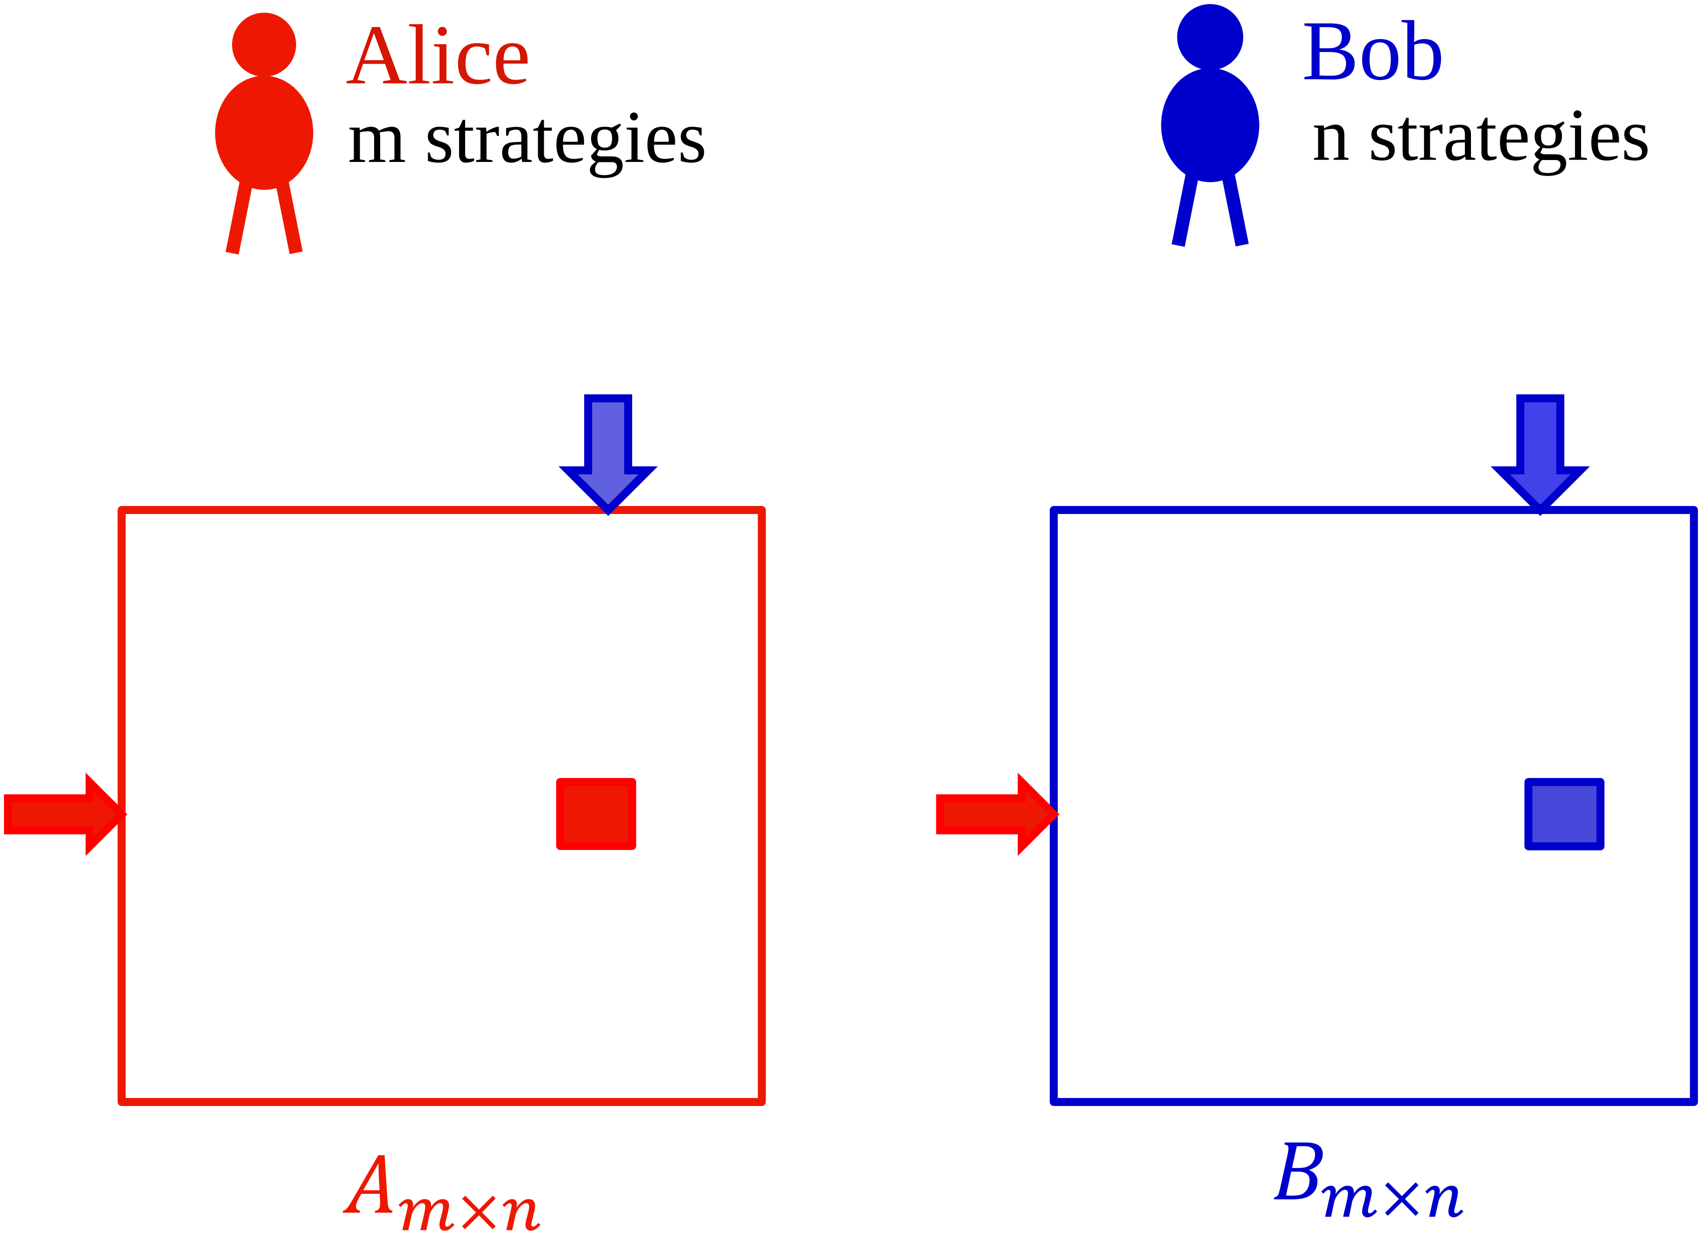

Alice
m strategies

Bob
n strategies

$A_{m\times n}$

$B_{m\times n}$

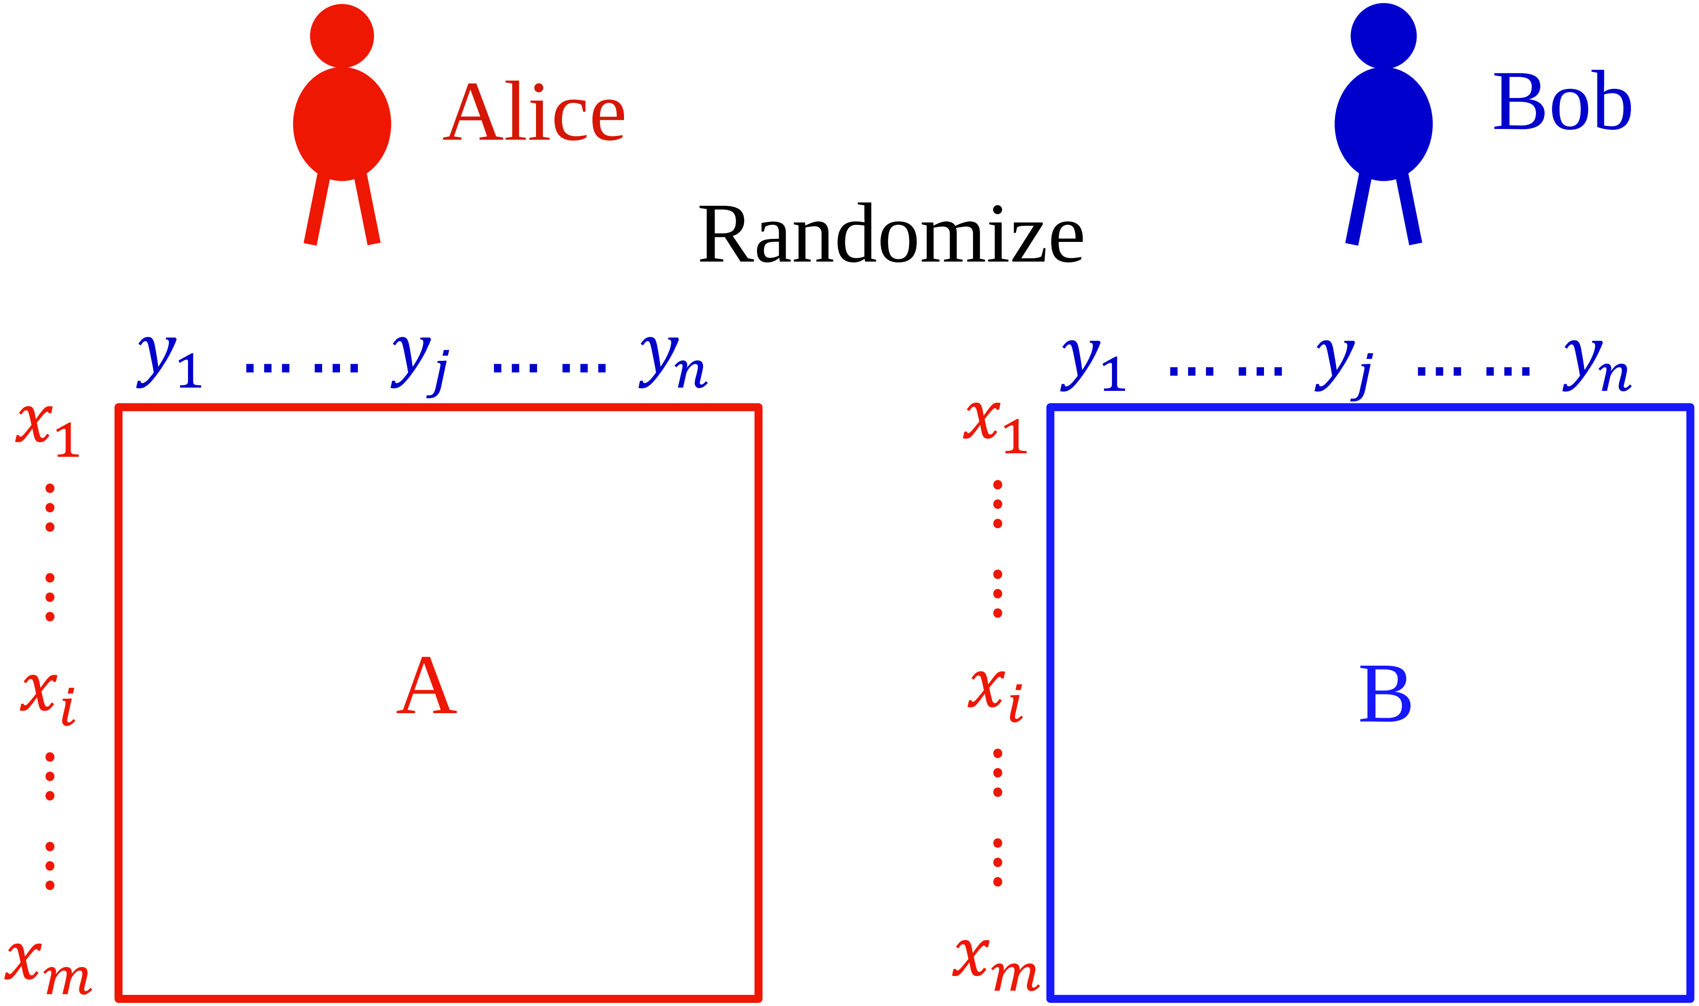

NE: $x^T A y \geq x'^T A y, \forall x'$ $\quad x^T B y \geq x^T A y', \forall y'$

ϵ-NE: $\geq x'^T A y - \epsilon$ $\quad \geq x^T A y' - \epsilon$

Example: Matching Penny

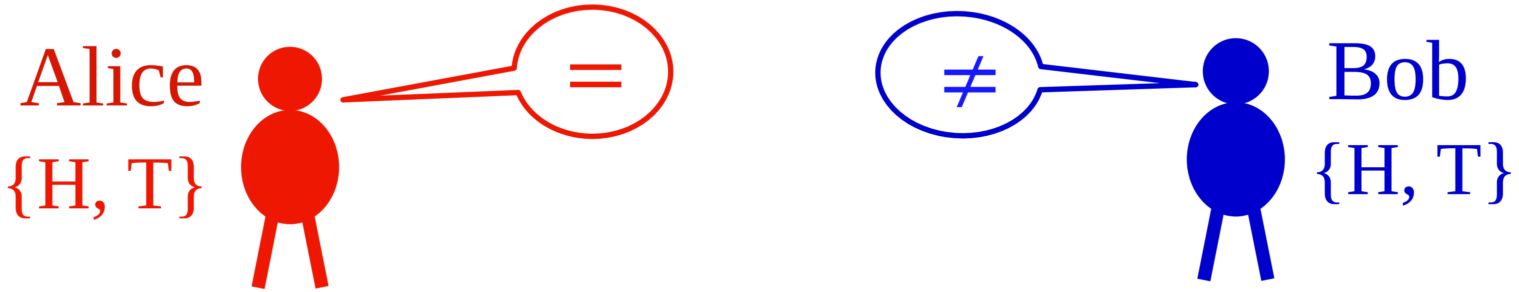

	H	T
H	**1** **-1**	**-1** **1**
T	**-1** **1**	**1** **-1**

Zero-sum:

A + **B** = **0**

von Neumann’28: Min-Max strategies are stable (NE)

Dantzig’51: That’s an LP!

Computational Complexity

- PPAD-complete. Even for win-loose, sparse, and low-rank games [AbbottKaneValiant'05, ChenDengTeng'06, Mehta'14]
- $\frac{1}{poly(n)}$-approximation is PPAD-complete [ChenDengTechng'06]
  - ☐ Smoothed complexity is not in P unless RP = PPAD.
- ϵ-approximation in $O(n^{\epsilon \log n})$ time [LiptonMarkakisMehta'03]
  - ☐ Best assuming exponential-time hypothesis for PPAD [Rubinstein'16]
- Decision versions, e.g., if ∃ more than one NE, NE with max-payoff
  - ☐ NP-complete. No constant approximation assuming ETH for 3-SAT [Gilboa-Zemel'89, Conitzer-Sandholm'08, HazanKrauthgamer'11, BravermanKoWeinstein'15, DeligkasFearnleySavani'16]
- Query complexity …

- i^{th} strategy gives Alice

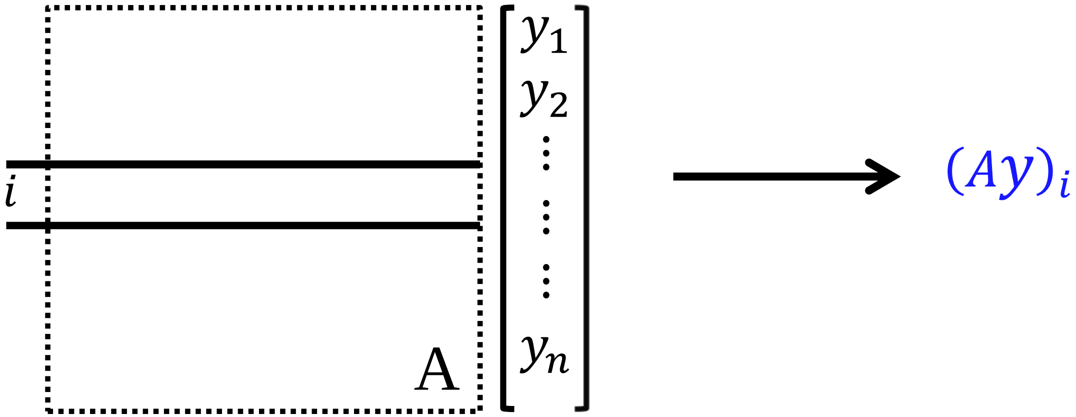

- Max payoff is $\max_i (Ay)_i$

- x achieves max payoff iff

$$\forall k, x_k > 0 \Rightarrow (Ay)_k = \max_i (Ay)_i$$

Given support of (x, y), $\exists$linear feasibility formulation

Efficient Algorithms

- Quasi-PTAS: ϵ-approximation in $O(n^{\epsilon \log n})$ time [LiptonMarkakisMehta'03]
    - ☐ Given NE (x, y), uniform strategy over $O(n^{\log n})$ sample as per (x, y) gives constant approximate NE.

Technique: Bound the search space, enumerate, and check.

Efficient Algorithms

- Quasi-PTAS: ϵ-approximation in $O(n^{\epsilon \log n})$ time [LiptonMarkakisMehta'03]
- Rank of A or B is a constant [JiangGargMehta'11]
  - ☐ If rank(A) is constant, then the row player has polynomialy many *valid* strategies.

Technique: Bound the search space, enumerate, and check.

Efficient Algorithms

- Quasi-PTAS: ϵ-approximation in $O(n^{\epsilon \log n})$ time [LiptonMarkakisMehta'03]
- Rank of A or B is a constant [JiangGargMehta'11]
- FPTAS for constant rank games; rank(A+B) is constant [KannanTheobald'05]
- (A+B) is sparse [Barman'15]

Technique: Bound the search space, enumerate, and check.

- Rank-1 games, i.e., rank(A+B)=1 [AdsulGargSohoniMehta'11]
  - ☐ Parameterized LP + binary search
- Multi-player succinct games …

This Talk

Games, Nash equilibrium, Algorithms, Complexity

Potential Games

- Network-flow, congestion

Extensive form games.

Commitment: Stackleberg equilibrium

- Application: Security games

Repeated games

(sessions 3B and 7B)

Potential Games

Session 7B

Potential function ϕ that captures progress of all the players

$$\phi(\boldsymbol{s}) - \phi(\color{red}{s_i'}\color{black}, \boldsymbol{s}_{-i}) = u_i(\boldsymbol{s}) - u_i(\color{red}{s_i'}\color{black}, \boldsymbol{s}_{-i}) \ \ \forall \text{players } i, \forall \boldsymbol{s}, \forall \color{red}{s_i'}$$

↑ Strategies of all the players

↑ Strategies of all players except i

	H	T
H	1 -	1 1
T	-1 1	1 -1

✗

	L	R
U	1 1	2 2
D	2 2	1 1

✓

$$\phi(s_1, s_2) = \color{red}{A_{s_1 s_2}}$$

Routing (network flow) games

- Directed (road) network given by a graph $G = (V, E)$
- Latency (delay) function on edge e is $l_e: R_+ \rightarrow R_+$, non-decreasing
- A set N of players. Player i wants to go from s_i to t_i
  - ☐ Each player wants to take the route that minimize her total delay.

Example

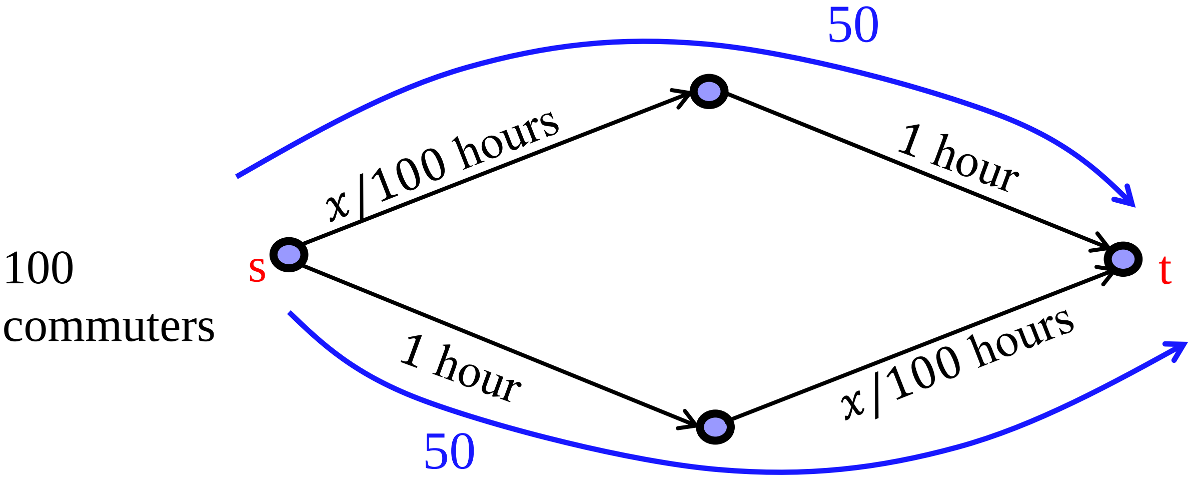

Commute time per person: 1.5 hours

Routing (network flow) games

Example: Braess' Paradox

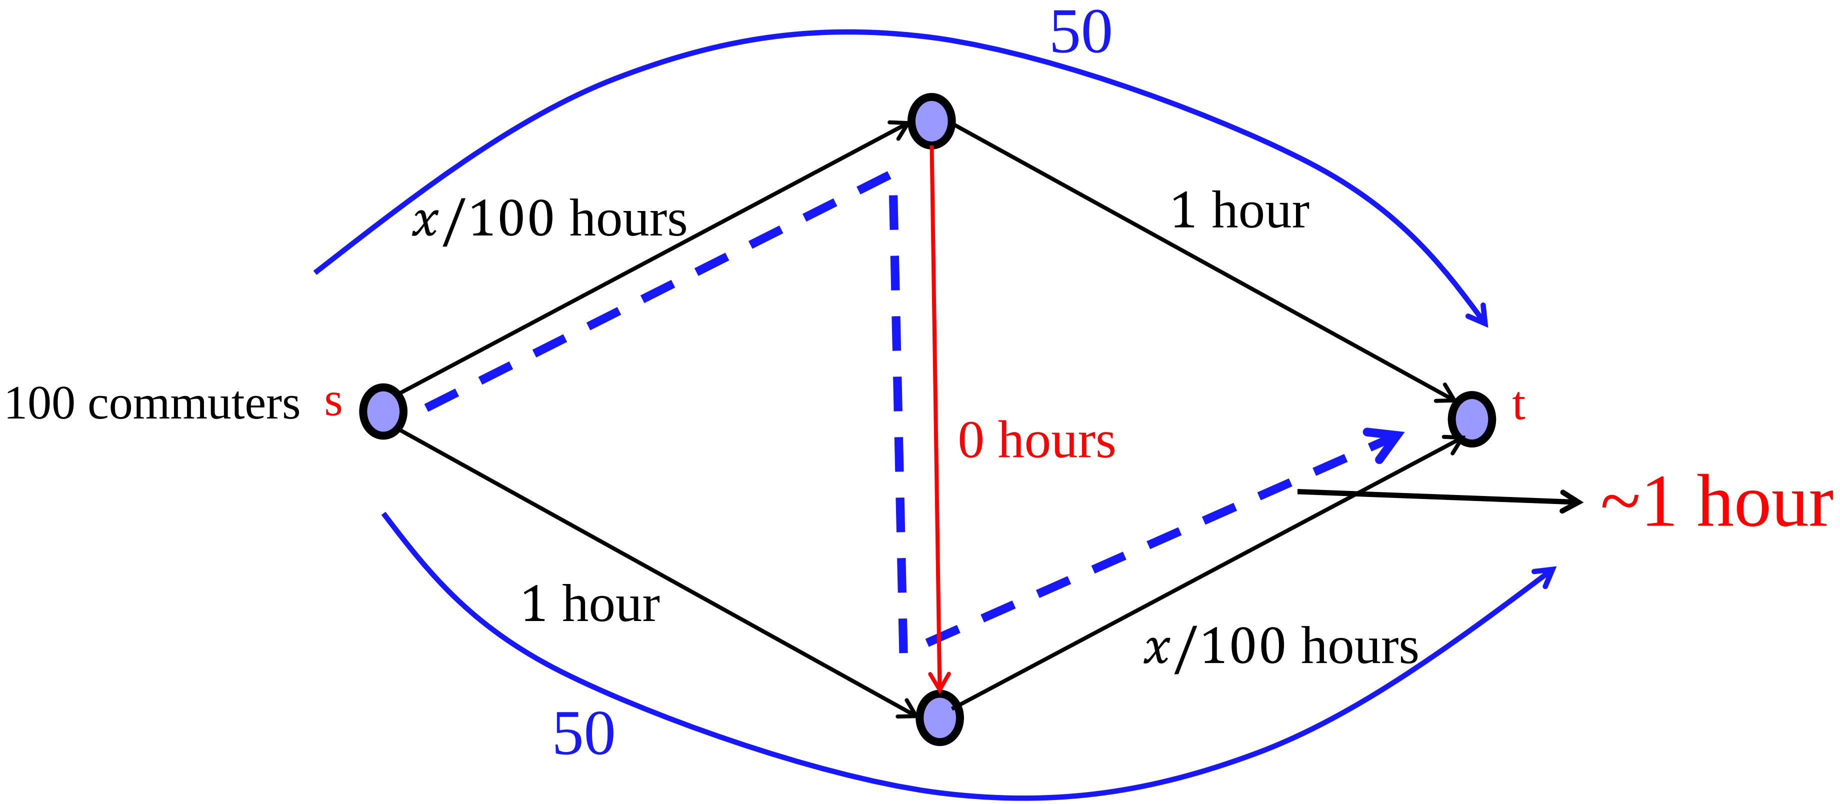

Commute time per person: 1.5 hours

Routing (network flow) games

Example: Braess' Paradox

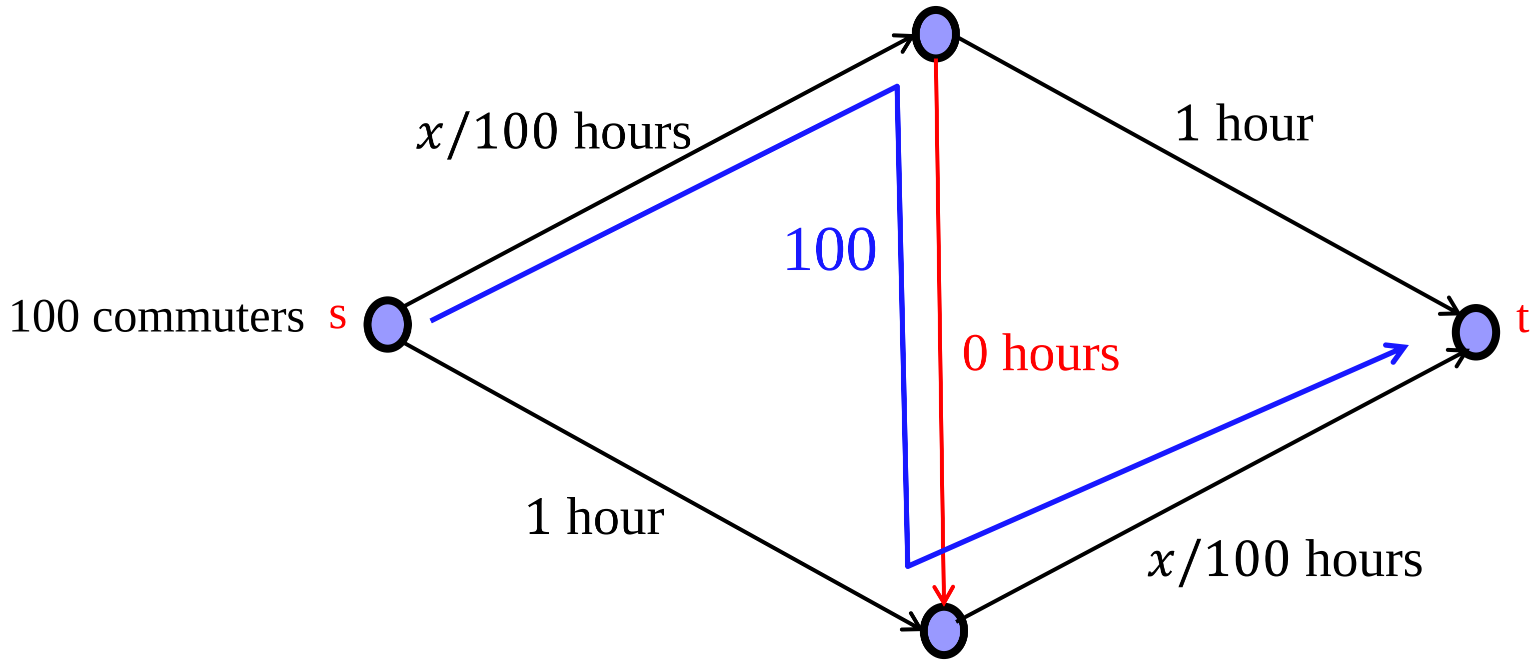

Commute time per person: 2 hours!

Routing games: Potential Function

- $P = (p_1, \dots, p_n)$ be the paths taken by players.
- n_e: players taking edge e as per P.

$$\phi(P) = \sum_{e \in E} \sum_{k=1}^{n_e} l_e(k)$$

$$u_i(P) - u_i(p_i', P_{-i}) = \sum\nolimits_{e \in p_i \setminus p_i'} \boxed{l_e(n_e + 1)} - \sum\nolimits_{e \in p_i' \setminus \mathrm{p}_i} \boxed{l_e(n_e)}$$

$$\left(\sum\nolimits_{k=1}^{n_e+1} l_e(k) - \sum\nolimits_{k=1}^{n_e} l_e(k)\right) \qquad \left(\sum\nolimits_{k=1}^{n_e} l_e(k) - \sum\nolimits_{k=1}^{n_e-1} l_e(k)\right)$$

$$= \phi(P) - \phi(p_i', P_{-i})$$

Congestion Games

Each player chooses some subset from a set of resources, and the cost of each resource depends on the number of other agents who select it.

- N players, R resources.
- Set of actions of player i, $A_i \subseteq 2^R$.
- Cost function for resource r is $l_r\colon \mathbb{N} \rightarrow \mathbb{R}$
- Given an action profile $a = (a_1, \dots, a_N)$, let $n_r = |\{i \mid r \in a_i\}|$
- Cost of player *i* at profile a is $c_i(a) = \sum_{r \in a_i} l_r(n_r)$

- Potential Function: $\phi(a) = \sum_r \sum_{k=1}^{n_r} l_r(k)$

Equivalent to Potential games.

Properties

- Existence of pure NE
  - Strategy profile with the best potential.
- Sequential best response always converges to a pure NE
  - Because the potential improves in every round.
- Finding pure NE is PLS-complete
  - Polynomial Local Search: Given a DAG, find a sink
- Finding mixed NE is in CLS
  - Continuous Local Search: Both PPAD and PLS like

This Talk

- Games, Nash equilibrium, Algorithms, Complexity
- Potential Games
  - ☐ Network-flow, congestion
- Extensive form games
- Commitment: Stackleberg equilibrium
  - ☐ Application: Security games
- Repeated games

(sessions 3B and 7B)

Following slides curtesy Vince Conitzer

Extensive-form Game

Session 3B

- Players move one after another
  - Chess, Poker, etc.
  - Tree representation.

Strategy of a player:
What to play at each of its node.

	I	O
F	-1, 1	2, 0
A	1, 1	2, 0

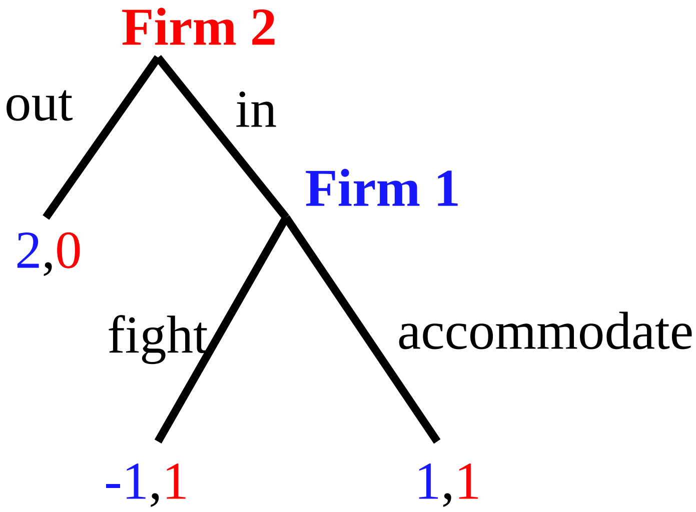

Entry game

A poker-like game

- Player 1 gets a card (King is a winning card, Jack a losing card)
- Player 1 decides to raise (add one to the pot) or check
- Player 2 decides to call (match) or fold (P1 wins)
- If player 2 called, player 1's card determines pot winner

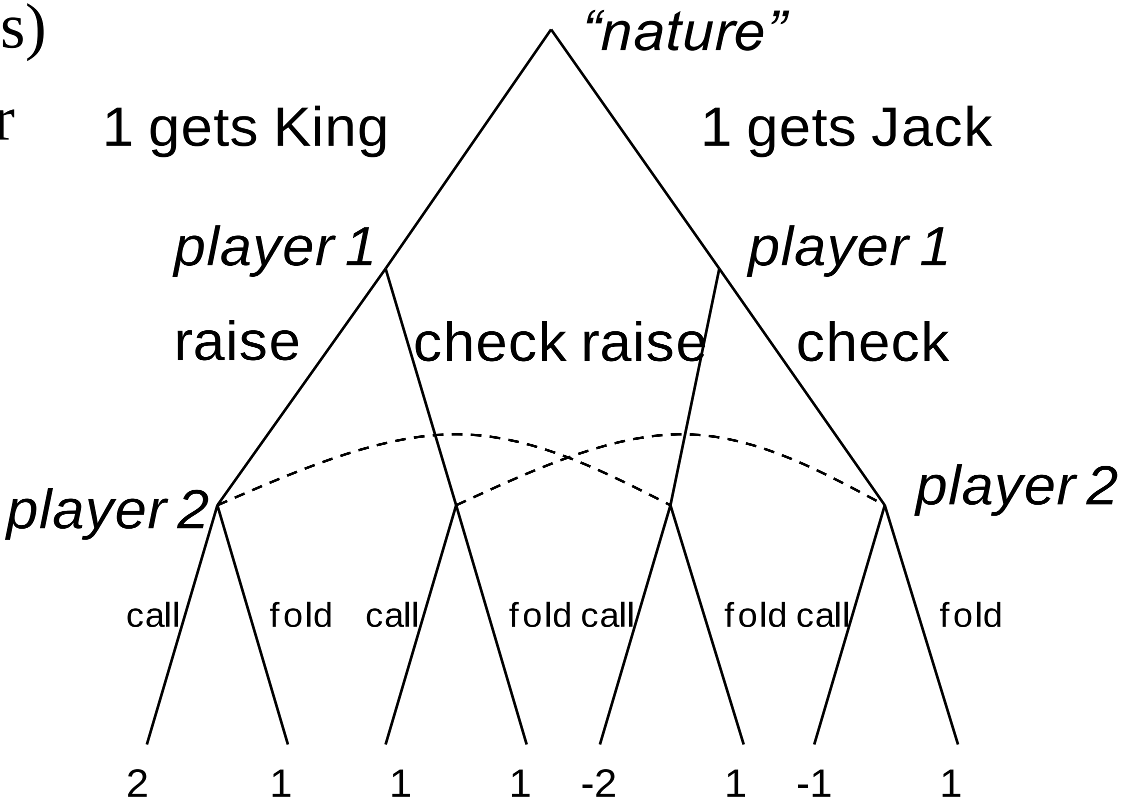

Poker-like game in normal form

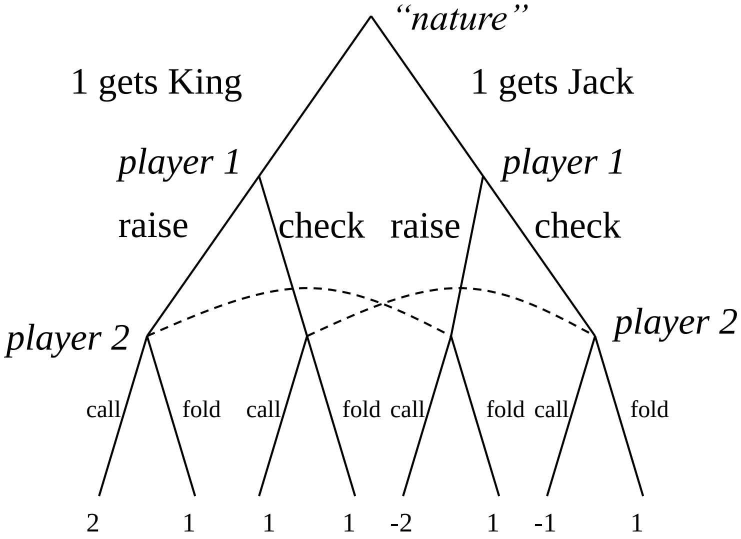

	cc	cf	fc	ff
rr	0, 0	0, 0	1, -1	1, -1
rc	.5, -.5	1.5, -1.5	0, 0	1, -1
cr	-.5, .5	-.5, .5	1, -1	1, -1
cc	0, 0	1, -1	0, 0	1, -1

Can be exponentially big!

Sub-Game Perfect Equilibrium

- Every sub-tree is at equilibrium
- Computation when perfect information (no nature/chance move): Backward induction

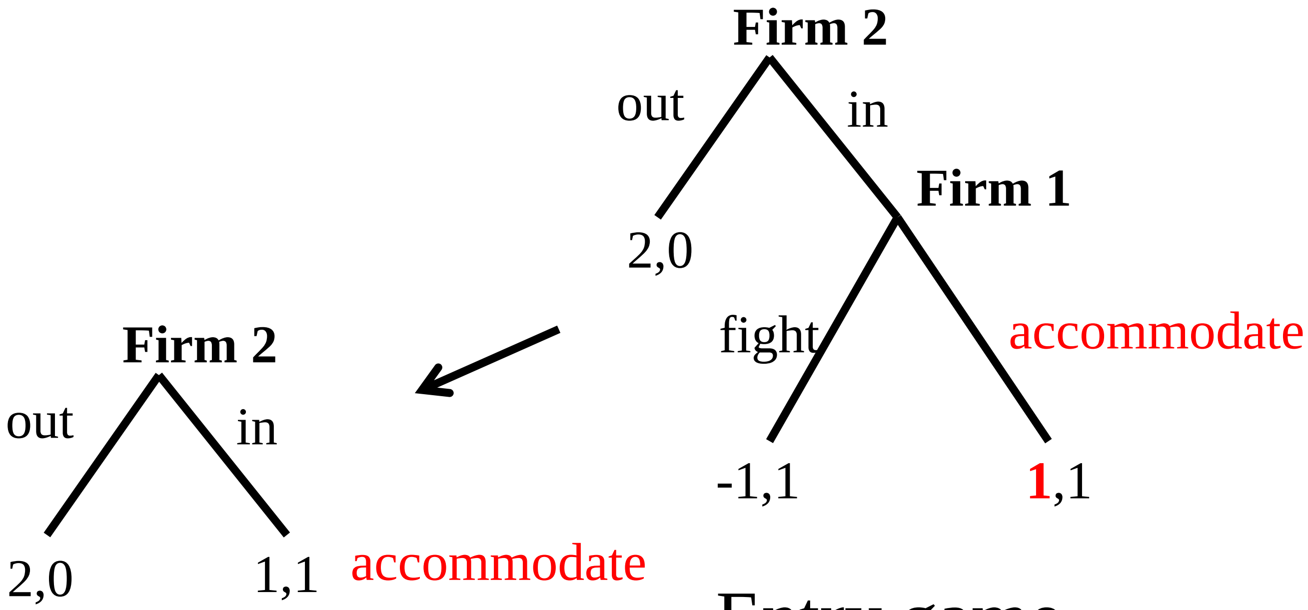

Entry game

Sub-Game Perfect Equilibrium

- Every sub-tree is at equilibrium
- Computation when perfect information (no nature/chance move): Backward induction

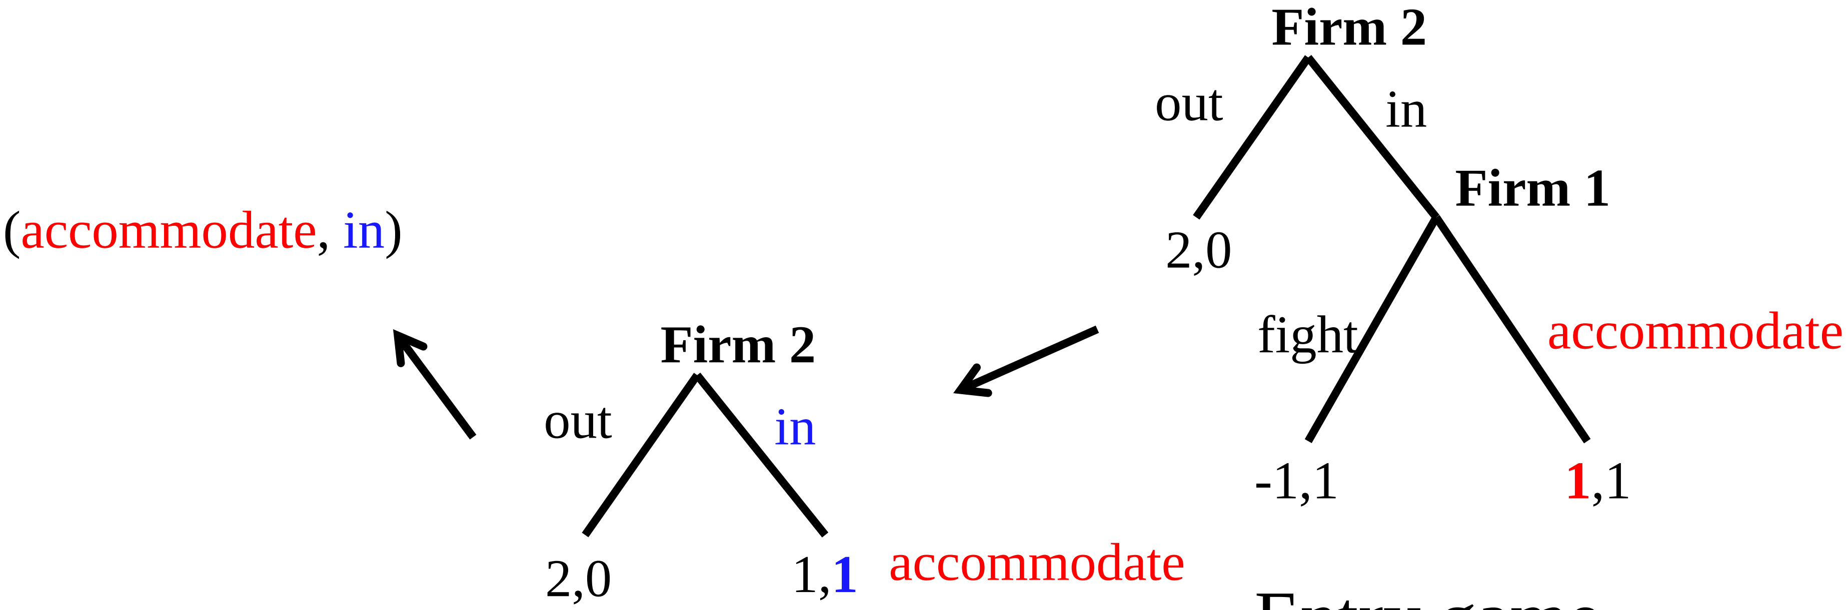

Entry game

This Talk

Games, Nash equilibrium, Algorithms, Complexity

Potential Games

- ☐ Network-flow, congestion

Extensive form games

Commitment: Stackleberg equilibrium

- ☐ Application: Security games

Repeated games

(sessions 3B and 7B)

Commitment

Session 7B

Unique Nash equilibrium

1, 1	3, 0
0, 0	2, 1

von Stackelberg

- Suppose the game is played as follows:
  - Player 1 commits to playing one of the rows,
  - Player 2 observes the commitment and then chooses a column
- Optimal strategy for player 1: commit to Down

Commitment: an extensive-form game

For the case of committing to a pure strategy:

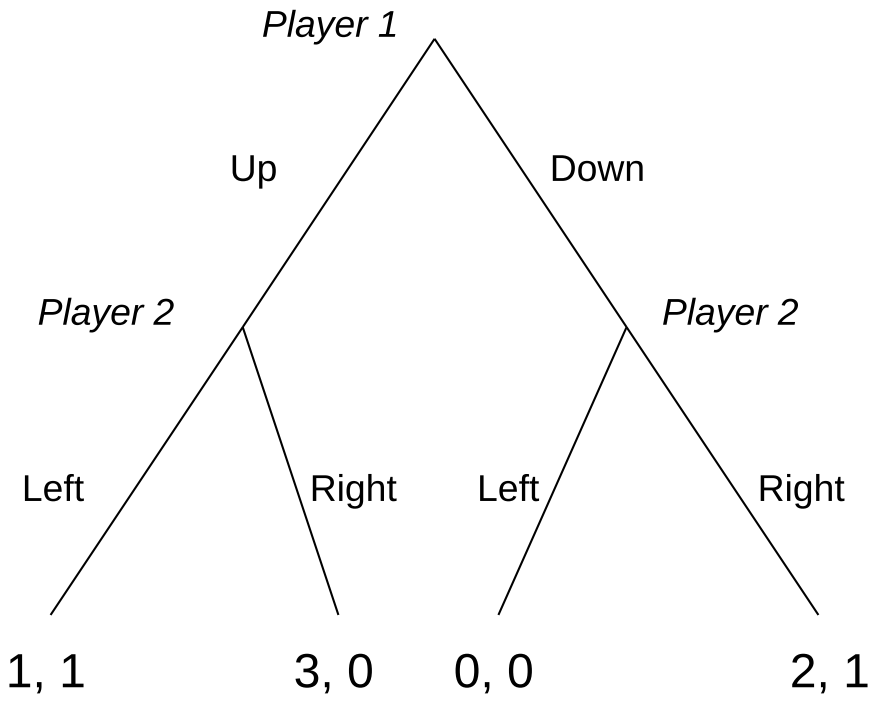

Commitment to mixed strategies

	0	1
.49	1, 1	3, 0
.51	0, 0	2, 1

Also called a Stackelberg (mixed) strategy

For the follower, pure best response always exist

Commitment: an extensive-form game

- … for the case of committing to a mixed strategy:

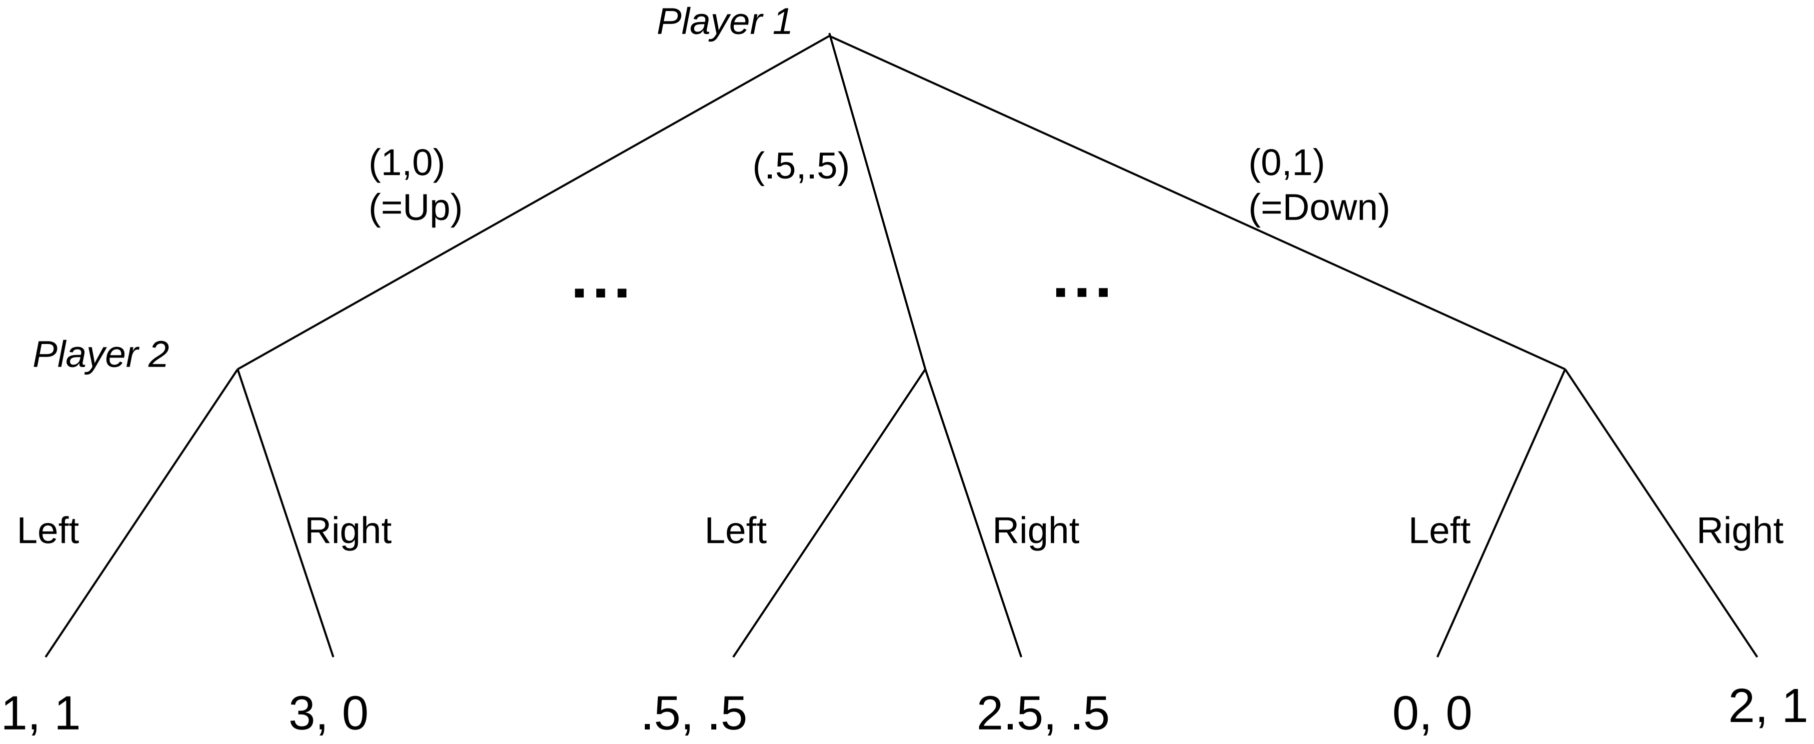

- Economist: Just an extensive-form game, nothing new here
- Computer scientist: Infinite-size game! Representation matters

Computing the optimal mixed strategy to commit to [Conitzer & Sandholm EC'06]

- Alice is a leader.
- Separate LP for every column $j^* \in S_2$ (actions of the column player

maximize $\sum_i x_i A_{ij^*}$ Row utility

subject to $\forall j,\ \ (x^T B)_{j^*} \geq (x^T B)_j$ j^* Column optimality

$\sum_i x_i = 1$ distributional constraint

Pick the one that gives max utility.

On the game we saw before

x_1	1, 1	3, 0
x_2	0, 0	2, 1

$$\textit{maximize } 1x_1 + 0x_2$$

$$\textit{subject to}$$

$$1\,x_1 + 0\,x_2 \geq 0\,x_1 + 1\,x_2$$

$$x_1 + x_2 = 1$$

$$x_1 \geq 0,\, x_2 \geq 0$$

$$\textit{maximize } 3\,x_1 + 2\,x_2$$

$$\textit{subject to}$$

$$0\,x_1 + 1\,x_2 \geq 1\,x_1 + 0\,x_2$$

$$x_1 + x_2 = 1$$

$$x_1 \geq 0,\, x_2 \geq 0$$

Generalizing beyond zero-sum games

Minimax, Nash, Stackelberg all agree in zero-sum games

0, 0	-1, 1
-1, 1	0, 0

minimax strategies

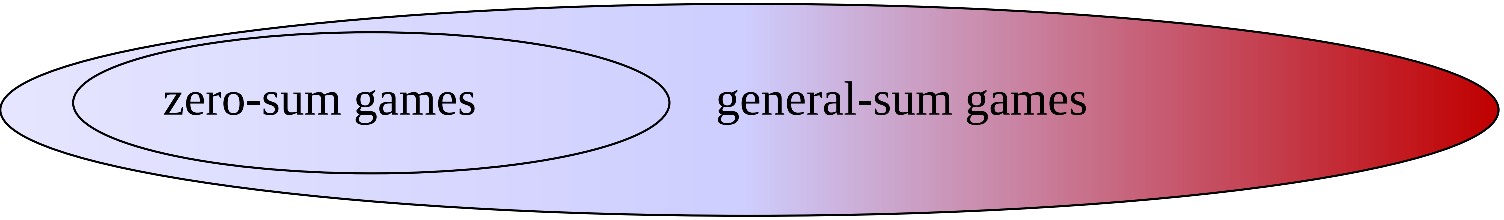

Nash equilibrium

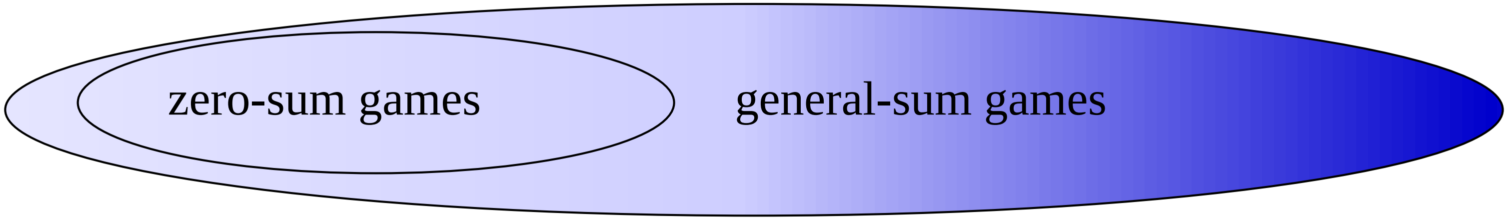

Stackelberg mixed strategies

Other nice properties of commitment to mixed strategies

- No equilibrium selection problem

0, 0	-1, 1
1, -1	-5, -5

- Leader’s payoff at least as good as any Nash eq. or even correlated eq. (von Stengel & Zamir [GEB ‘10])

≥

Applications

Session 7B

Security Games

- Players: Defender team, Attacker team
- Defender's goal: Design a security strategy such that even if attacker has some idea, it can not gain much.
  - Defender is a natural leader, and attacker the follower.

- LAX security, NYC Coast guards, Poaching, etc.
  [Teamcore, USC]

This Talk

- Games, Nash equilibrium, Algorithms, Complexity
- Potential Games
  - Network-flow, congestion
- Extensive form games
- Commitment: Stackleberg equilibrium
  - Application: Security games
- Repeated games

(sessions 3B and 7B)

Repeated games

Session 3B
Talk 3

- In a (typical) repeated game,
  - players play a normal-form game (aka. the stage game),
  - then they see what happened (and get the utilities),
  - then they play again,
  - etc.
- Can be repeated finitely or infinitely many times
- Really, an extensive form game
  - Would like to find subgame-perfect equilibria
- One subgame-perfect equilibrium: keep repeating some Nash equilibrium of the stage game
- But are there other equilibria?

Finitely repeated Prisoner’s Dilemma

- Two players play the Prisoner’s Dilemma k times

	cooperate	defect
cooperate	2, 2	0, 3
defect	3, 0	1, 1

- In the last round, it is dominant to defect
- Hence, in the second-to-last round, there is no way to influence what will happen
- So, it is optimal to defect in this round as well
- Etc.
- So the only equilibrium is to always defect

Infinitely repeated games

- First problem: are we just going to add up the utilities over infinitely many rounds?
  - Everyone gets infinity!
- (Limit of) average payoff: $\lim_{n \to \infty} \Sigma_{1 \leq t \leq n} u(t)/n$
  - Limit may not exist…
- Discounted payoff: $\Sigma_t \delta^t u(t)$ for some $\delta < 1$

Infinitely repeated Prisoner's Dilemma

	cooperate	defect
cooperate	2, 2	0, 3
defect	3, 0	1, 1

- Tit-for-tat strategy:
  - Cooperate the first round,
  - In every later round, do the same thing as the other player did in the previous round
- Is both players playing this a Nash/subgame-perfect equilibrium? Does it depend on δ?
- Trigger strategy:
  - Cooperate as long as everyone cooperates
  - Once a player defects, defect forever
- Is both players playing this a subgame-perfect equilibrium?
- What about one player playing tit-for-tat and the other playing trigger?

THANK YOU