Introduction to Game Theory

Lirong Xia

Rensselaer

Voting: manipulation

(ties are broken alphabetically)

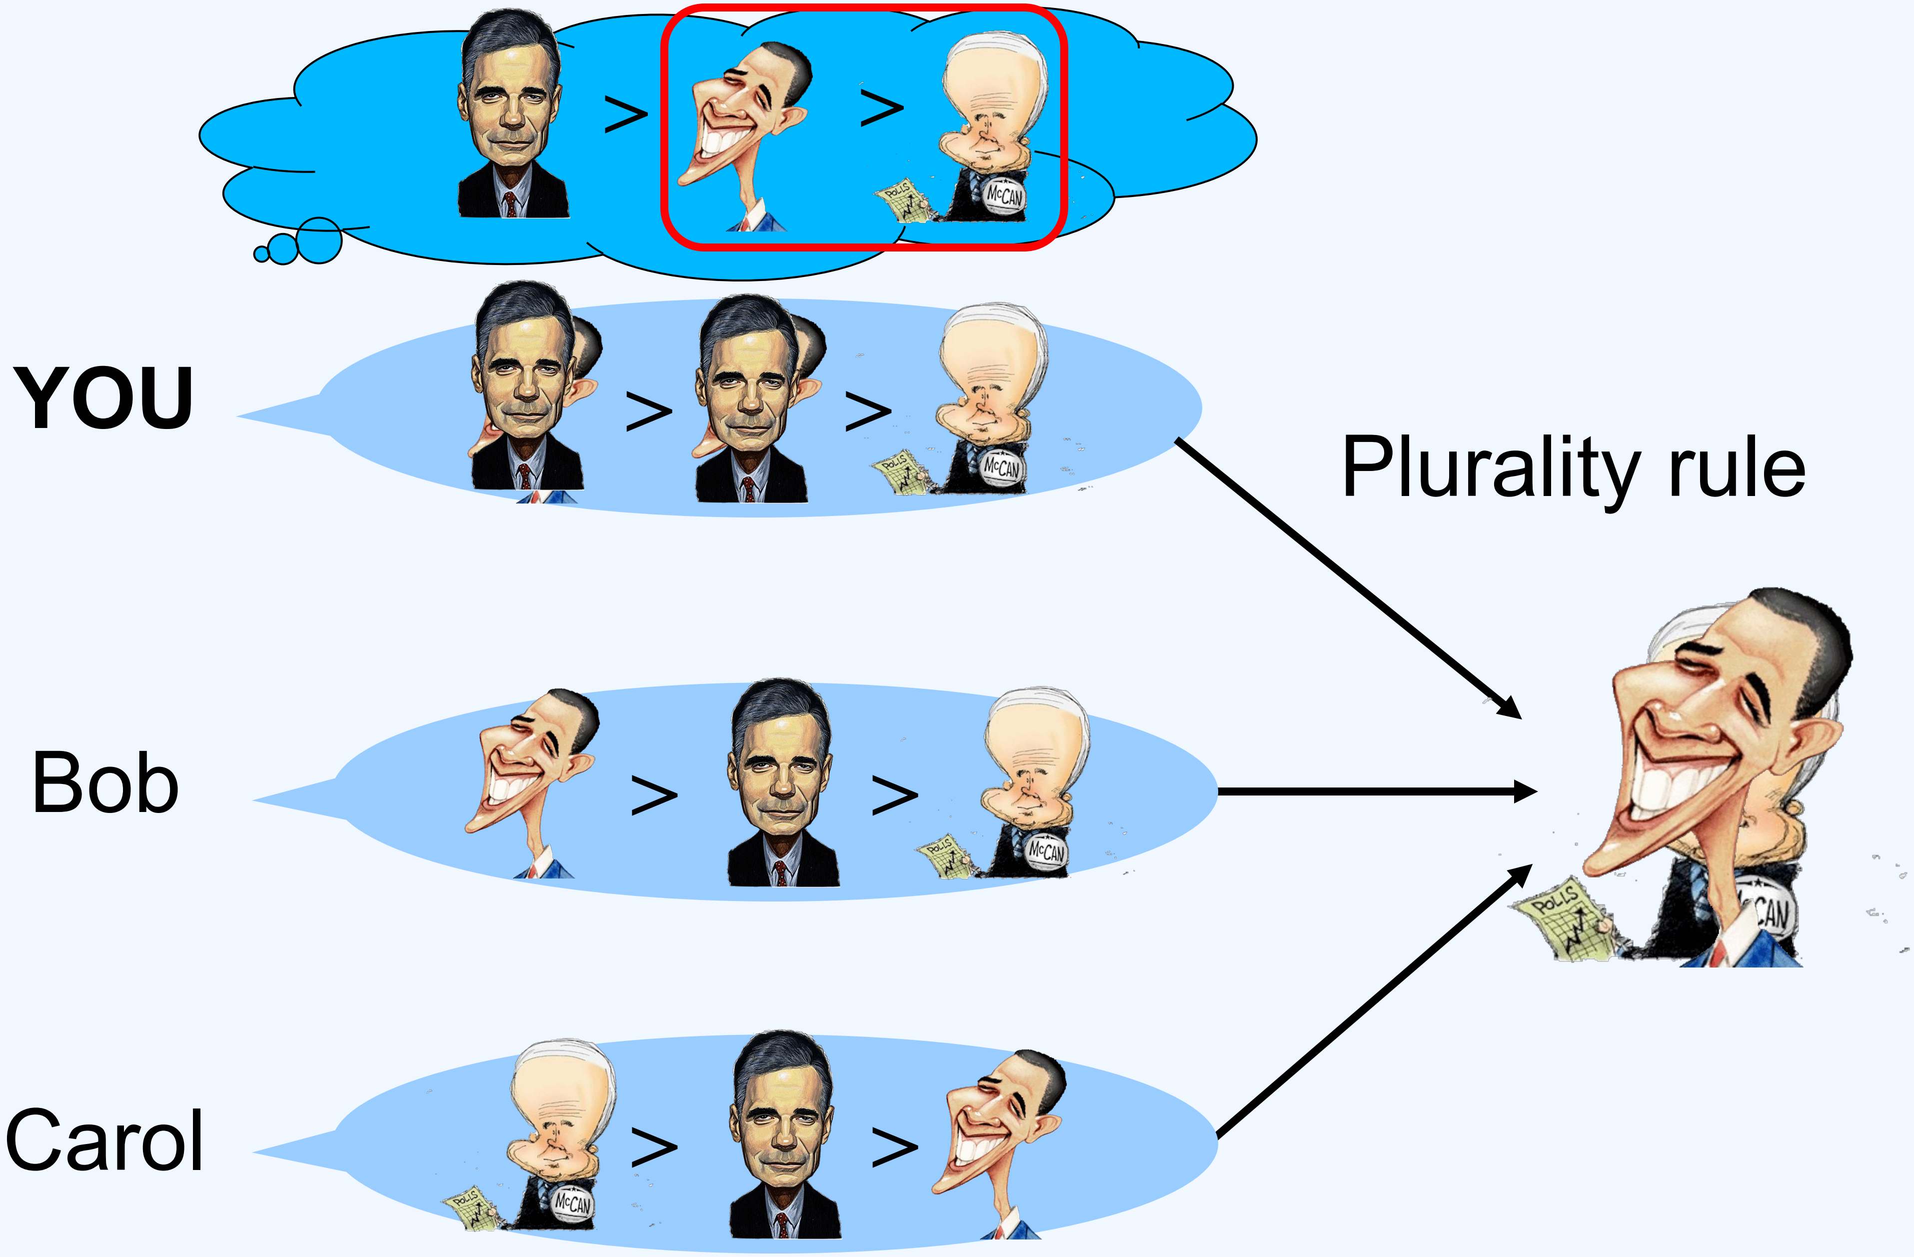

What if everyone is incentivized to lie?

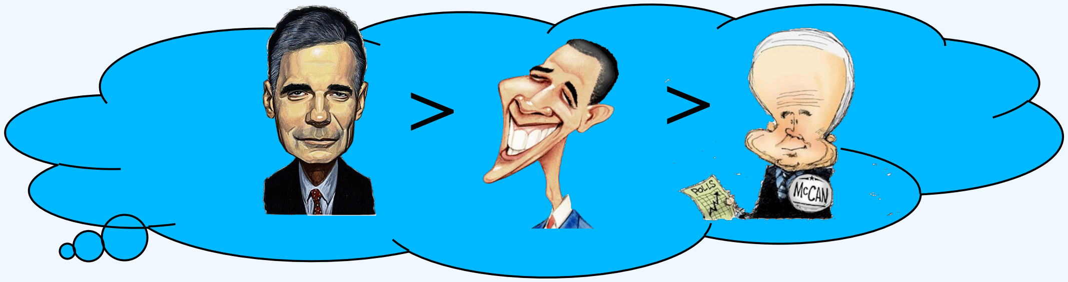

YOU

Plurality rule

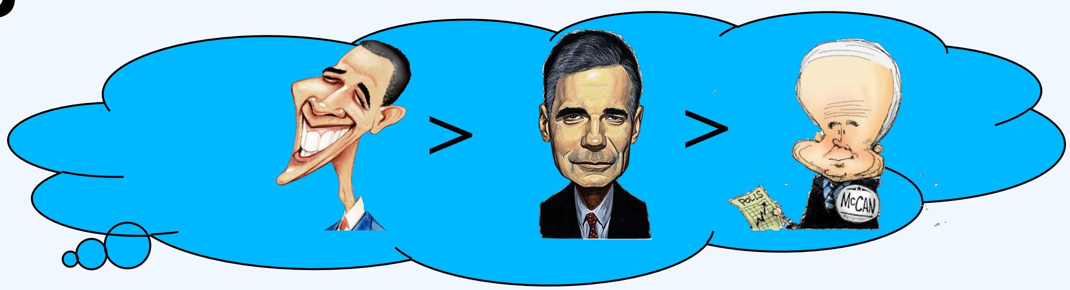

Bob

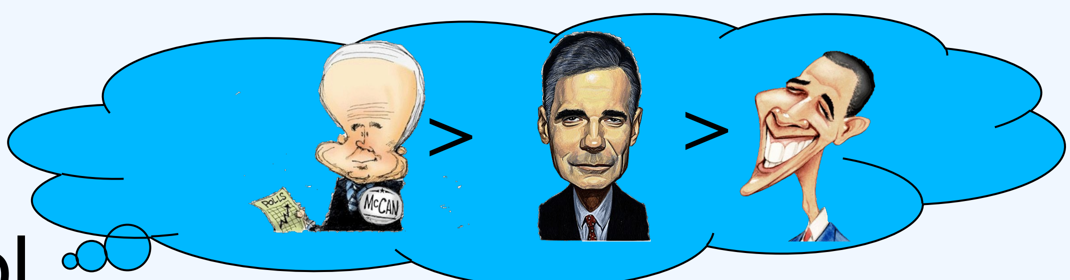

Carol

History of Game Theory

- *On the Theory of Games of Strategy.* Mathematische Annalen, 1928.
  - John von Neumann

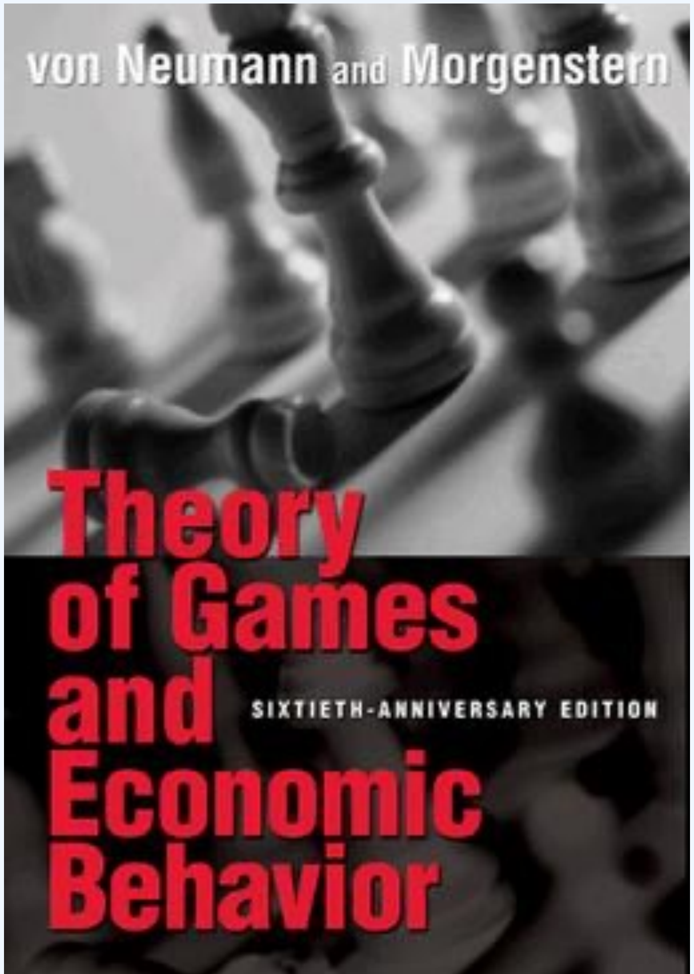

Nobel Prize Winners

- 1994:
  - Nash (Nash equilibrium)
  - Selten (Subgame pefect equilibrium)
  - Harsanyi (Bayesian games)

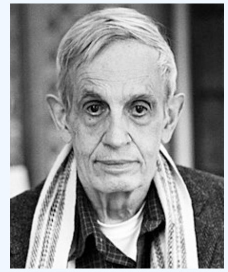

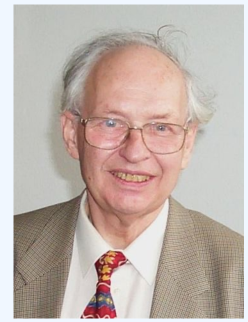

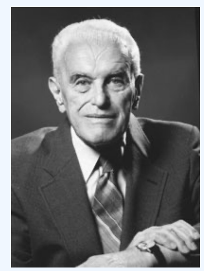

- 2005
  - Schelling (evolutionary game theory)
  - Aumann (correlated equilibrium)

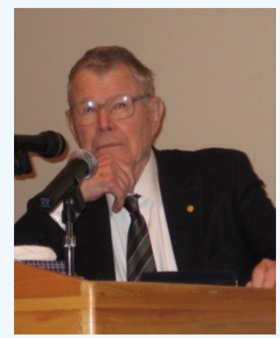

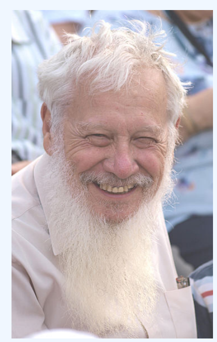

- 2014
  - Jean Tirole

A game of two prisoners

Column player

Row player

	Cooperate	**Defect**
Cooperate	(-1 , -1)	(-3 , 0)
Defect	(0 , -3)	(-2 , -2)

- Players: 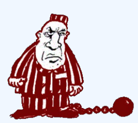 
- Strategies: { Cooperate, Defect }
- Outcomes: {(-2 , -2), (-3 , 0), (0 , -3), (-1 , -1)}
- Preferences: self-interested 0 > -1 > -2 > -3
  - 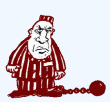 : (0 , -3) > (-1 , -1) > (-2 , -2) > (-3 , 0)
  -  : (-3 , 0) > (-1 , -1) > (-2 , -2) > (0 , -3)
- Mechanism: the table

Formal Definition of a Game

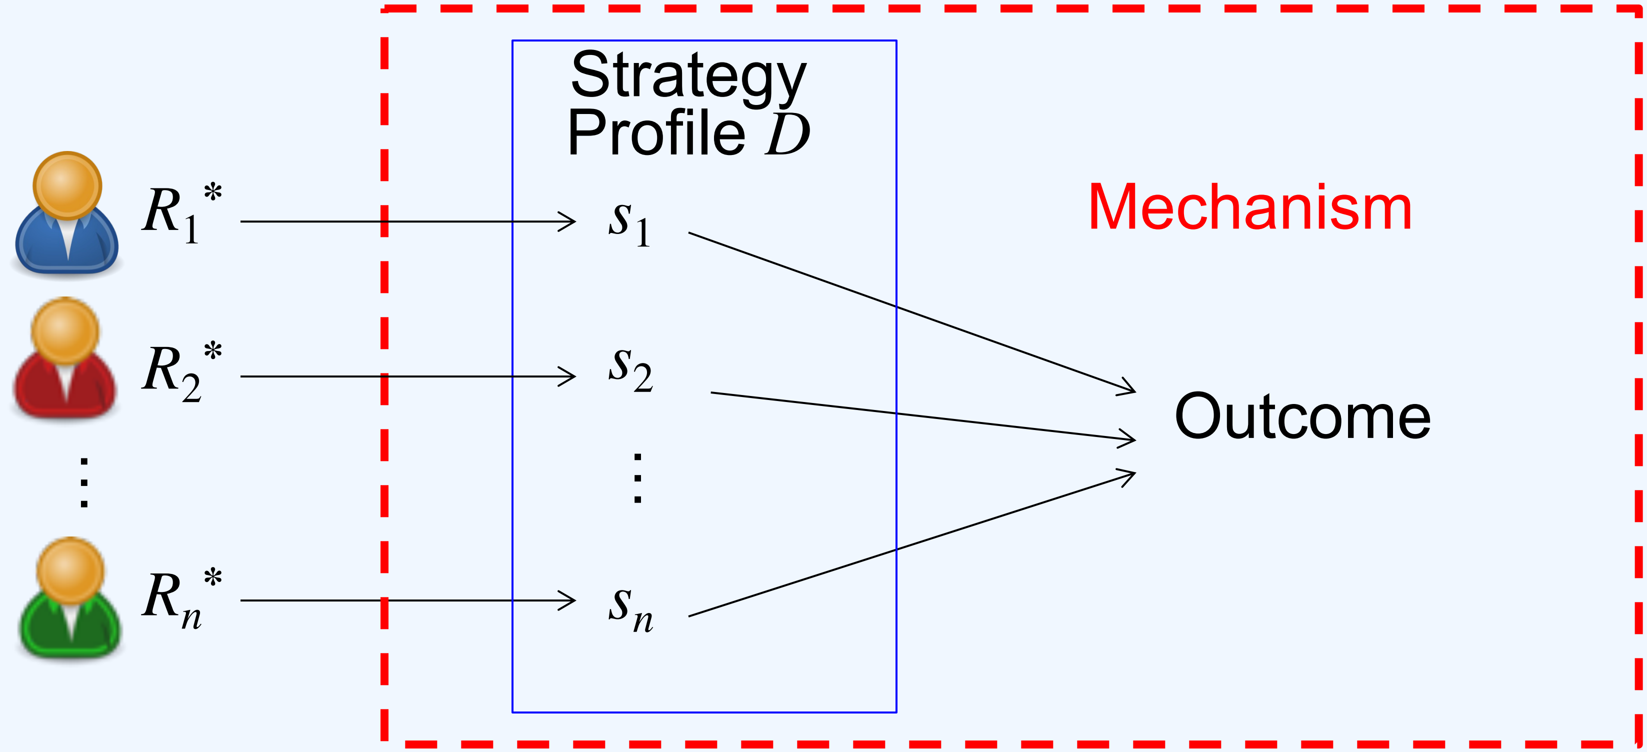

- Players: $N=\{1,\ldots,n\}$
- Strategies (actions):
  - S_j for agent j, $s_j \in S_j$
  - $(s_1,\ldots,s_n)$ is called a strategy profile.
- Outcomes: O
- Mechanism $f : \Pi_j S_j \rightarrow O$
- Preferences: total preorders (full rankings with ties) over O
  - often represented by a utility function $u_i : O \rightarrow R$

A game of plurality elections

YOU

Bob

Carol

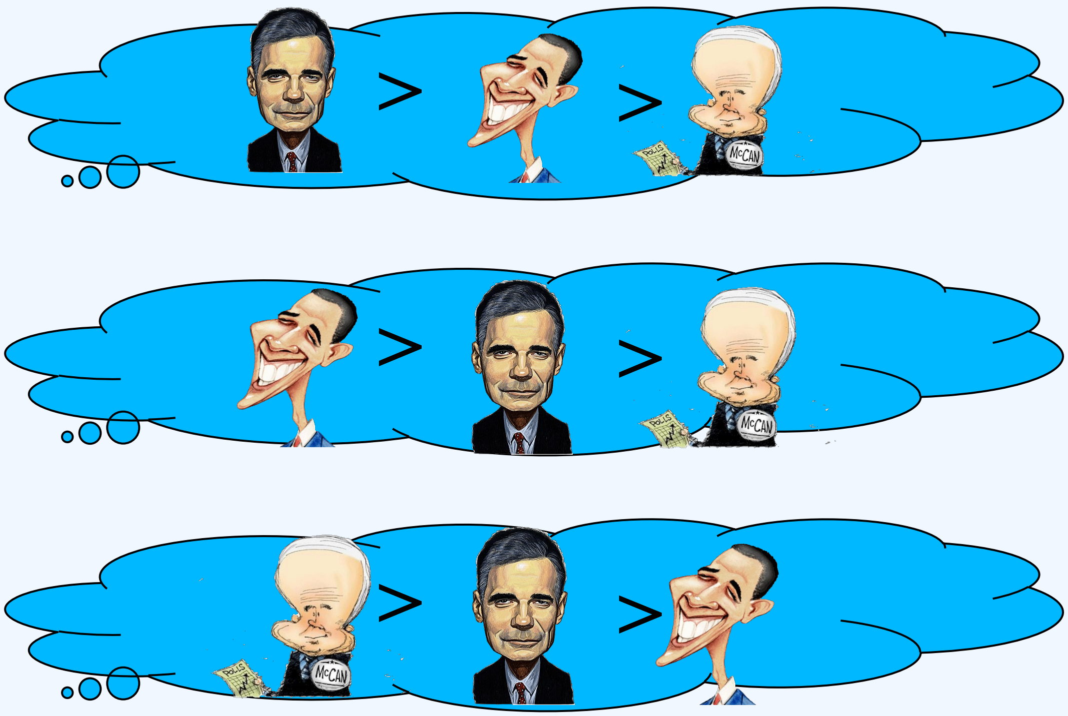

Plurality rule

- Players: { YOU, Bob, Carol }
- Outcomes: O = {  }
- Strategies: S_j = Rankings(O)
- Preferences: See above
- Mechanism: the plurality rule

Solving the game

- Suppose
  - every player wants to make the outcome as preferable (to her) as possible by controlling her own strategy (but not the other players')
- What is the outcome?
  - No one knows for sure
  - A "stable" situation seems reasonable
- A Nash Equilibrium (NE) is a strategy profile $(s_1,\ldots,s_n)$ such that
  - For every player j and every $s_j' \in S_j$,

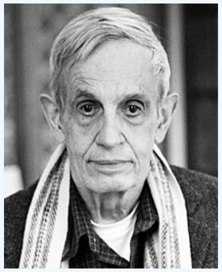

$$f(s_j, s_{-j}) \geq_j f(s_j', s_{-j}) \text{ or equivalently } u_j(s_j, s_{-j}) \geq u_j(s_j', s_{-j})$$

  - $s_{-j} = (s_1,\ldots,s_{j-1}, s_{j+1},\ldots,s_n)$
  - no single player can be better off by unilateral deviation

Prisoner’s dilemma

Column player

Row player

	Cooperate	**Defect**
Cooperate	(-1 , -1)	(-3 , 0)
Defect	(0 , -3)	(-2 , -2)

The Game of Chicken

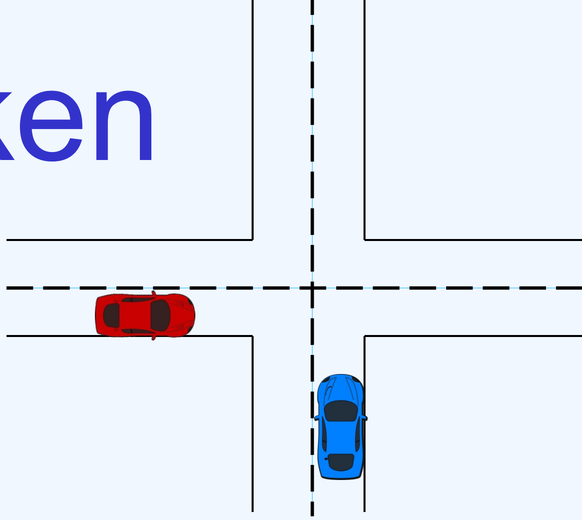

- Two drivers arrives at a cross road
  - each can either (D)air or (C)hicken out
  - If both choose D, then crash.
  - If one chooses C and the other chooses D, the latter "wins".
  - If both choose C, both are survived

Column player

Row player

	Dare	Chicken
Dare	(0 , 0)	(7 , 2)
Chicken	(2 , 7)	(6 , 6)

NE

A beautiful mind

- “If everyone competes for the blond, we block each other and no one gets her. So then we all go for her friends. But they give us the cold shoulder, because no one likes to be second choice. Again, no winner. But what if none of us go for the blond. We don’t get in each other’s way, we don’t insult the other girls. That’s the only way we win. That’s the only way we all get [a girl.]”

A beautiful mind: the bar game

Hansen Column player

Nash Row player

	Blond	Another girl
Blond	(0 , 0)	(5 , 1)
Another girl	(1 , 5)	(2 , 2)

- Players: { Nash, Hansen }
- Strategies: { Blond, another girl }
- Outcomes: {(0 , 0), (5 , 1), (1 , 5), (2 , 2)}
- Preferences: self-interested
- Mechanism: the table

Does an NE always exists?

- Not always (matching pennis game)

Column player

Row player

	H	T
H	(-1 , 1)	(1 , -1)
T	(1 , -1)	(-1 , 1)

- But an NE exists when every player has a dominant strategy
  - s_j is a dominant strategy for player j, if for every $s_j' \in S_j$,
    1. for every s_{-j}, $f(s_j, s_{-j}) \geq_j f(s_j', s_{-j})$
    2. the preference is strict for some s_{-j}

Dominant-strategy NE

- For player j, strategy s_j dominates strategy s_j', if
  1. for every s_{-j}, $u_j(s_j, s_{-j}) \geq u_j(s_j', s_{-j})$
  2. the preference is strict for some s_{-j}
  3. strict dominance: inequality is strict for every s_{-j}
- Recall that an NE exists when every player has a dominant strategy s_j, if
  - s_j dominates other strategies of the same agent
- A dominant-strategy NE (DSNE) is an NE where
  - every player takes a dominant strategy
  - may not exists
  - if strict DSNE exists, then it is the unique NE

Prisoner's dilemma

Column player

Row player

	Cooperate	Defect
Cooperate	(-1 , -1)	(-3 , 0)
Defect	(0 , -3)	(-2 , -2)

Defect is the dominant strategy for both players

Rock Paper Scissors

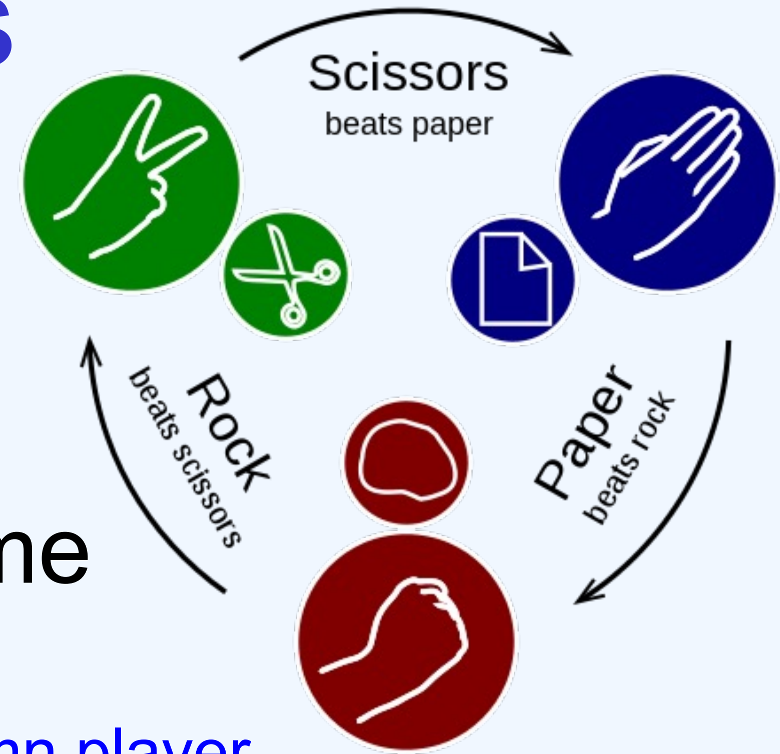

- Actions: {R, P, S}
- Two-player zero sum game

No pure NE

Column player

Row player

	R	P	S
R	(0 , 0)	(-1 , 1)	(1 , -1)
P	(1 , -1)	(0 , 0)	(-1 , 1)
S	(-1 , 1)	(1 , -1)	(0 , 0)

Rock Paper Scissors: Lirong vs. young Daughter

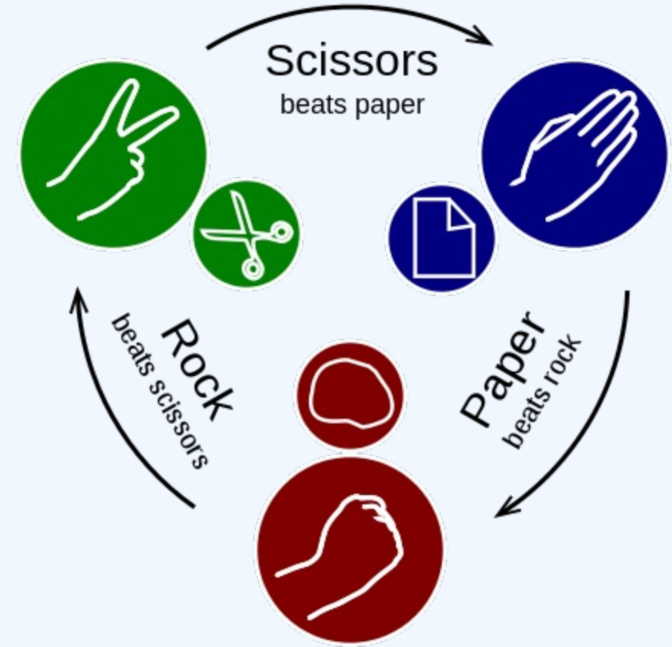

- Actions
  - Lirong: {R, P, S}
  - Daughter: {mini R, mini P}
- Two-player zero sum game

No pure NE

Daughter

Lirong

	mini R	mini P
R	(0 , 0)	(-1 , 1)
P	(1 , -1)	(0 , 0)
S	(-1 , 1)	(1 , -1)

Computing NE: Iterated Elimination

➢ Eliminate dominated strategies sequentially

Column player

Row player

	L	M	R
U	(1 , 0)	(1 , 2)	(0 , 1)
D	(0 , 3)	(0 , 1)	(2 , 0)

Normal form games

- Given pure strategies: S_j for agent j

Normal form games

- Players: $N=\{1,\ldots,n\}$
- Strategies: lotteries (distributions) over S_j
  - $L_j \in \text{Lot}(S_j)$ is called a mixed strategy
  - $(L_1,\ldots, L_n)$ is a mixed-strategy profile
- Outcomes: $\Pi_j \text{ Lot}(S_j)$
- Mechanism: $f(L_1,\ldots,L_n) = p$
  - $p(s_1,\ldots,s_n) = \Pi_j L_j(s_j)$
- Preferences:
  - Soon

Column player

Row player

	L	R
U	(0 , 1)	(1 , 0)
D	(1 , 0)	(0 , 1)

Preferences over lotteries

- Option 1 vs. Option 2
  - Option 1: $0@50%+$30@50%
  - Option 2: $5 for sure
- Option 3 vs. Option 4
  - Option 3: $0@50%+$30M@50%
  - Option 4: $5M for sure

Lotteries

- There are m objects. Obj=$\{o_1,\ldots,o_m\}$
- Lot(Obj): all lotteries (distributions) over Obj
- In general, an agent's preferences can be modeled by a preorder (ranking with ties) over Lot(Obj)
  - But there are infinitely many outcomes

Utility theory

- Utility function: u: Obj $\rightarrow \mathbb{R}$

➢ For any $p \in$ Lot(Obj)

  - $u(p) = \sum_{o \in \text{Obj}} p(o)u(o)$

➢ u represents a total preorder over Lot(Obj)

  - $p_1 > p_2$ if and only if $u(p_1) > u(p_2)$

Example

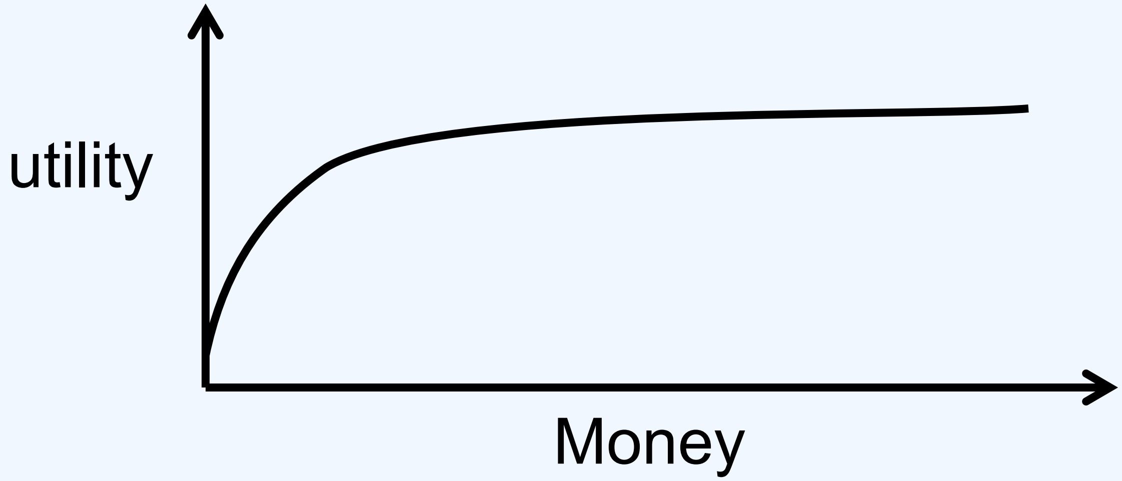

Money	0	5	30	5M	30M
Utility	1	3	10	100	150

- u(Option 1) = $u(0) \times 50\% + u(30) \times 50\% = 5.5$
- u(Option 2) = $u(5) \times 100\% = 3$
- u(Option 3) = $u(0) \times 50\% + u(30\text{M}) \times 50\% = 75.5$
- u(Option 4) = $u(5\text{M}) \times 100\% = 100$

Normal form games

- Pure strategies: S_j for agent j
- Players: $N=\{1,\ldots,n\}$
- (Mixed) Strategies: lotteries (distributions) over S_j
  - $L_j \in \mathrm{Lot}(S_j)$ is called a mixed strategy
  - $(L_1,\ldots, L_n)$ is a mixed-strategy profile
- Outcomes: Π_j Lot(S_j)
- Mechanism: $f(L_1,\ldots,L_n) = p$, such that
  - $p(s_1,\ldots,s_n) = \Pi_j L_j(s_j)$
- Preferences: represented by utility functions $u_1,\ldots,u_n$

Mixed-strategy NE

- Mixed-strategy Nash Equilibrium is a mixed strategy profile $(L_1,\dots, L_n)$ s.t. for every j and every $L_j' \in \text{Lot}(S_j)$

$$u_j(L_j, L_{-j}) \geq u_j(L_j', L_{-j})$$

- Any normal form game has at least one mixed-strategy NE [Nash 1950]
- Any L_j with $L_j(s_j)=1$ for some $s_j \in S_j$ is called a pure strategy
- Pure Nash Equilibrium
  - a special mixed-strategy NE $(L_1,\dots, L_n)$ where all strategies are pure strategy

Example: mixed-strategy NE

Column player

Row player

	H	T
H	(-1 , 1)	(1 , -1)
T	(1 , -1)	(-1 , 1)

- (**H**@0.5+**T**@0.5, **H**@0.5+**T**@0.5)

Row player's strategy — Column player's strategy

Best responses

- For any agent j, given any other agents' strategies L_{-j}, the set of best responses is
  - BR(L_{-j}) = $\text{argmax}_{s_j}\, u_j(s_j, L_{-j})$
  - It is a set of pure strategies
- A strategy profile L is an NE if and only if
  - for all agent j, L_j only takes positive probabilities on BR(L_{-j})

Proof of Nash's Theorem

- Idea: Brouwer's fixed point theorem
  - for any continuous function f mapping a compact convex set to itself, there is a point x such that $f(x) = x$
- The setting for n players
  - The compact convex set: $\Pi_{j=1}{}^{n}$ Lot (S_j)
  - f: $L_{ji} \rightarrow \frac{L_{ji}+g_{ji}(L)}{1+\sum_i g_{ji}(L)}$
  - $g_{ji}(L) = \max(u_j(L_{-j}, a_{ji}) - u_j(L), 0)$= improvement if switching to a_{ji}
- Fixed point L^* must be an NE
  - if not, there exists j s.t. $\sum_i g_{ji}(L)$>0
  - $L_{ji} > 0 \Leftrightarrow g_{ji}(L) > 0$
  - Improvement on all support, impossible

Computing NEs by guessing supports

- Step 1. "Guess" the support sets Supp_j for all players
- Step 2. Check if there are ways to assign non-negative probabilities to Supp_j s.t.
  - for all $s_j, t_j \in \text{Supp}_j$, $u_j(s_j, L_{-j}) = u_j(t_j, L_{-j})$
  - for all $s_j, \in \text{Supp}_j, t_j \notin \text{Supp}_j$, $u_j(s_j, L_{-j}) \geq u_j(t_j, L_{-j})$

Example

	Column player H	Column player T
Row player H	(-1 , 1)	(1 , -1)
Row player T	(1 , -1)	(-1 , 1)

- Hypothetical $Supp_{Row}$={H,T}, $Supp_{Col}$={H,T}
  - Pr_{Row} (H)=p, Pr_{Col} (H)=q
    - Row player: 1-q-q=q-(1-q)
    - Column player: 1-p-p=p-(1-p)
    - p=q=0.5
- Hypothetical $Supp_{Row}$={H,T}, $Supp_{Col}$={H}
  - Pr_{Row} (H)=p
    - Row player: -1 = 1
    - Column player: p-(1-p)>=-p+(1-p)
    - No solution

Mixed-Strategy NE
The Game of Chicken

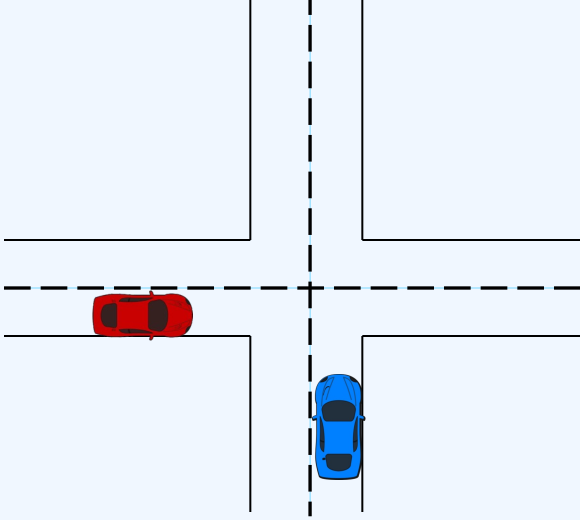

- Participation

Column player

	Dare	Chicken
Dare	(0 , 0)	(7 , 2)
Chicken	(2 , 7)	(6 , 6)

Row player

Finding all mixed NE

- Step 0. Iteratively eliminate pure strategies that are strictly dominated
  - If just finding one mixed NE, then weak dominance suffices
- Step 1. "Guess" the support sets Supp$_j$ for all players
- Step 2. Check if there are ways to assign non-negative probabilities

Dominated by mixed strategies

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

- Row player
  - 0.5 U + 0.5 D = (2.5, 2.5, 4) > (2, 2, 3) = M
- Remaining is homework

Rock Paper Scissors: Lirong vs. young Daughter

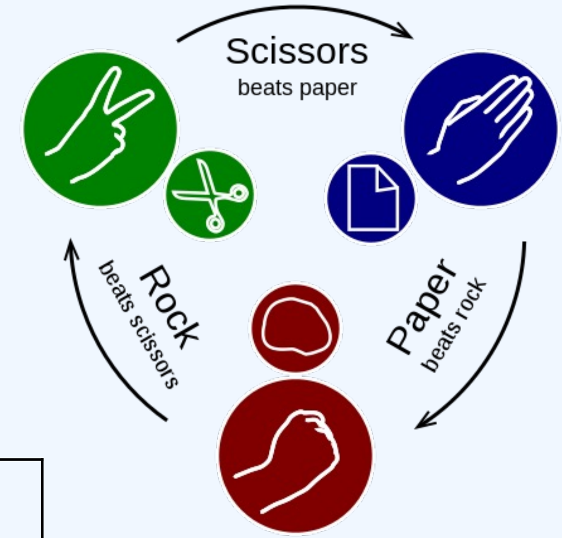

Daughter

Lirong

	mini R	mini P
R	(0 , 0)	(-1 , 1)
P	(1 , -1)	(0 , 0)
S	(-1, 1)	(1 , -1)

- Hypothetical $Supp_L$={P,S}, $Supp_D$: {mini R, mini P}
  - Pr_L (P)=p, Pr_D (mini R) = q
  - Lirong: q = (1-q)-q
  - Daughter: -1p+(1-p) = -1(1-p)
  - p=2/3, q=1/3

Correlated Equilibrium

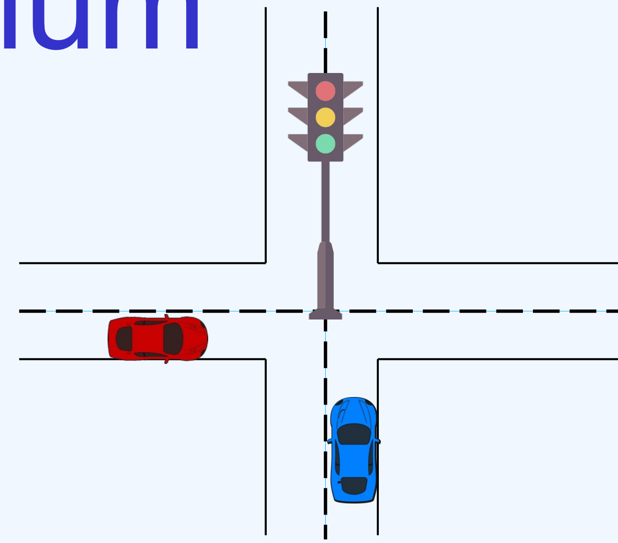

- Solution: Traffic light
  - Tell each play what to do
  - No incentive to deviate

	Dare	Chicken
Dare	(0 , 0)	(7 , 2)
Chicken	(2 , 7)	(6 , 6)

- Signal: (C,C)@1/3 + (C,D)@1/3 + (D,C)@1/3
  - When seeing C, u(C) = 4 > u(D) = 3.5
  - When seeing D, u(D) = 7 > u(C) = 6

Correlated Equilibrium: formal definition

- A correlated equilibrium x is a distribution over $\Pi_j S_j$
- For all players j, all s_j , $s_j' \in S_j$

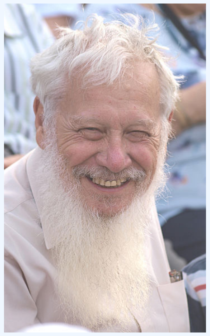

Belief about instruction of other players

$$E_{s_{-j} | x, s_j} u_j (s_j, s_{-j}) \geq E_{s_{-j} | x, s_j} u_j (s_j', s_{-j})$$

Follow the instruction   Does not follow the instruction

Computing CE: Linear Programming

- Variables: the distribution x
- Objective: any
- Constraints: incentive constraints
- Example: chicken game

	D	C
D	x_{DD}	x_{DC}
C	x_{CD}	x_{CC}

	D	C
D	(0 , 0)	(7 , 2)
C	(2 , 7)	(6 , 6)

- Obj: $9x_{DC} + 9x_{CD} + 12x_{CC}$
- Constraints for row player
  - Receiving signal D: $7\, x_{DC} \geq 2\, x_{DD} + 6\, x_{DC}$
  - Receiving signal C: $2\, x_{CD} + 6\, x_{CC} \geq 7\, x_{CC}$
- Constraints for column player
  - Receiving signal D: $7\, x_{CD} \geq 2\, x_{DD} + 6\, x_{CD}$
  - Receiving signal C: $2\, x_{DC} + 6\, x_{CC} \geq 7\, x_{CC}$

Extensive-form games

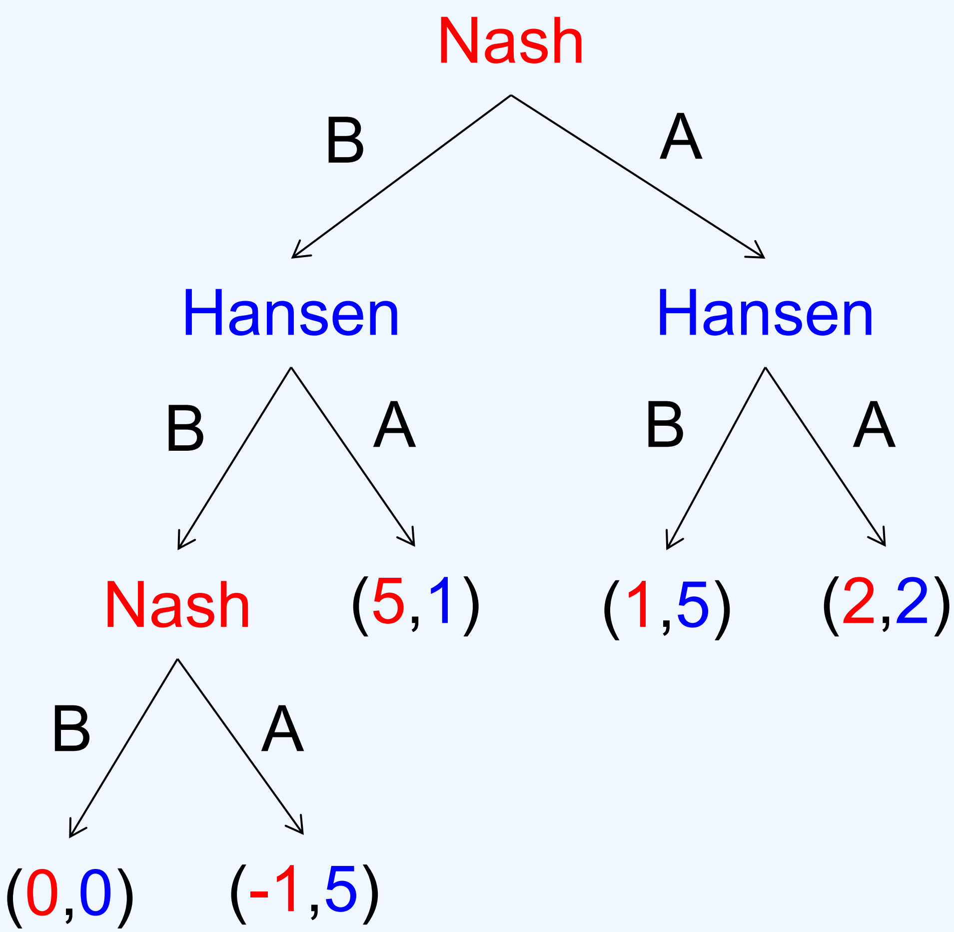

leaves: utilities (Nash,Hansen)

- Players move sequentially
- Outcomes: leaves
- Preferences are represented by utilities
- A strategy of player j is a combination of all actions at her nodes
- All players know the game tree (complete information)
- At player j's node, she knows all previous moves (perfect information)

Convert to normal-form

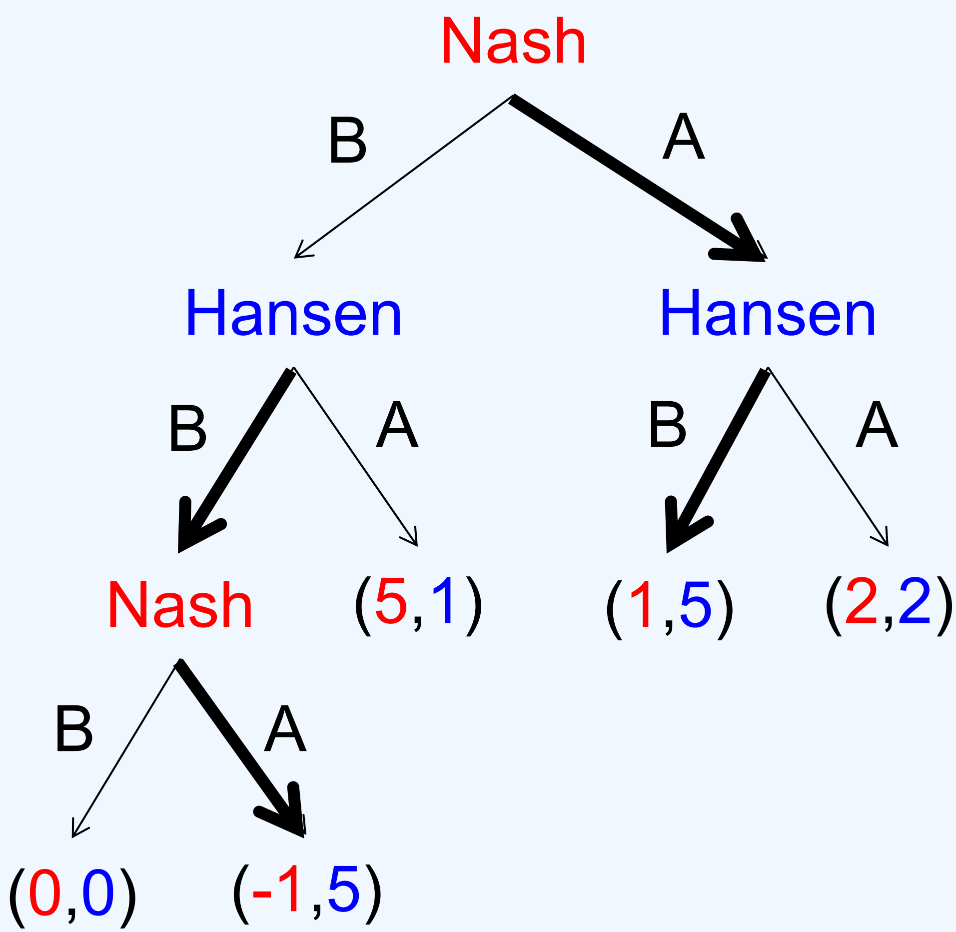

Hansen

Nash	(B,B)	(B,A)	(A,B)	(A,A)
(B,B)	(0,0)	(0,0)	(5,1)	(5,1)
(B,A)	(-1,5)	(-1,5)	(5,1)	(5,1)
(A,B)	(1,5)	(2,2)	(1,5)	(2,2)
(A,A)	(1,5)	(2,2)	(1,5)	(2,2)

Nash: (Up node action, Down node action)

Hansen: (Left node action, Right node action)

Subgame perfect equilibrium

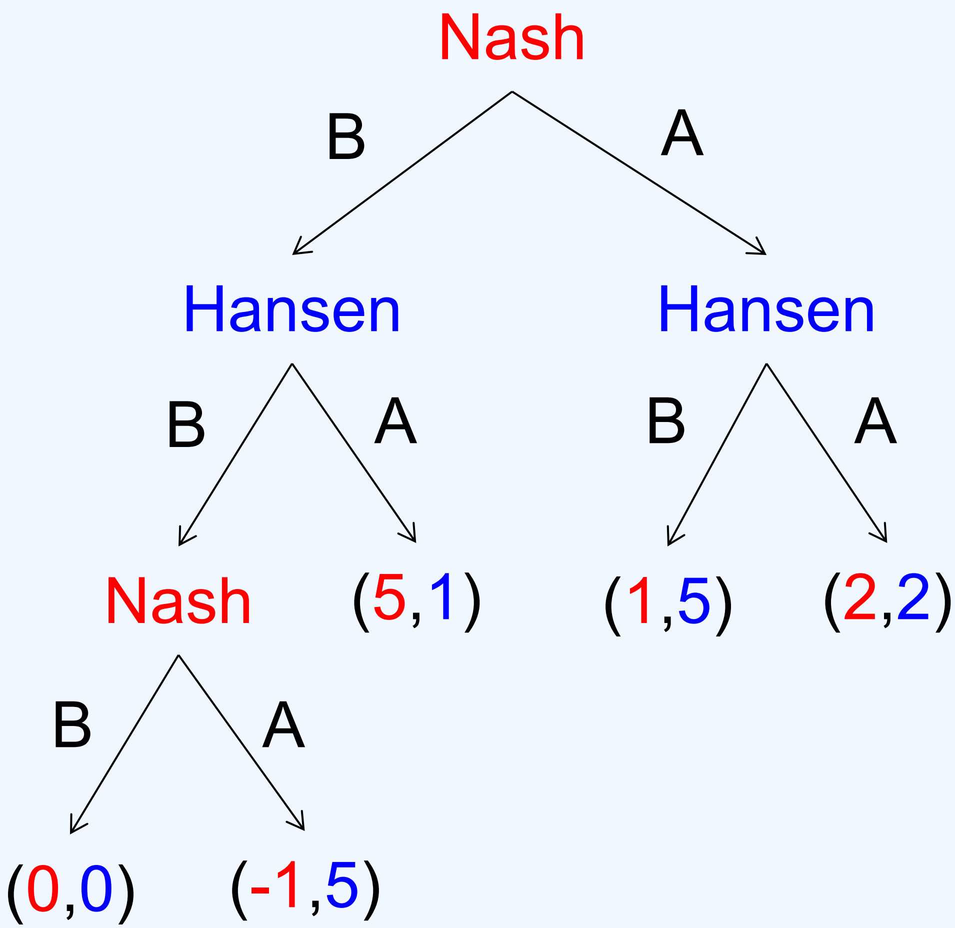

- Usually too many NE
- (pure) SPNE
  - a refinement (special NE)
  - also an NE of any subgame (subtree)

Backward induction

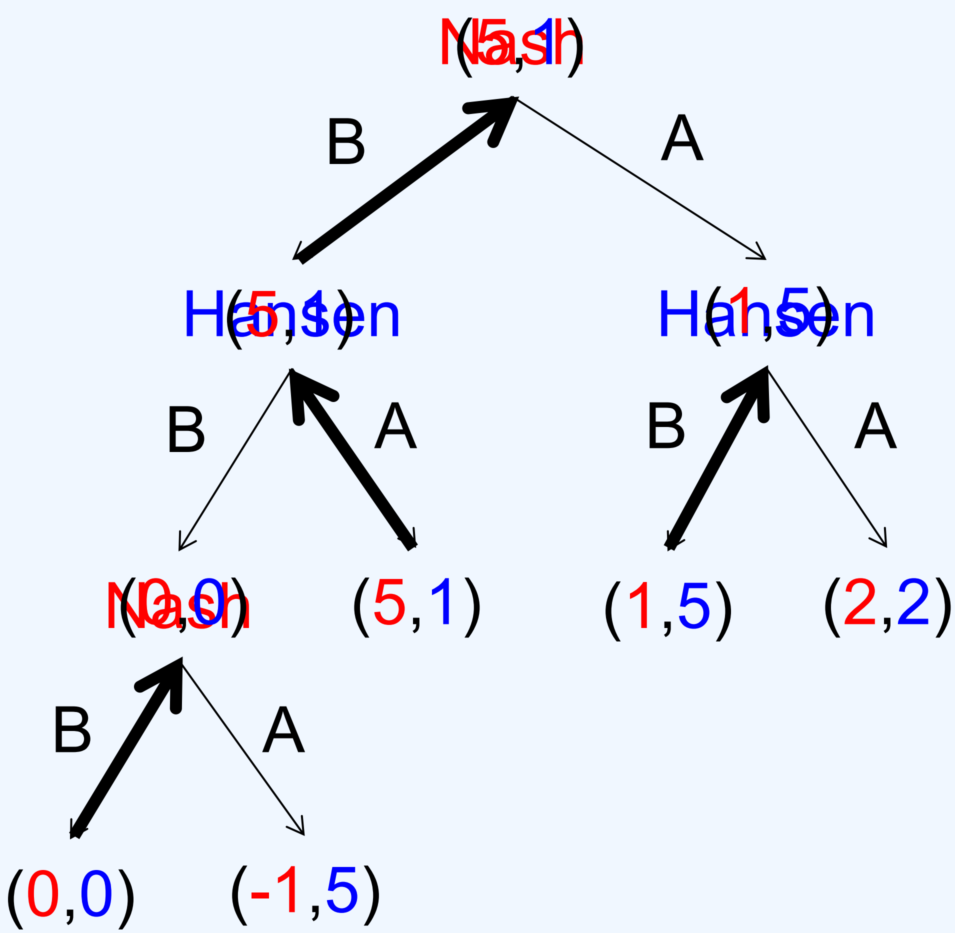

- Determine the strategies bottom-up
- Unique if no ties in the process
- All SPNE can be obtained, if
  - the game is finite
  - complete information
  - perfect information

Algorithmic Game Theory

- Algorithmic game theory is an area in the intersection of game theory and computer science, whose objective is to understand and design algorithms in strategic environments ---wiki

- Complexity of computing NE
  - **P**a**PAD**imitriou complete
    - *Polynomial parity argument on a directed graph*
  - Conjecture P != PPAD

Topic: Price of Anarchy
[Koutsoupias & Papadimitriou STACS 99]

- SW(S): social welfare of strategy profile S
- Price of Anarchy = $\frac{\text{OPT SW}}{\text{Worst equilibrium SW}}$
  - measures the worst-case loss of strategic behavior
  - Game of Chicken 12/9

	D	C
D	(0 , 0)	(7 , 2)
C	(2 , 7)	(6 , 6)

- Price of Stability = $\frac{\text{OPT SW}}{\text{Best equilibrium SW}}$

Review: Game Theory

- What?
  - Self-interested agents may behave strategically
- Why?
  - Hard to predict the outcome for strategic agents
- How?
  - A general framework for games
    - Solution concept: Nash equilibrium
    - Improvement: Correlated equilibrium
  - Preferences: utility theory
  - Special games
    - Normal form games: mixed Nash equilibrium
    - Extensive form games: subgame-perfect equilibrium