Introduction to Game Theory

Lirong Xia

Rensselaer

Fall, 2016

Homework 1

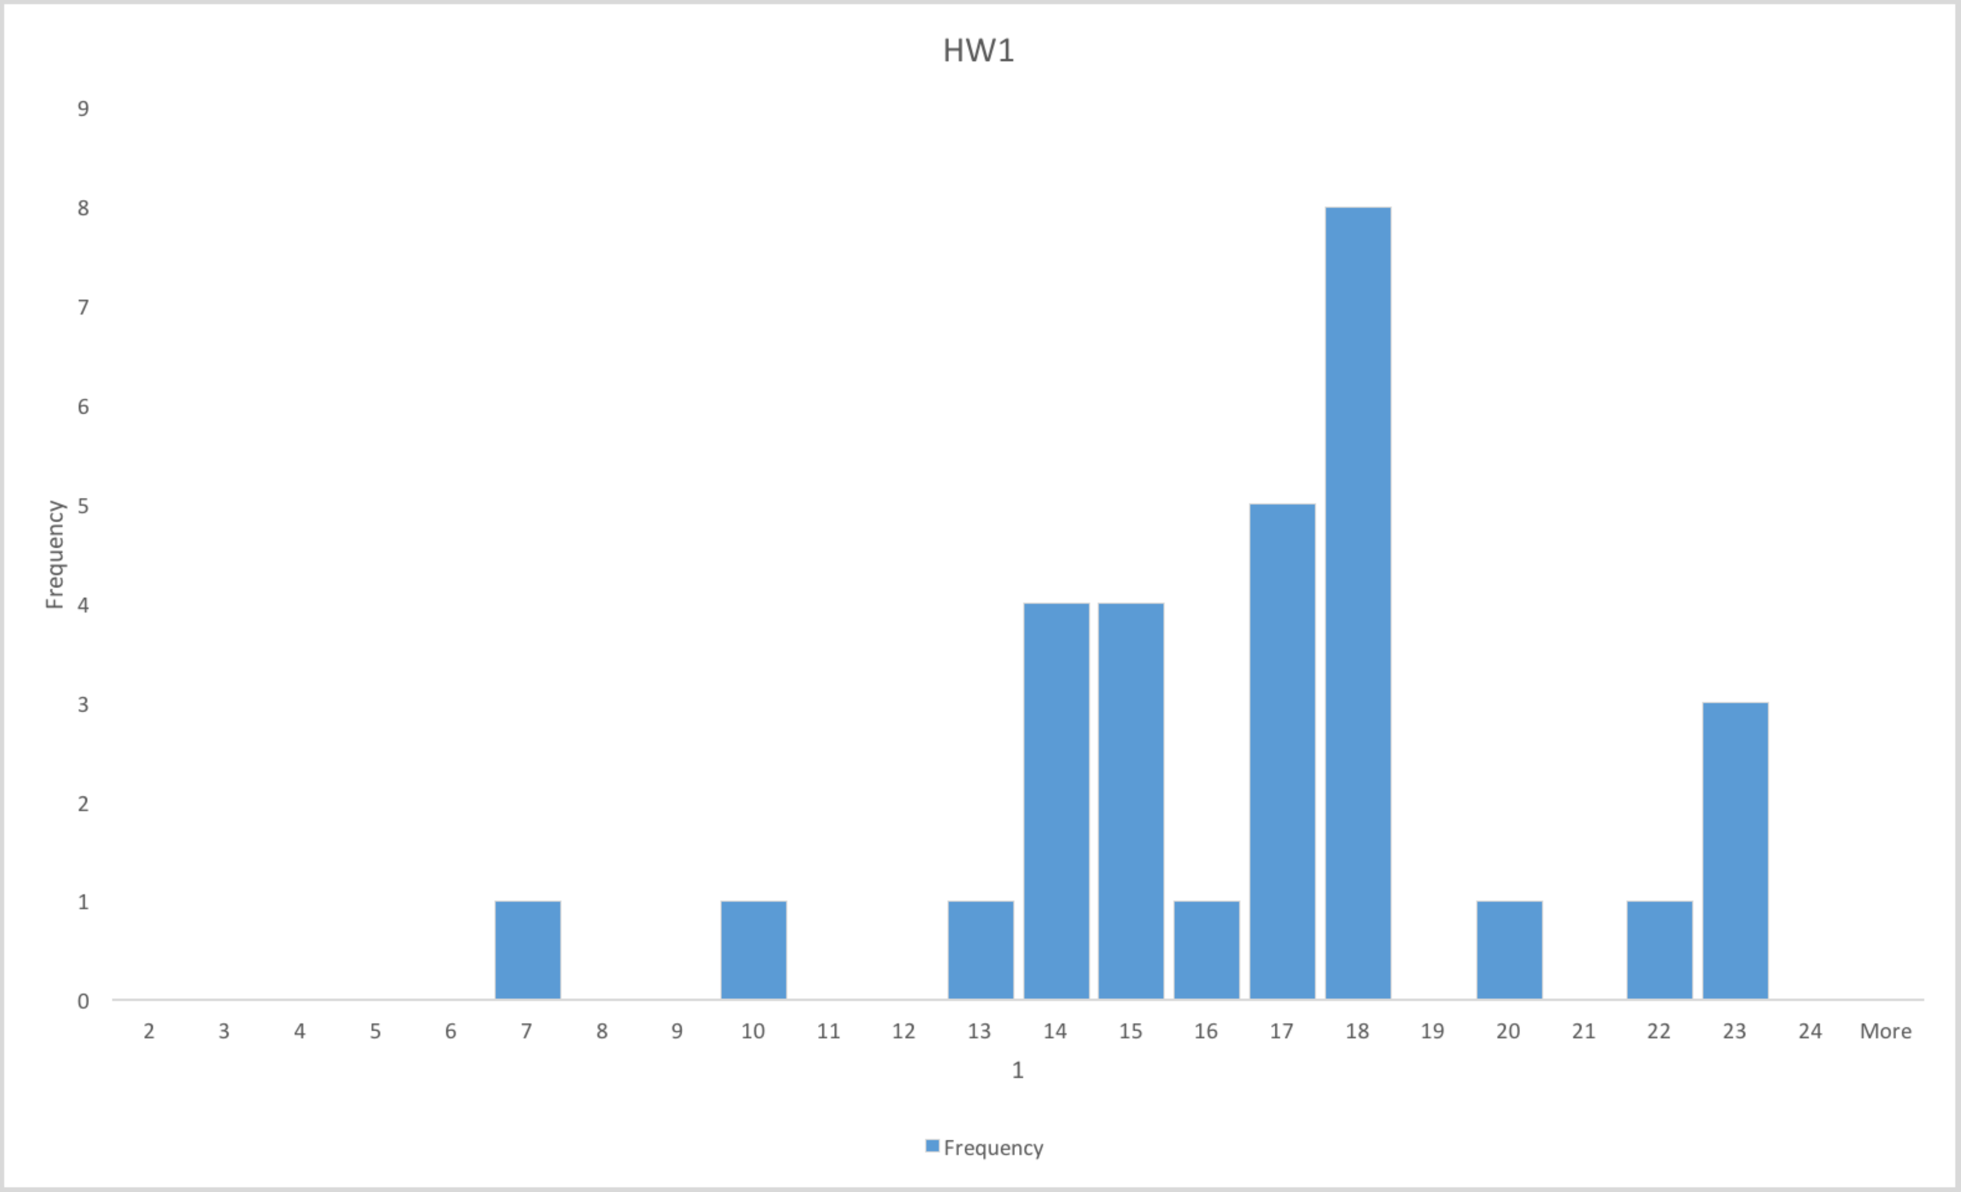

HW1
Frequency
9
8
7
6
5
4
3
2
1
0
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
More
1
Frequency

Announcements

- We will use LMS for submission and grading
- Please just submit one copy
- Please acknowledge your team mates

Remarks

- Show the math and formal proof
  - No math/steps, no points (esp. in midterm)
  - Especially Problem 1, 4, 5
- Problem 1
  - Must use u(1M) etc.
  - Must hold for all utility function
- Problem 2
  - must show your calculation
  - For Schulze, if you have already found one strict winner, no need to check other alternatives
  - Kemeny outputs a single winner, unless otherwise mentioned
- Problem 3.2
  - b winning itself is not a paradox
  - people can change the outcome by not voting is not a paradox

Last class

- Mallows' model
- MLE and MAP
- P = {a>b>c, 2@c>b>a}
- Likelihood
- Prior distribution
  - Pr(a>b>c)=Pr(a>c>b)=0.3
  - all other linear orders have prior 0.1
- Posterior distribution
  - proportional to Likelihood*prior

Last class

- Plackett-Luce model
  - Example
    - alternatives {a,b,c}
    - parameter space {(4,3,3), (3,4,3), (3,3,4)}
- MLE and MAP
- P = {a>b>c, 2@c>b>a}
- Likelihood
- Prior distribution
  - Pr(4,3,3)=0.8
  - all others have prior 0.1
- Posterior distribution
  - proportional to Likelihood*prior

Review: manipulation

(ties are broken alphabetically)

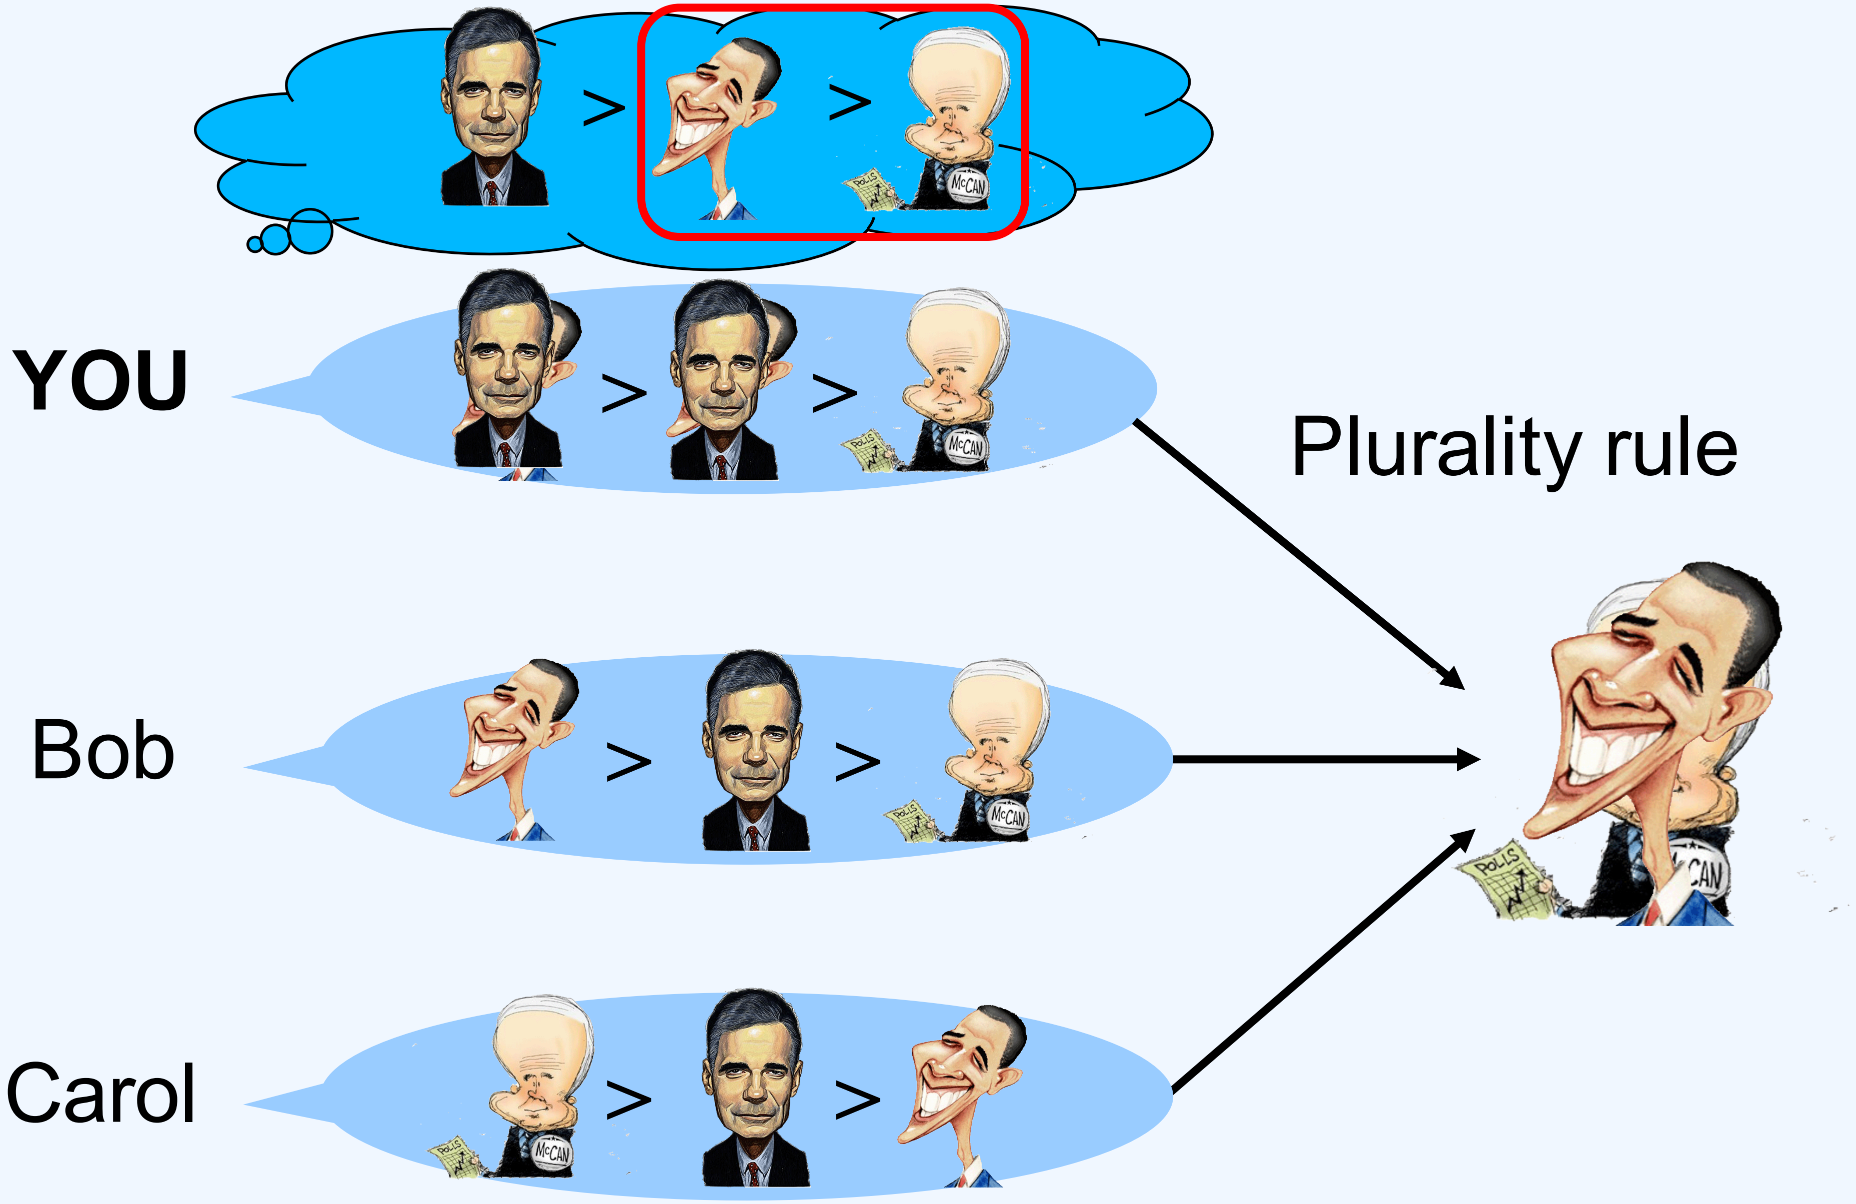

YOU

Bob

Carol

Plurality rule

What if everyone is incentivized to lie?

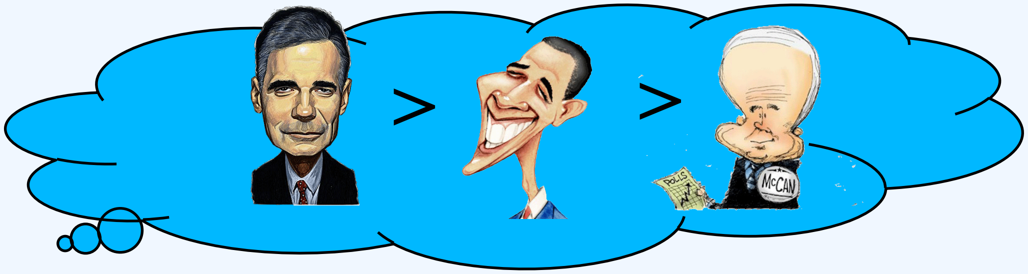

YOU

Plurality rule

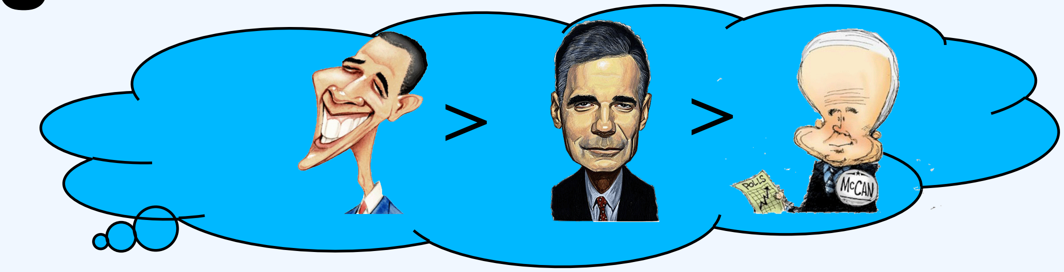

Bob

?

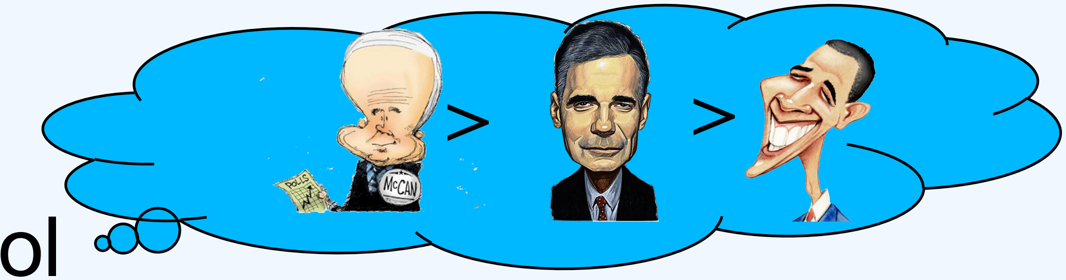

Carol

Today's schedule: game theory

- What?
  - Agents may have incentives to lie
- Why?
  - Hard to predict the outcome when agents lie
- How?
  - A general framework for games
    - Solution concept: Nash equilibrium
  - Modeling preferences and behavior: utility theory
  - Special games
    - Normal form games: mixed Nash equilibrium
    - Extensive form games: subgame-perfect equilibrium

A game

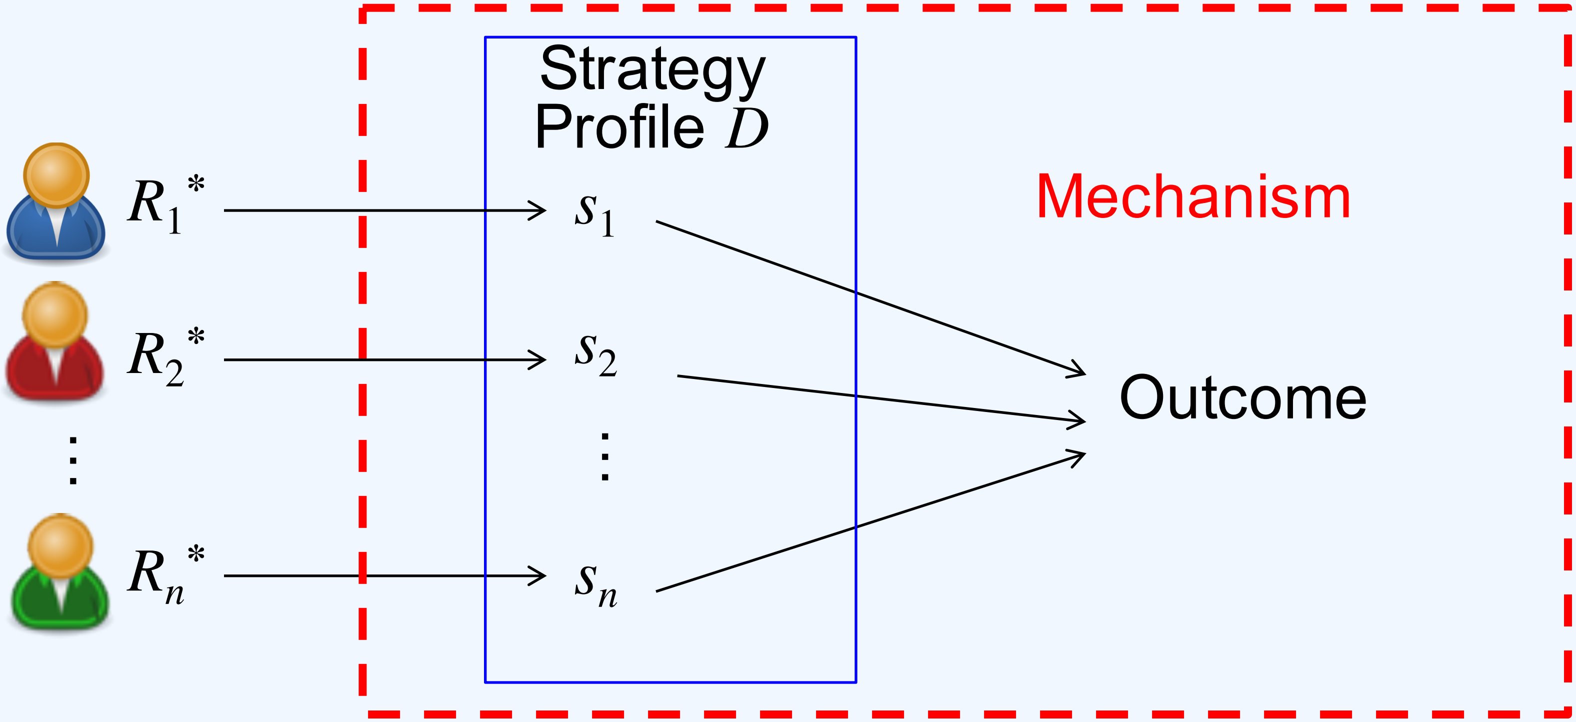

Strategy Profile D

Mechanism

R_1^* → s_1

R_2^* → s_2

$\vdots$

R_n^* → s_n

Outcome

- 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
- Preferences: total preorders (full rankings with ties) over O
- often represented by a utility function $u_i : \Pi_j\ S_j \rightarrow R$
- Mechanism $f : \Pi_j\ S_j \rightarrow O$

A game of plurality elections

YOU

Bob

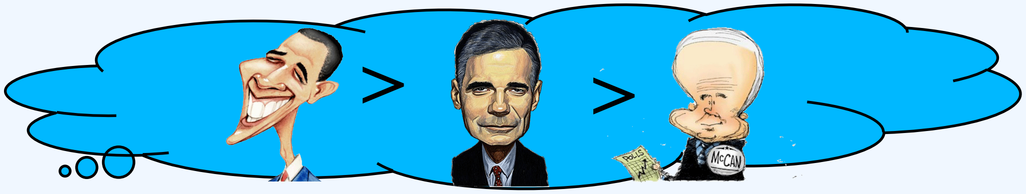

Carol

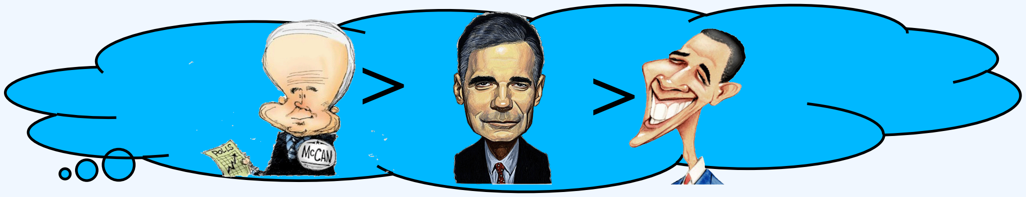

Plurality rule

?

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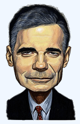 , 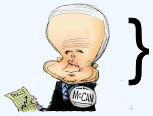 }
- Strategies: S_j = Rankings(O)
- Preferences: See above
- Mechanism: the plurality rule

A game of two prisoners

Column player

Row player

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

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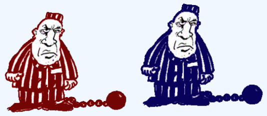
- Strategies: { Cooperate, Defect }
- Outcomes: {(-2 , -2), (-3 , 0), (0 , -3), (-1 , -1)}
- Preferences: self-interested 0 > -1 > -2 > -3
  -  : (0 , -3) > (-1 , -1) > (-2 , -2) > (-3 , 0)
  -  : (-3 , 0) > (-1 , -1) > (-2 , -2) > (0 , -3)
- Mechanism: the table

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 $s_j' \in S_j$,

    $f(s_j, s_{-j}) \geq_j f(s_j', s_{-j})$ or $u_j(s_j, s_{-j}) \geq u_j(s_j', s_{-j})$
  - $s_{-j} = (s_1,\ldots,s_{j-1}, s_{j+1},\ldots,s_n)$
  - no single player can be better off by deviating

Prisoner's dilemma

Column player

Row player

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

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
Nash Row player	**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

Column player

Row player

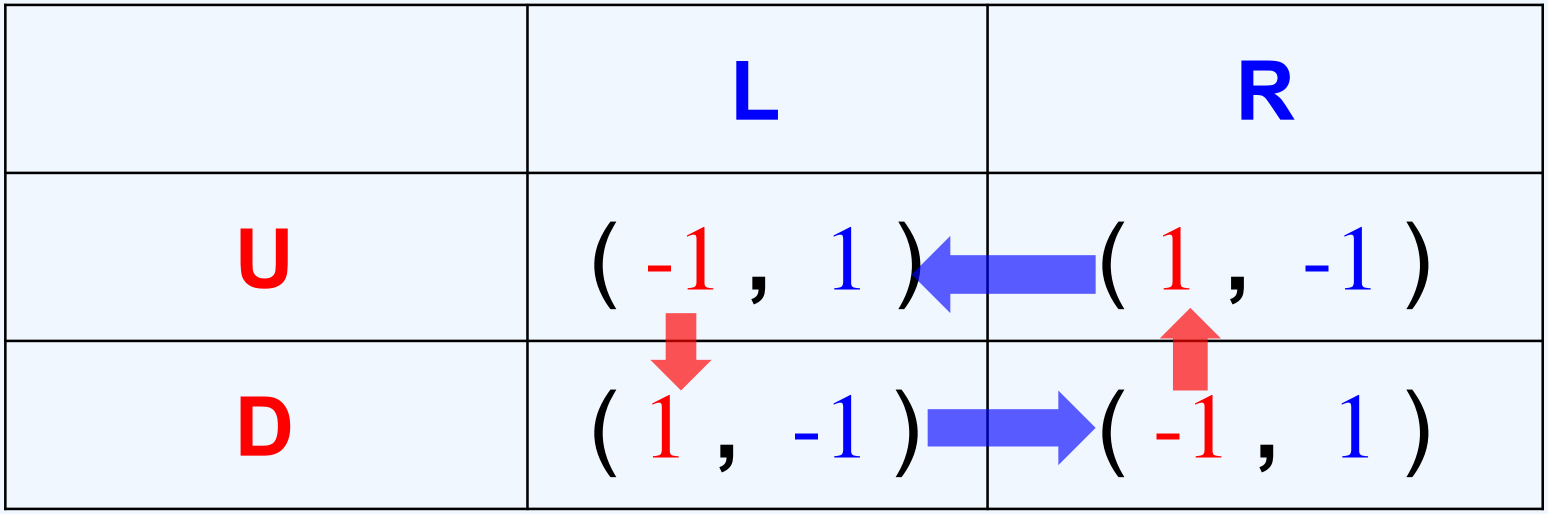

	L	**R**
U	(-1 , 1)	(1 , -1)
D	(1 , -1)	(-1 , 1)

- 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}
- 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, but if exists, then must be unique

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

The Game of Chicken

- Two drivers for a single-lane bridge from opposite directions and each can either (S)traight or (A)way.
  - If both choose S, then crash.
  - If one chooses A and the other chooses S, the latter “wins”.
  - If both choose A, both are survived

Column player

Row player

	A	**S**
A	(0 , 0)	(0 , 1)
S	(1 , 0)	(-10 , -10)

NE

Rock Paper Scissors

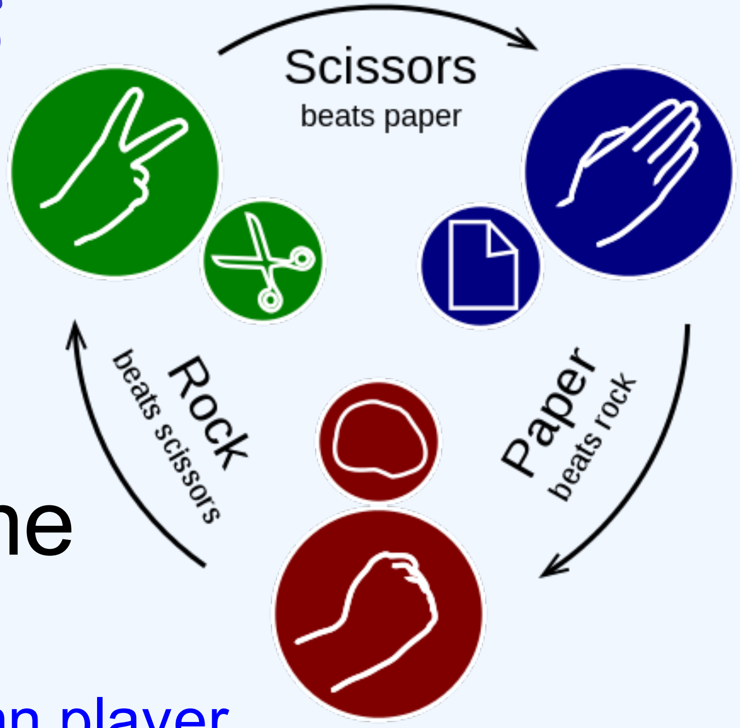

- ➢Actions: {R, P, S}
- ➢Two-player zero sum game

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

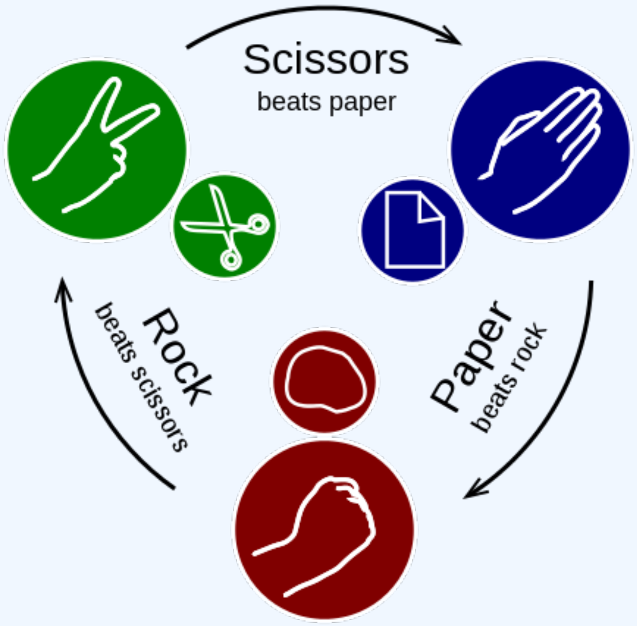

- Actions
  - Lirong: {R, P, S}
  - Daughter: {mini R, mini P}
- Two-player zero sum game

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)~~

Iterated Elimination: Lirong vs. young Daughter

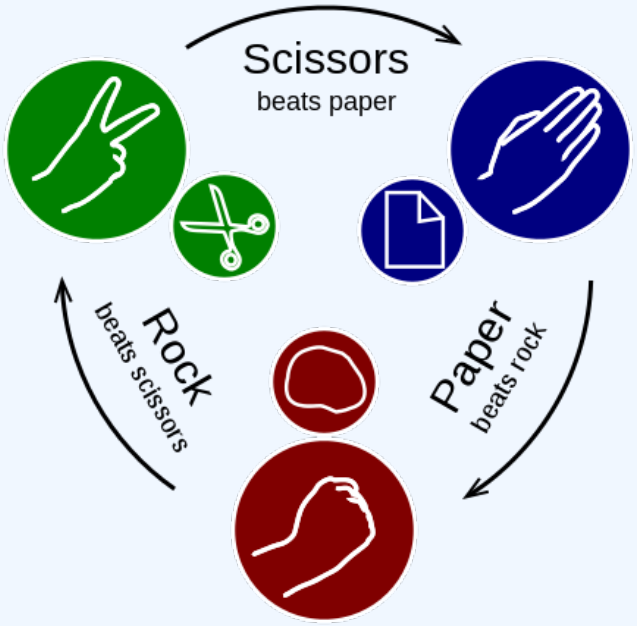

- Actions
  - Lirong: {R, P, S}
  - Daughter: {mini R, mini P}
- Two-player zero sum game

No pure NE

Daughter

Lirong

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

Normal form games

- Given pure strategies: S_j for agent j

Normal form games

- Players: N={1,…,n}
- Strategies: lotteries (distributions) over S_j
  - $L_j \in$ Lot(S_j) is called a mixed strategy
  - (L_1,…, L_n) is a mixed-strategy profile
- Outcomes: Π_j Lot(S_j)
- Mechanism: f (L_1,…,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

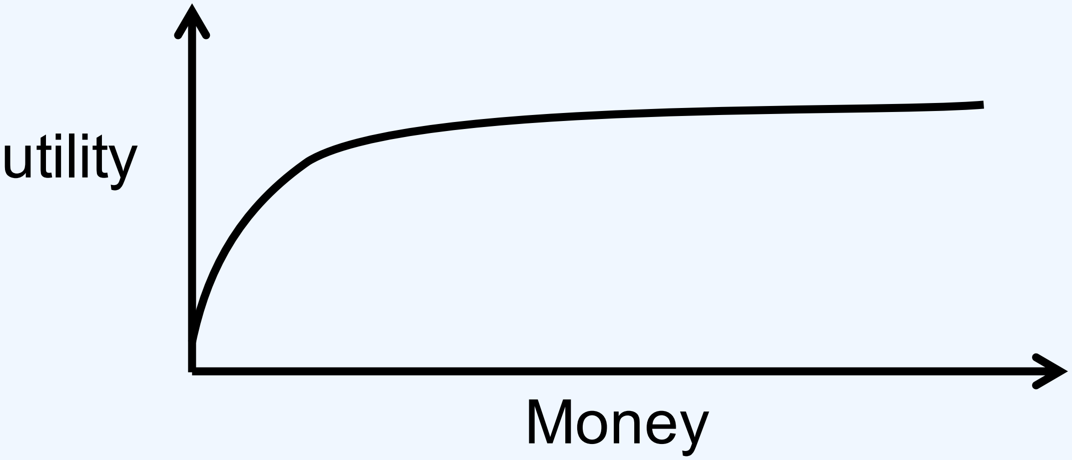

Money	0	5	30	5M	30M
Utility	1	3	10	100	150

- 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

- Given pure strategies: S_j for agent j
- 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: $f\,(L_1,\ldots,L_n) = p$, 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 $(L_1,\ldots, L_n)$ 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 $(L_1,\ldots, L_n)$ where all strategies are pure strategy

Example: mixed-strategy NE

Column player

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

- $(\underbrace{\mathbf{H}@0.5+\mathbf{T}@0.5}_{\text{Row player's strategy}},\ \underbrace{\mathbf{H}@0.5+\mathbf{T}@0.5}_{\text{Column player's strategy}})$

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})

Computing NEs by guessing best responses

- Step 1. “Guess” the best response sets BR_j for all players
- Step 2. Check if there are ways to assign probabilities to BR_j to make them actual best responses

Example

Column player

Row player

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

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

Rock Paper Scissors: Lirong vs. young Daughter

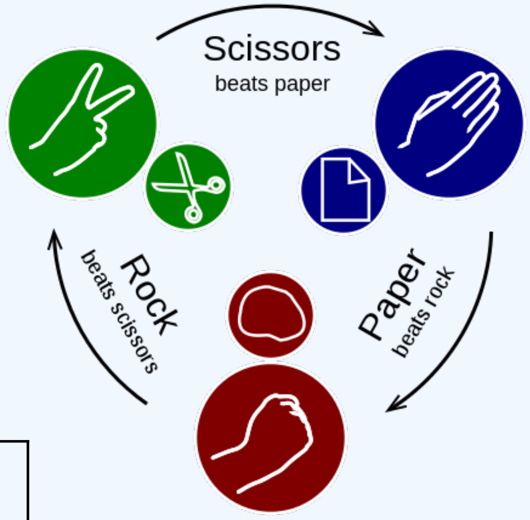

Daughter

Lirong

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

- Hypothetical BR_L={P,S}, BR_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

Extensive-form games

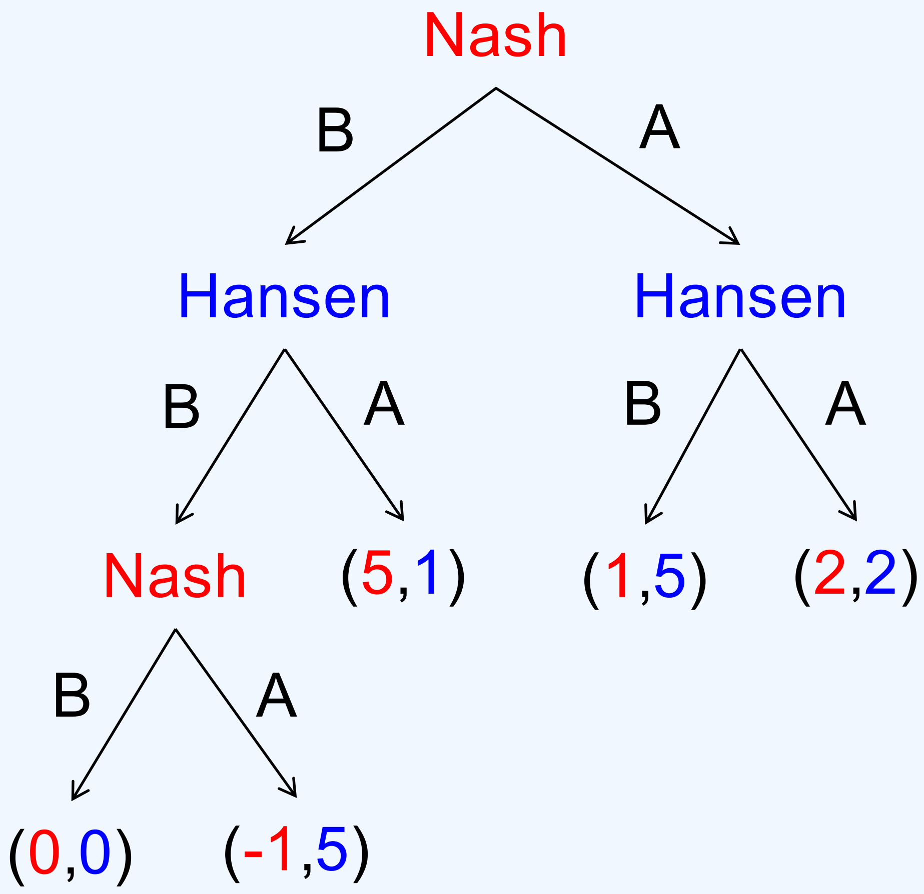

leaves: utilities (Nash,Hansen)

- 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

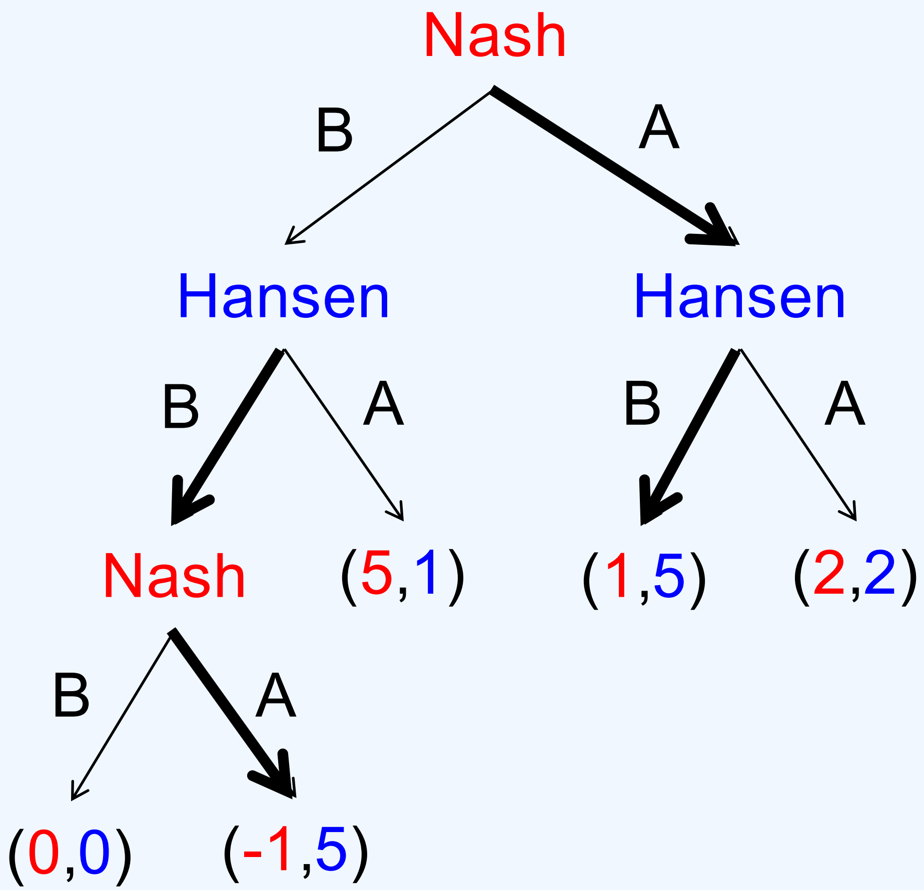

Hansen

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

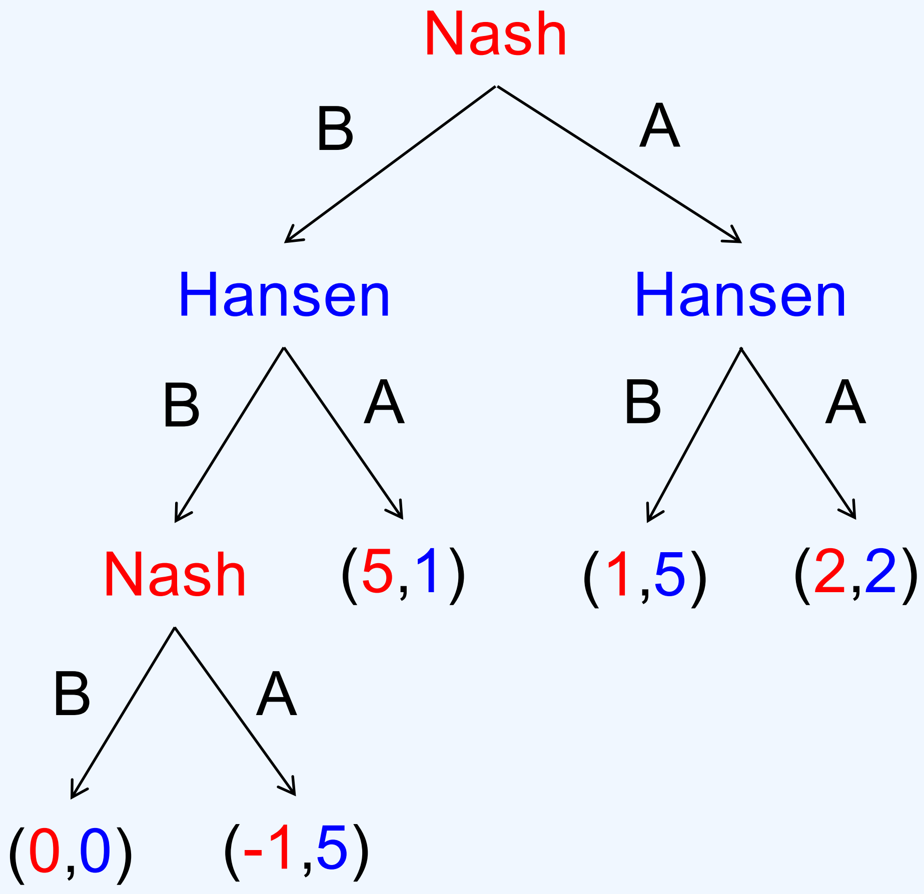

- Usually too many NE
- (pure) SPNE
  - a refinement (special NE)
  - also an NE of any subgame (subtree)

Backward induction

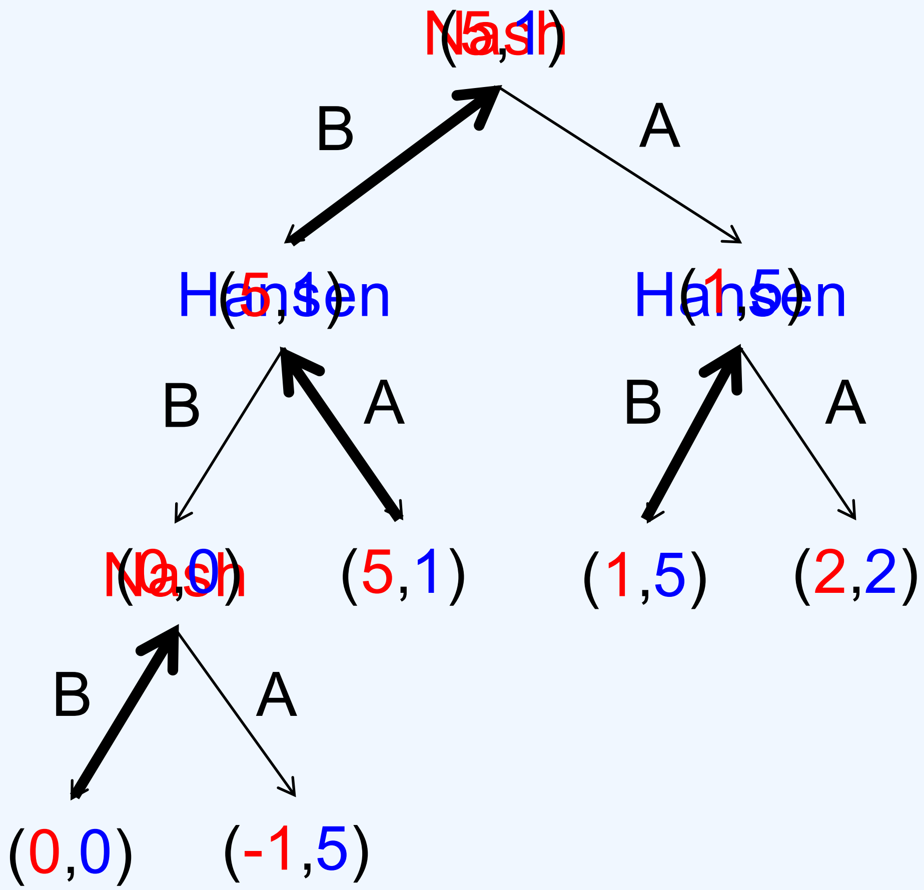

- Determine the strategies bottom-up
- Unique if no ties in the process
- All SPNE can be obtained, if
  - the game is finite
  - complete information
  - perfect information

A different angle

- How good is SPNE as a solution concept?
  - At least one
  - In many cases unique
  - is a refinement of NE (always exists)

Wrap up

	Preferences	Solution concept	How many	Computation
General game	total preorders	NE	0-many	
Normal form game	utilities	mixed-strategy NE pure NE	mixed: 1-many pure: 0-many	
Extensive form game	utilities	Subgame perfect NE	1 (no ties) many (ties)	Backward induction

The reading questions

- What is the problem?
  - agents may have incentive to lie
- Why we want to study this problem? How general it is?
  - The outcome is hard to predict when agents lie
  - It is very general and important
- How was problem addressed?
  - by modeling the situation as a game and focus on solution concepts, e.g. Nash Equilibrium
- Appreciate the work: what makes the work nontrivial?
  - It is by far the most sensible solution concept. Existence of (mixed-strategy) NE for normal form games
- Critical thinking: anything you are not very satisfied with?
  - Hard to justify NE in real-life
  - How to obtain the utility function?

Looking forward

- ➢ So far we have been using game theory for prediction
- ➢ How to design the mechanism?
  - when every agent is self-interested
  - as a whole, works as we want
- ➢ The next class: mechanism design

NE of the plurality election game

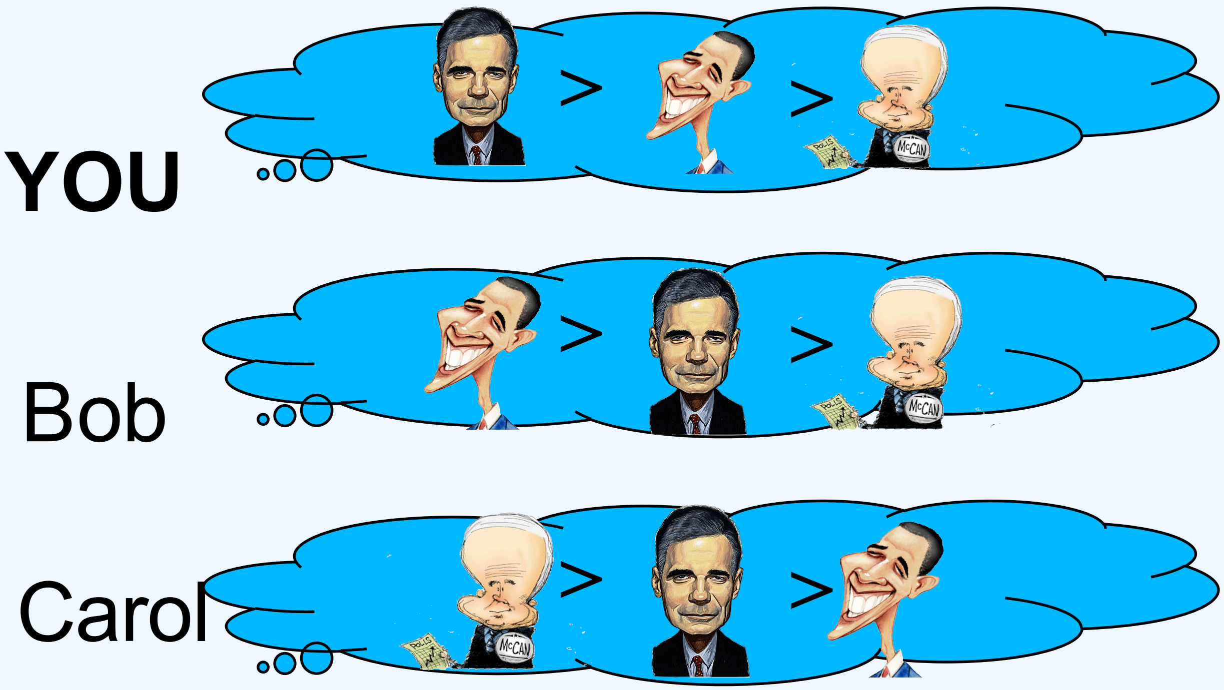

Plurality rule

?

- Players: { YOU, Bob, Carol}, n=3
- Outcomes: O = {  , 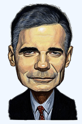 , 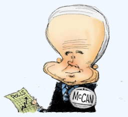 }
- Strategies: S_j = Rankings(O)
- Preferences: Rankings(O)
- Mechanism: the plurality rule