Summary of Piazza discussions

- More social choice problems
  - Friends choose a common activity, for democracy: Yuriy, Xin, Yuan
  - Legal system: Yuriy
  - sport tournaments: Kevin

Last class: Voting

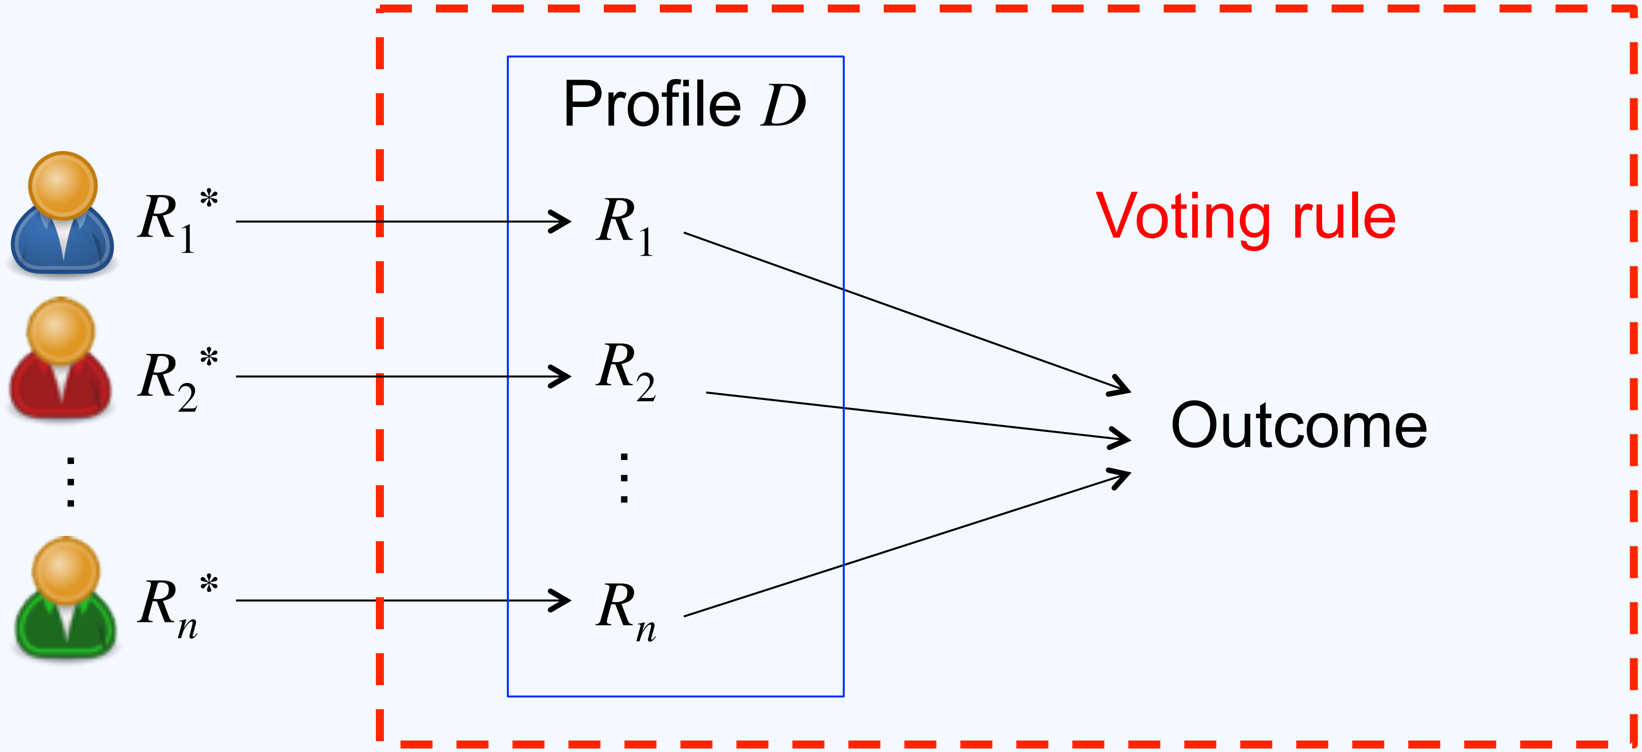

- Agents: n voters, $N=\{1,\dots,n\}$
- Alternatives: m candidates, $A=\{a_1,\dots,a_m\}$
- Outcomes:
  - winners (alternatives): $O=A$. Social choice function
  - rankings over alternatives: $O=\text{Rankings}(A)$. Social welfare function
- Preferences: R_j^* and R_j are full rankings over A
- Voting rule r: $(\text{Rankings}(A))^N \rightarrow O$

What if agents lie?

plurality rule

(ties are broken in favor of 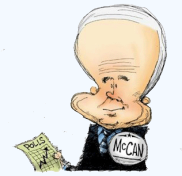)

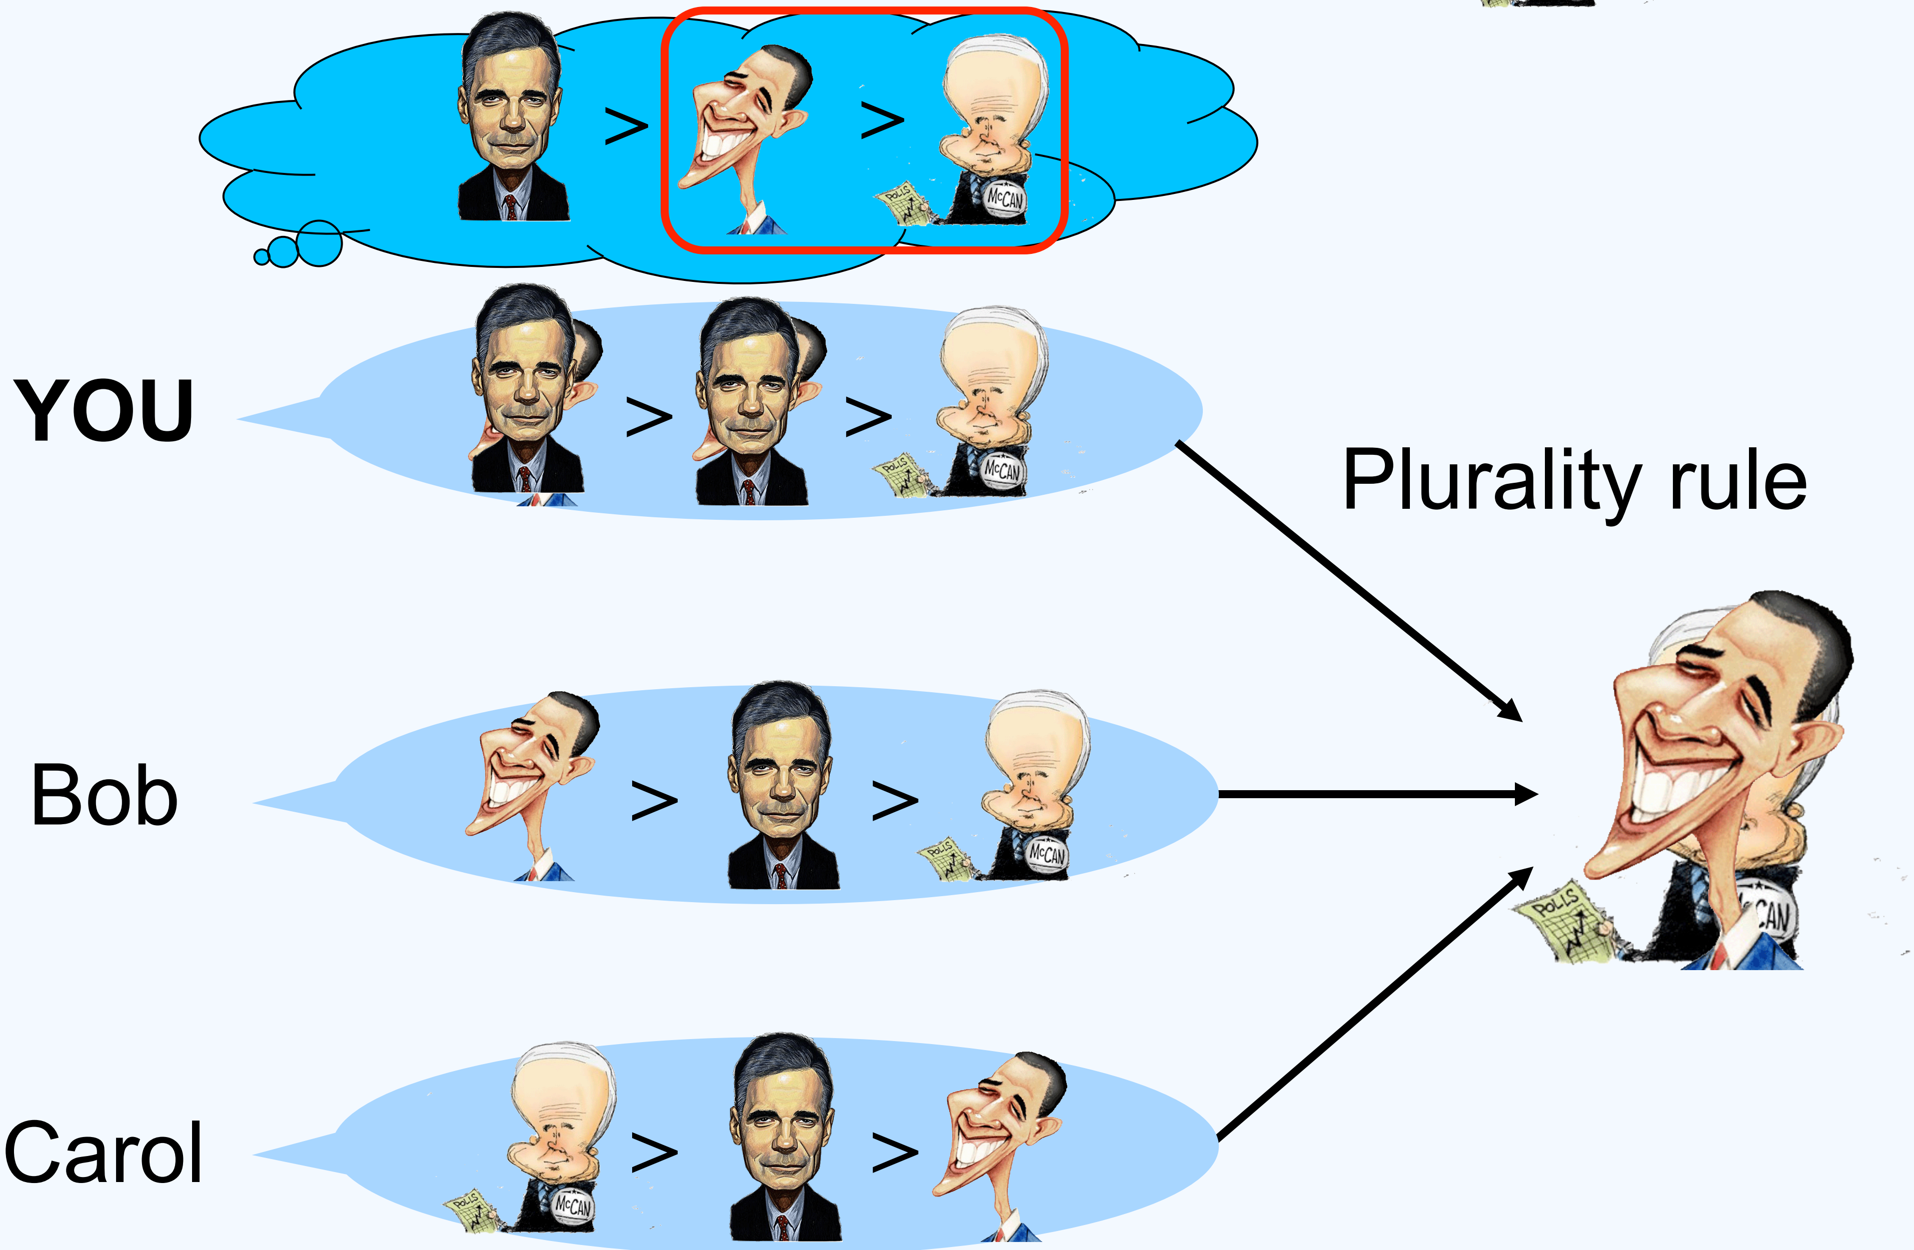

What if everyone is incentivized to lie?

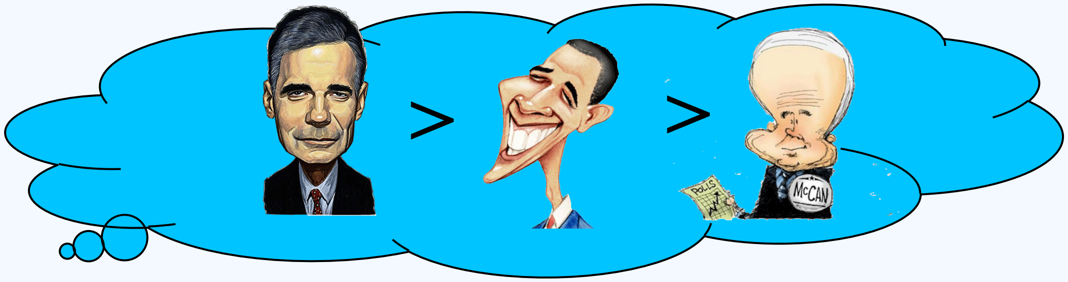

YOU

Plurality rule

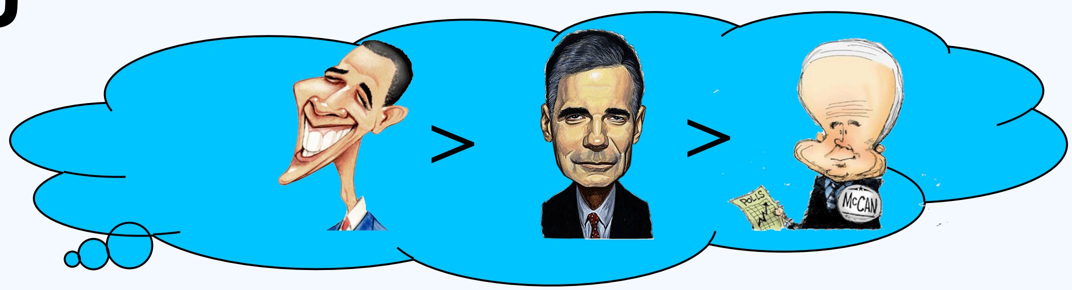

Bob

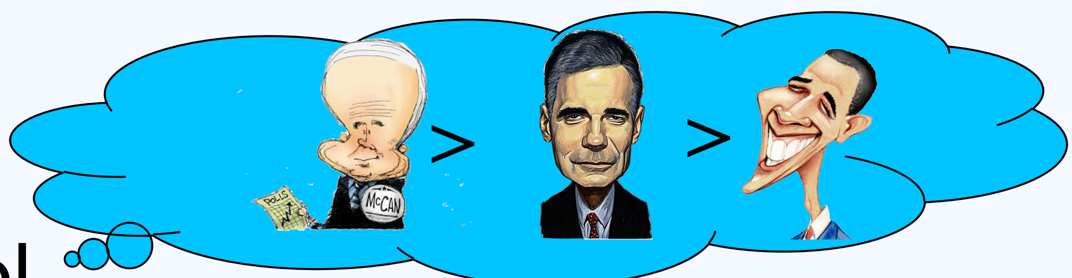

Carol

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

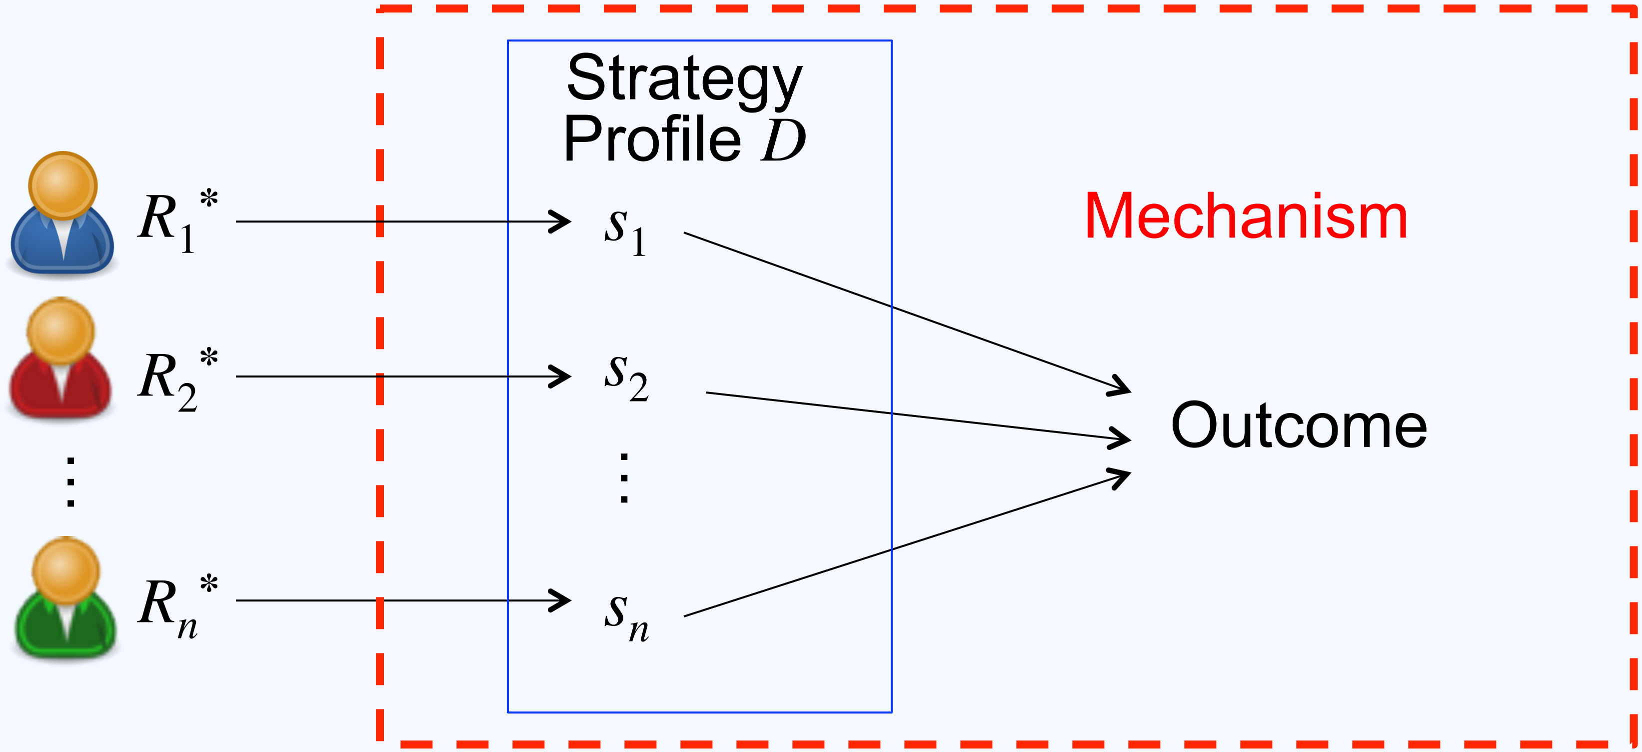

- Players: $N=\{1,\ldots,n\}$
- Strategies (actions):
  - S_j for agent j, $s_j \in S_j$
  - $(s_1,\ldots,s_n)$ is called a strategy profile.
- Outcomes: O
- Preferences: total preorders (full rankings with ties) over O
- Mechanism $f : \Pi_j S_j \to O$

A game of plurality elections

YOU

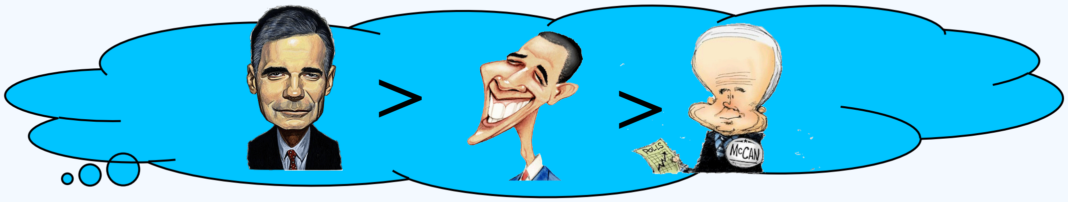

Bob

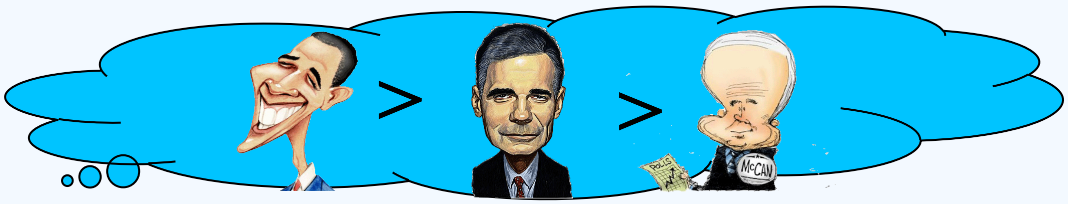

Carol

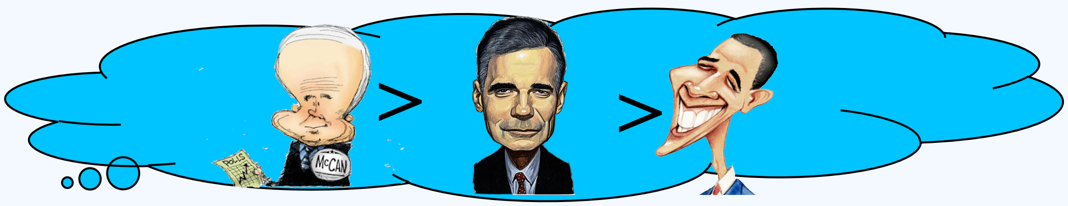

Plurality rule

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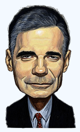, }
- Strategies: S_j = Rankings(O)
- Preferences: See above
- Mechanism: the plurality rule

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 $s_j' \in S_j$,

$$f(s_j, s_{-j}) \geq_j f(s_j', s_{-j})$$

  - $s_{-j} = (s_1,\ldots,s_{j-1}, s_{j+1},\ldots,s_n)$
  - 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
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

Research 104

- Lesson 4: scientific skepticism (critical thinking)
  - default for any argument should be “wrong”
  - it is the authors’ responsibility to prove the correctness
  - Once you find an error, try to correct and improve it
- Example: Nash equilibrium in “A beautiful mind”
  - really?

Does an NE always exists?

- Not always

Column player

Row player

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

End of story?

- How to evaluate this solution concept?
  - Does it really model real-world situations?
  - What if an NE does not exist?
    - approximate NE (beyond this course)
  - What if there are too many NE?
    - Equilibrium selection
    - refinement: a "good" NE
- Cases where an NE always exists
  - Normal form games
  - Strategy space: lotteries over pure strategies
  - Outcome space: lotteries over atom outcomes

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 $u(p_1) > u(p_2)$
- Utility is virtual, preferences are real
  - Preferences represented by utility theory have a neat characterization

Example

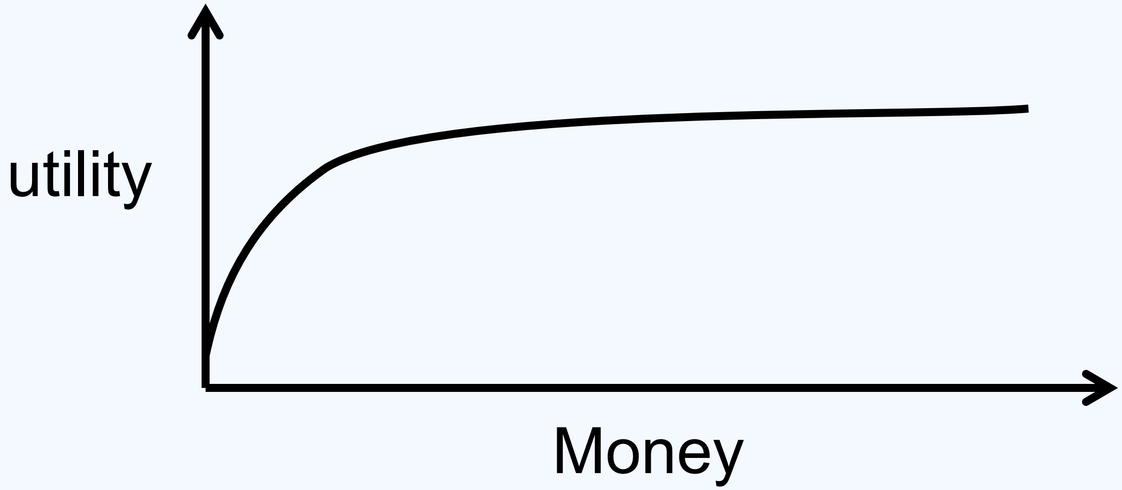

Money	0	5	30	5M	30M
Utility	1	3	10	100	150

- 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 $u_1,\ldots,u_n$

Solution concepts for normal form games

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

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

- 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	
		L	**R**
Row player	**U**	(0 , 1)	(1 , 0)
	D	(1 , 0)	(0 , 1)

- $(\underbrace{\mathbf{U}@0.5+\mathbf{D}@0.5}_{\text{Row player's strategy}},\ \underbrace{\mathbf{L}@0.5+\mathbf{R}@0.5}_{\text{Column player's strategy}})$

A different angle

- In normal-form games
  - Mixed-strategy NE = NE in the general framework
  - pure NE = a refinement of (mixed-strategy) NE
- How good is mixed-strategy NE as a solution concept?
  - At least one
  - Maybe many
  - Can use pure NE as a refinement (may not exist)

Extensive-form games

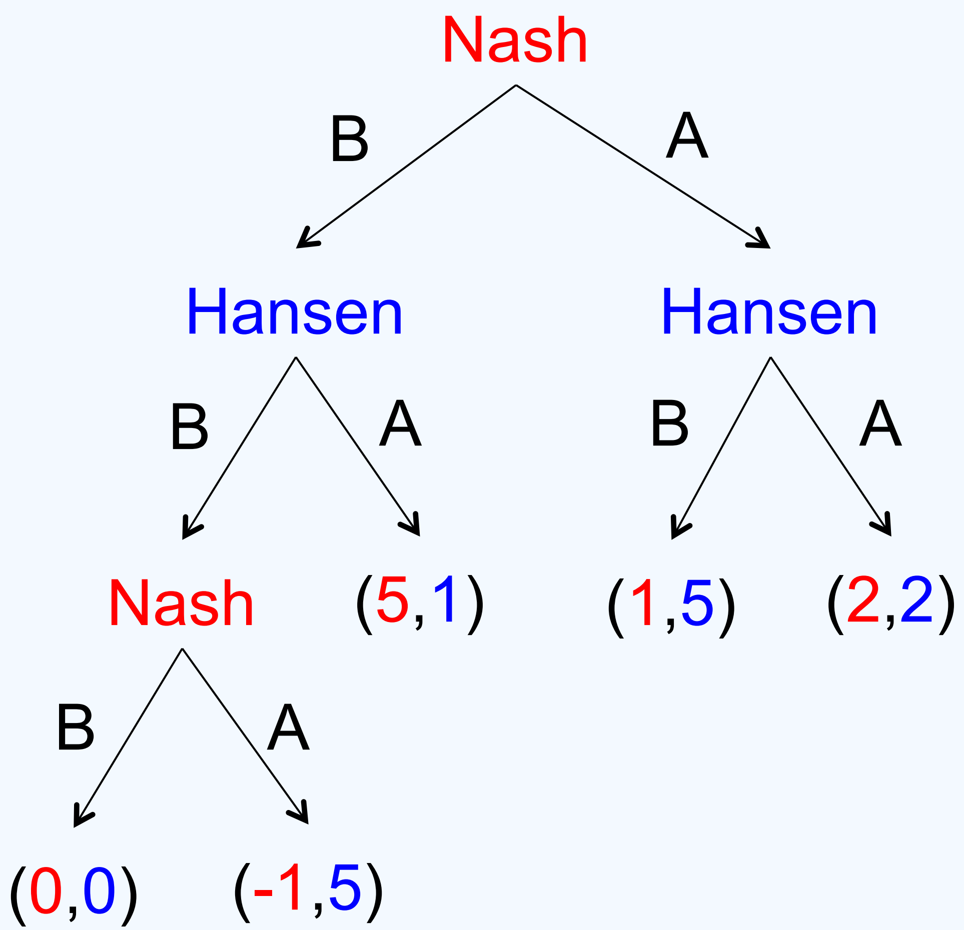

leaves: utilities (Nash,Hansen)

- 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

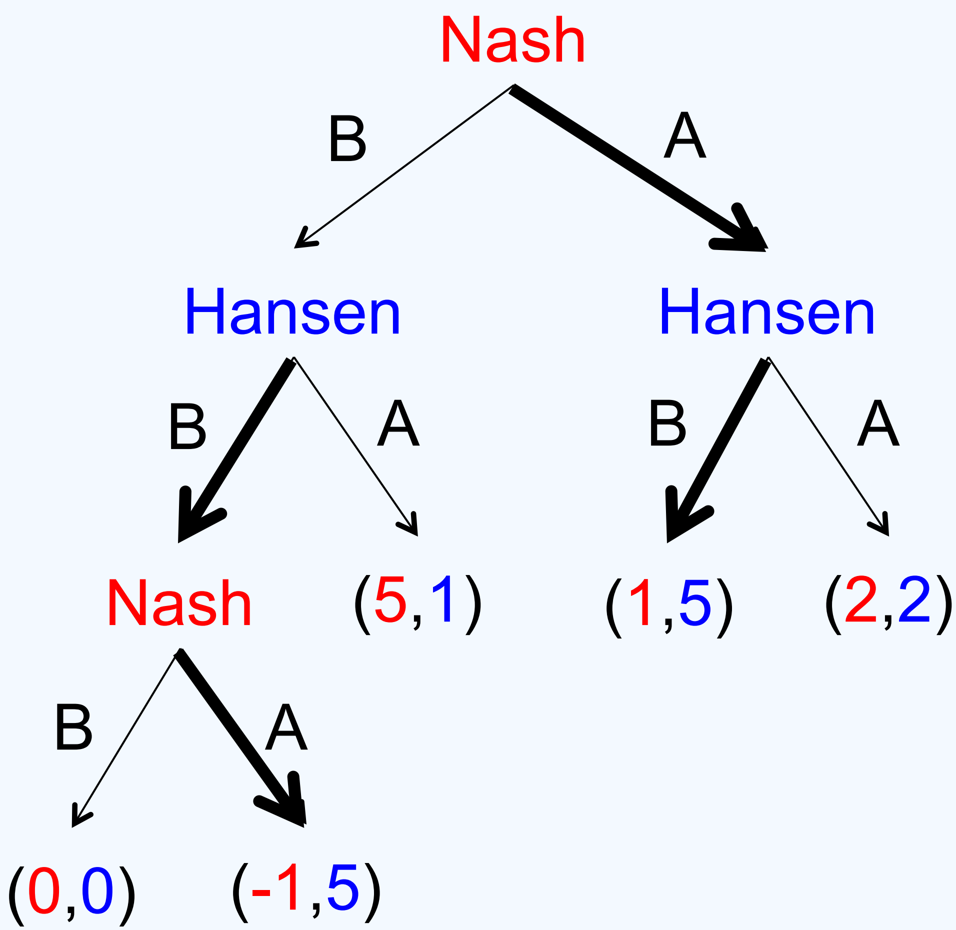

Hansen

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

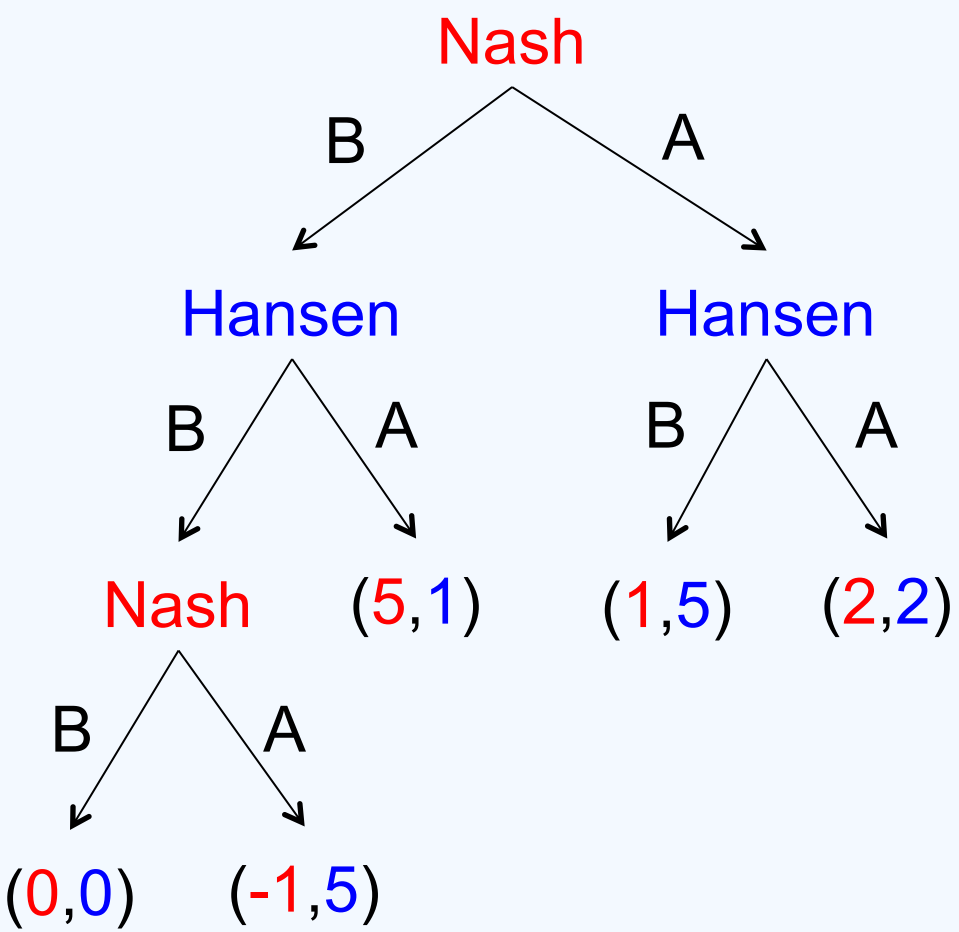

- Usually too many NE
- (pure) SPNE
  - a refinement (special NE)
  - also an NE of any subgame (subtree)

Backward induction

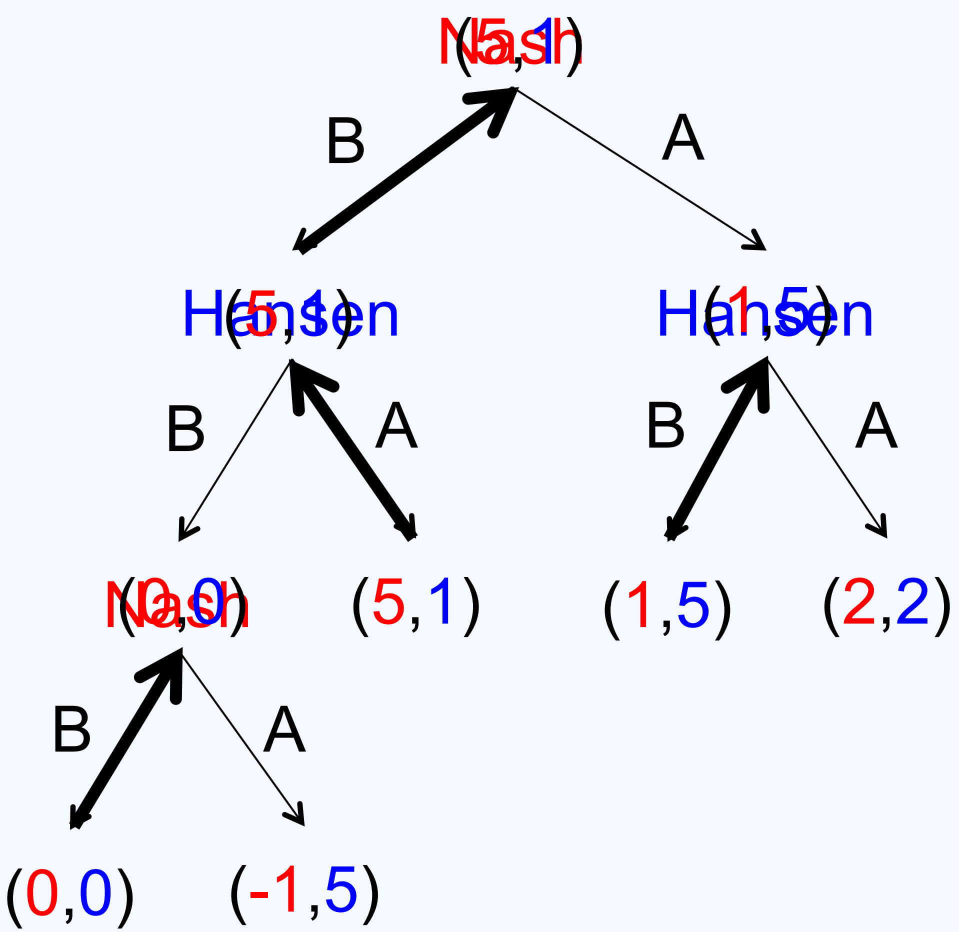

- Determine the strategies bottom-up
- Unique if no ties in the process
- All SPNE can be obtained, if
  - the game is finite
  - complete information
  - perfect information

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

YOU

Bob

Carol

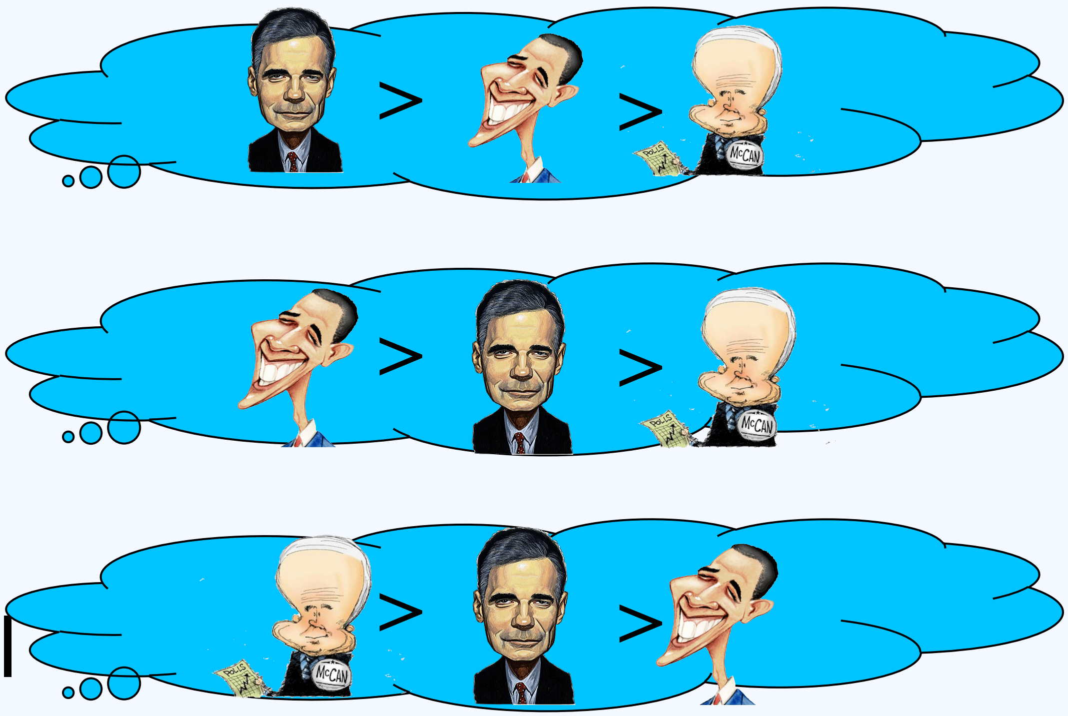

Plurality rule

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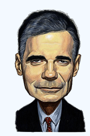, 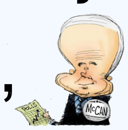}
- Strategies: S_j = Rankings(O)
- Preferences: Rankings(O)
- Mechanism: the plurality rule