Introduction to Social Choice

Lirong Xia

Rensselaer

Jan 29, 2016

Last class: Two goals for social choice

GOAL1: democracy

GOAL2: truth

Change the world: 2011 UK Referendum

- The second nationwide referendum in UK history
  - The first was in 1975
- Member of Parliament election:

  Plurality rule ➔ Alternative vote rule
- 68% No vs. 32% Yes
- In 10/440 districts more voters said yes
  - 6 in London, Oxford, Cambridge, Edinburgh Central, and Glasgow Kelvin
- Why change?
- Why failed?
- Which voting rule is the best?

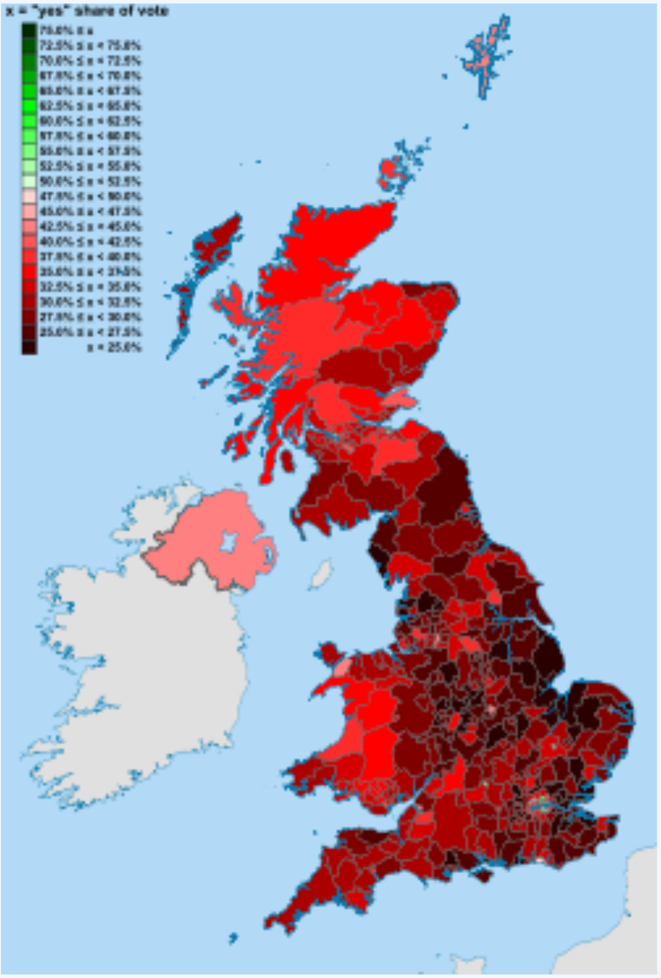

Today's schedule: memory challenge

- Topic: Voting
- We will learn
  - How to aggregate preferences?
    - A large variety of voting rules
  - How to evaluate these voting rules?
    - Democracy: A large variety of criteria (axioms)
    - Truth: an axiom related to the Condorcet Jury theorem
  - Characterize voting rules by axioms
    - impossibility theorems
- Home 1 out

Social choice: Voting

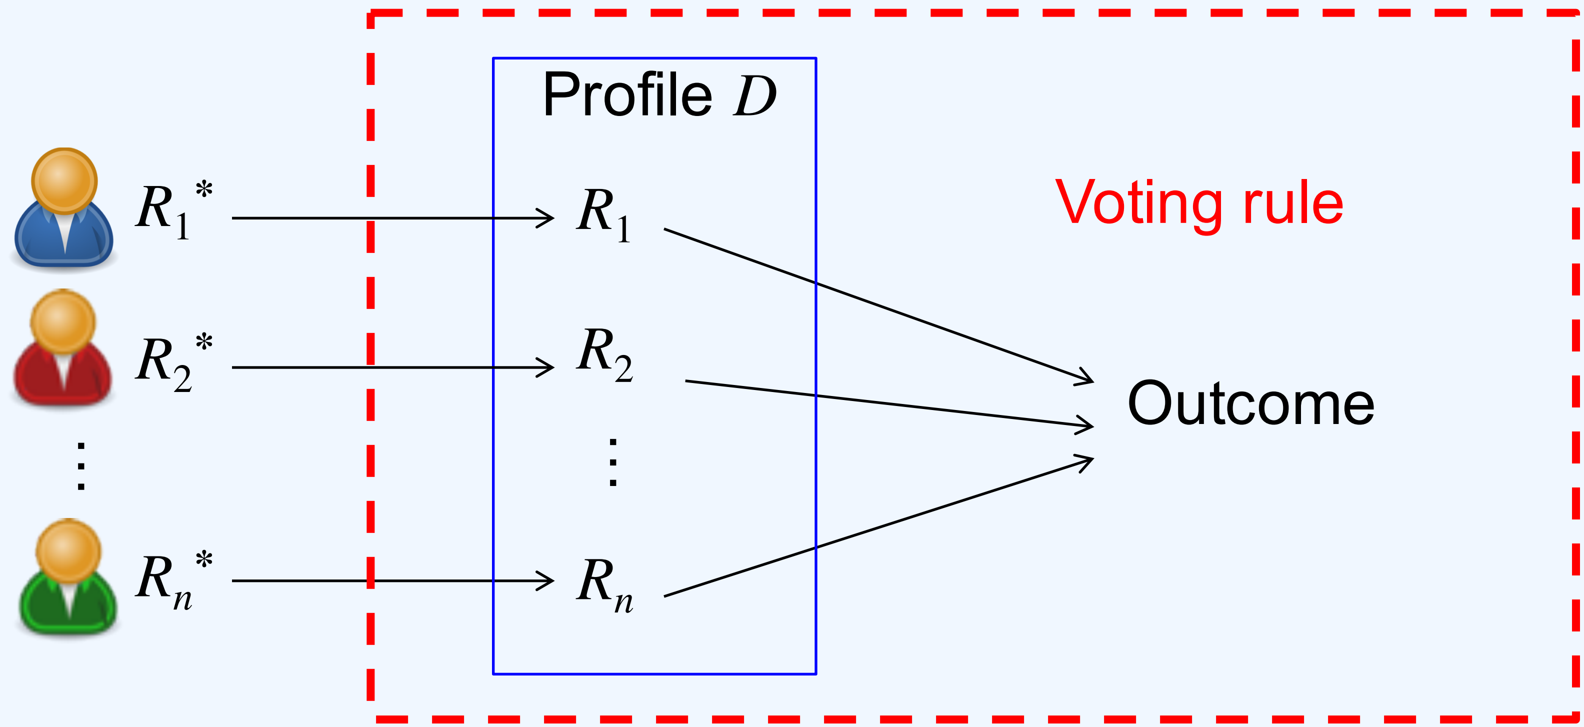

- Agents: n voters, $N=\{1,\ldots,n\}$
- Alternatives: m candidates, $A=\{a_1,\ldots,a_m\}$ or $\{a, b, c, d,\ldots\}$
- Outcomes:
  - winners (alternatives): $O=A$. Social choice function
  - rankings over alternatives: O=Rankings(A). Social welfare function
- Preferences: R_j^* and R_j are full rankings over A
- Voting rule: a function that maps each profile to an outcome

Popular voting rules

(a.k.a. what people have done in the past two centuries)

The Borda rule

$P=\{$  $>$ 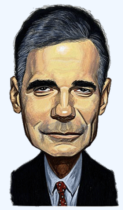 $>$ 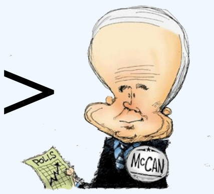 $\times 4$, 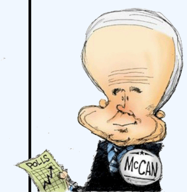 $>$ 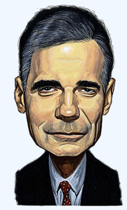 $>$  $\times 3$

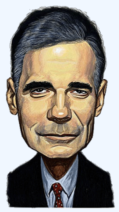 $>$  $>$ 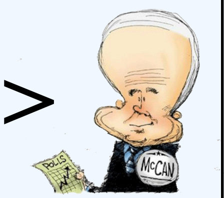 $\times 2$, 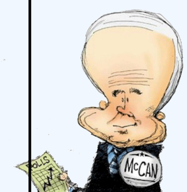 $>$  $>$ 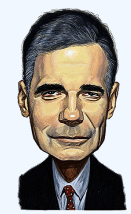 $\times 2\}$

Borda(P)= 

Borda scores  : $2\times 4+4=12$ 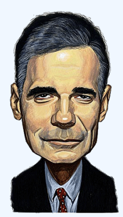 : 2*2+7=11  : 2*5=10

Positional scoring rules

- Characterized by a score vector $s_1,\ldots,s_m$ in non-increasing order
- For each vote R, the alternative ranked in the i-th position gets s_i points
- The alternative with the most total points is the winner
- Special cases
  - Borda: score vector $(m\text{-}1, m\text{-}2, \ldots, 0)$ [French academy of science 1784-1800, Slovenia, Naru]
  - k-approval: score vector $(\underbrace{1\ldots1}_{k}, 0\ldots0)$
  - Plurality: score vector $(1, 0\ldots0)$ [UK, US]
  - Veto: score vector $(1...1, 0)$

Example

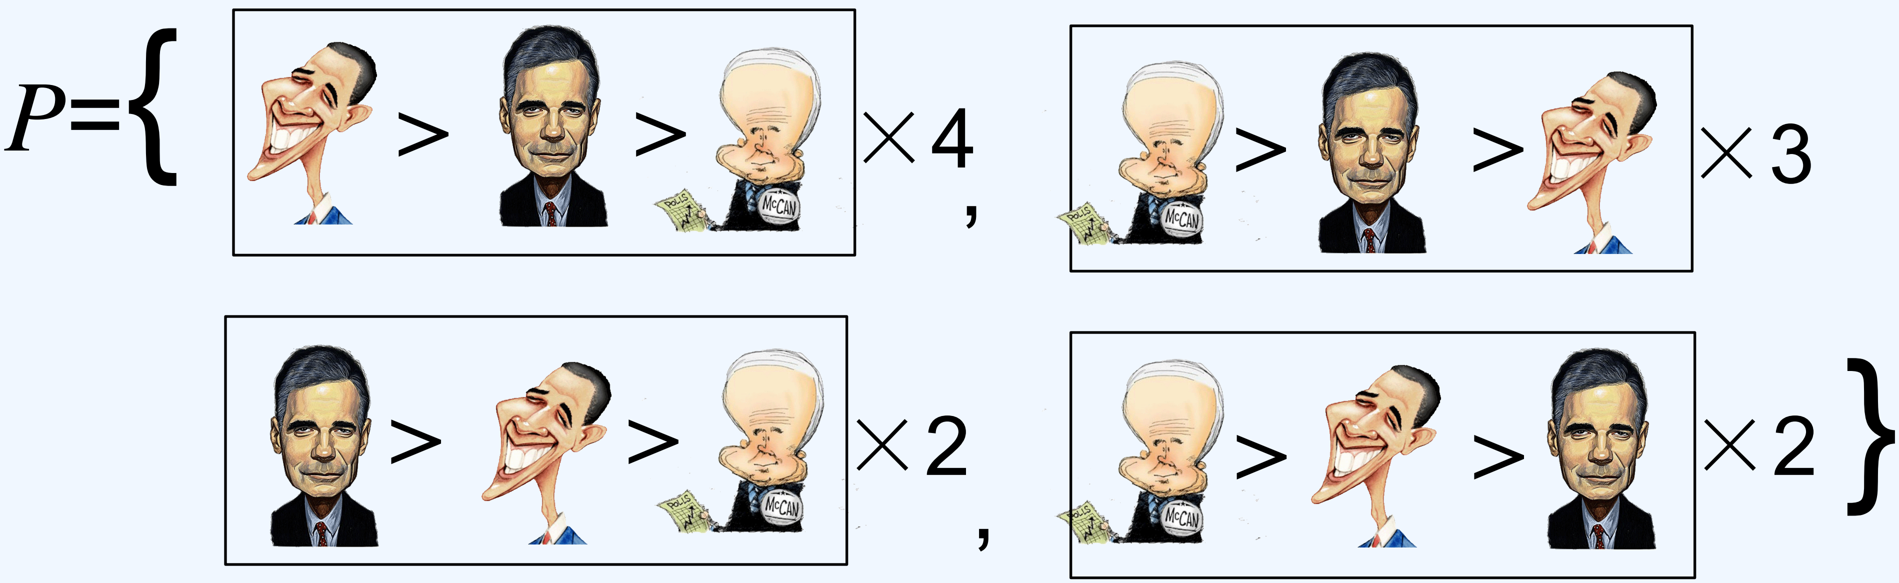

Borda

Plurality
(1- approval)

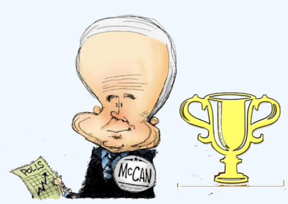

Veto
(2-approval)

Off topic: different winners for the same profile?

Research 101

- Lesson 1: generalization
- Conjecture: for any $m \geq 3$, there exists a profile P such that
  - for different $k \leq m\text{-}1$, k-approval chooses a different winner

Research 102

- Lesson 2: open-mindedness
  - “*If we knew what we were doing, it wouldn't be called research, would it?*”

    ---Albert Einstein

- Homework: Prove or disprove the conjecture

Research 103

- Lesson 3: inspiration in simple cases
- Hint: look at the following example for m=3
  - 3 voters: $a_1 > a_2 > a_3$
  - 2 voters: $a_2 > a_3 > a_1$
  - 1 voter: $a_3 > a_1 > a_2$

It never ends!

- You can apply Lesson 1 again to generalize your observation, e.g.
  - If the conjecture is true, then can you characterize the smallest number of votes in P? How about adding Borda? How about any combination of voting rules?
  - If the conjecture is false, then can you characterize the set of k-approvals to make it true?

Plurality with runoff

- The election has two rounds
  - First round, all alternatives except the two with the highest plurality scores drop out
  - Second round, the alternative preferred by more voters wins
- [used in France, Iran, North Carolina State]

Example: Plurality with runoff

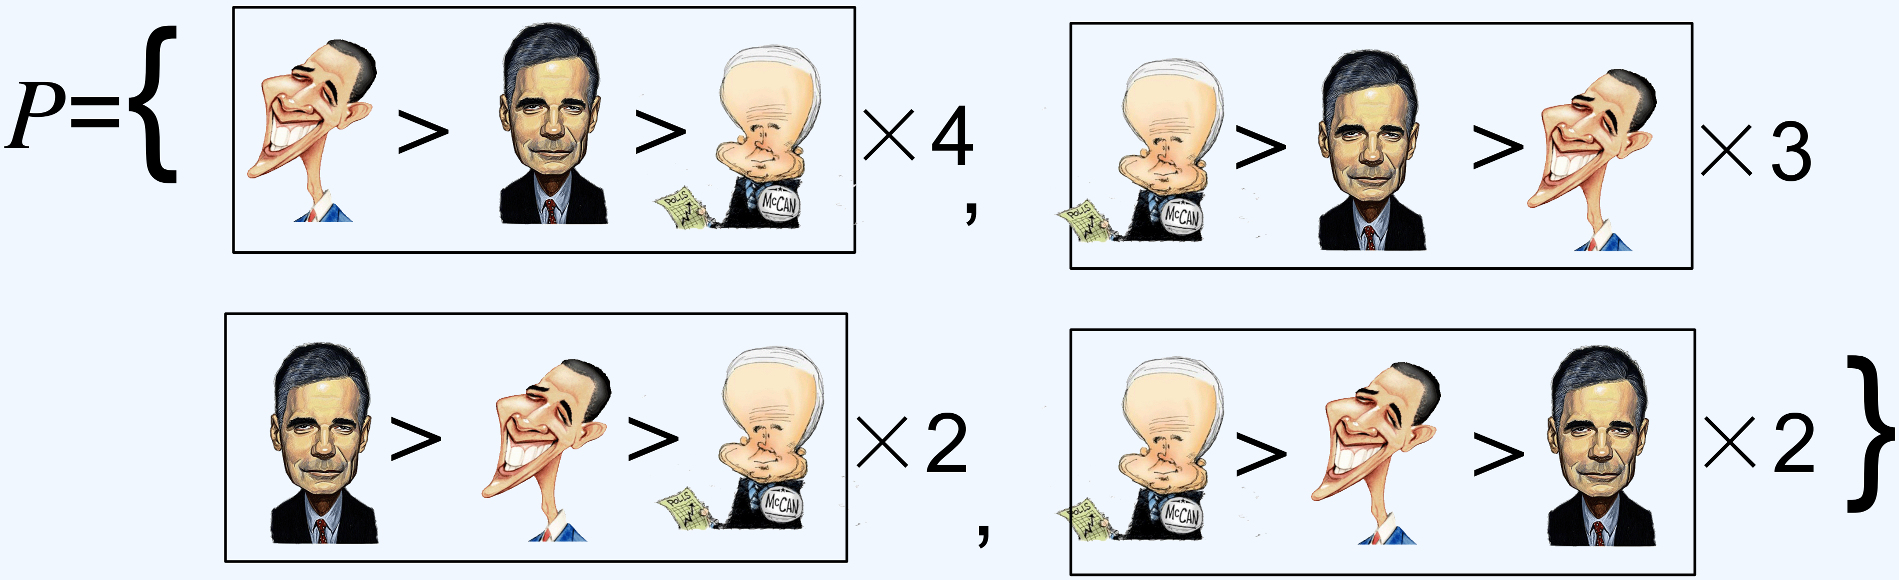

- First round: 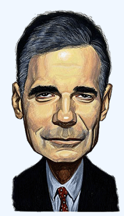 drops out
- Second round:  defeats 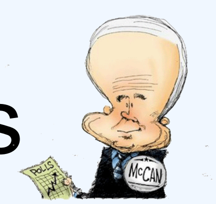

Different from Plurality!

Single transferable vote (STV)

- Also called instant run-off voting or alternative vote
- The election has m-1 rounds, in each round,
  - The alternative with the lowest plurality score drops out, and is removed from all votes
  - The last-remaining alternative is the winner
- [used in Australia and Ireland]

$a > b > c > d$	$d > a > b > c$	$c > d > a > b$	$b > c > d > a$
10	7	6	3


a

Other multi-round voting rules

- Baldwin's rule
  - Borda+STV: in each round we eliminate one alternative with the lowest Borda score
  - break ties when necessary
- Nanson's rule
  - Borda with multiple runoff: in each round we eliminate all alternatives whose Borda scores are below the average
  - [Marquette, Michigan, U. of Melbourne, U. of Adelaide]

Weighted majority graph

- Given a profile P, the weighted majority graph WMG(P) is a weighted directed complete graph (V,E,w) where
  - $V = A$
  - for every pair of alternatives (a, b)
  - $w(a\rightarrow b)$ = #{$a > b$ in P} - #{$b > a$ in P}
  - $w(a\rightarrow b) = -w(b\rightarrow a)$
- WMG (only showing positive edges} might be cyclic
  - Condorcet cycle: { $a>b>c$, $b>c>a$, $c>a>b$}

$$
\begin{tikzcd}
 & a \arrow[dl, "1"'] & \\
b \arrow[rr, "1"'] & & c \arrow[ul, "1"']
\end{tikzcd}
$$

Example: WMG

P=

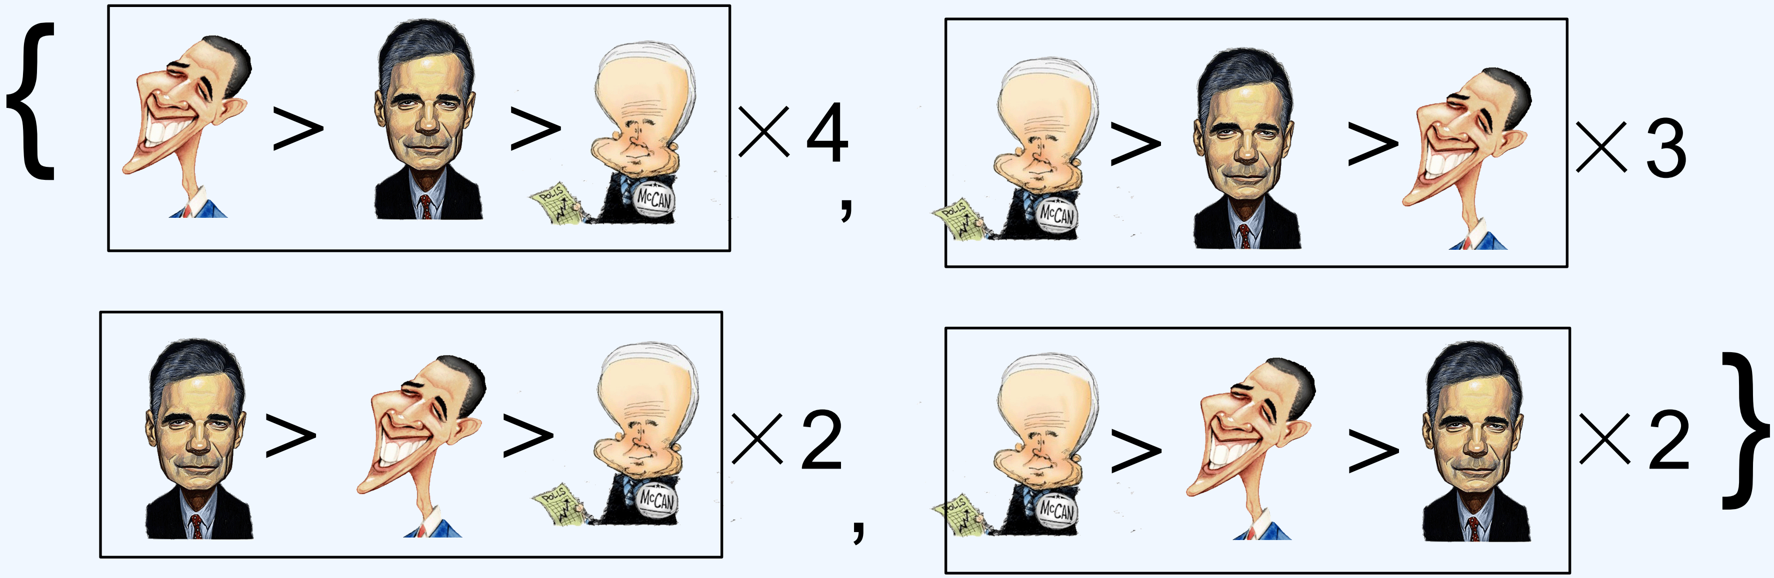

WMG(P) =

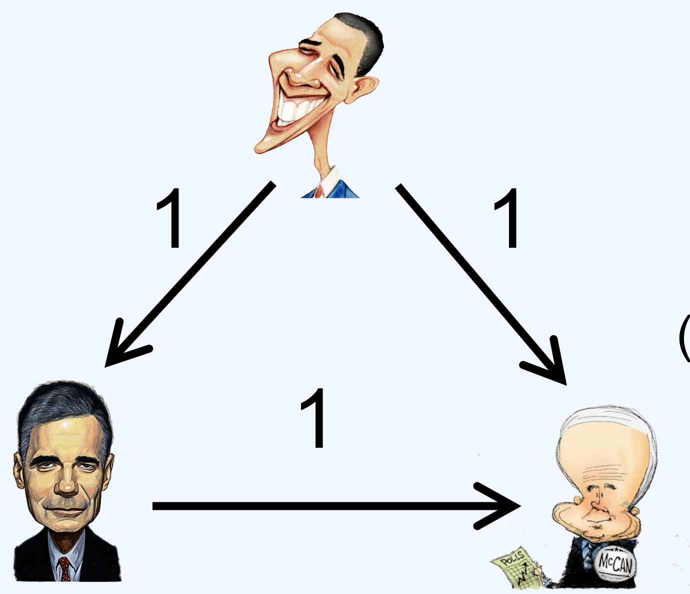

(only showing positive edges)

WGM-based voting rules

- A voting rule r is based on weighted majority graph, if for any profiles P_1, P_2,

  $[\text{WMG}(P_1)=\text{WMG}(P_2)] \Rightarrow [r(P_1)=r(P_2)]$
- WMG-based rules can be redefined as a function that maps {WMGs} to {outcomes}
- Example: Borda is WMG-based
  - Proof: the Borda winner is the alternative with the highest sum over outgoing edges.

The Copeland rule

- The Copeland score of an alternative is its total “pairwise wins”
  - the number of positive outgoing edges in the WMG
- The winner is the alternative with the highest Copeland score
- WMG-based

Example: Copeland

P={ 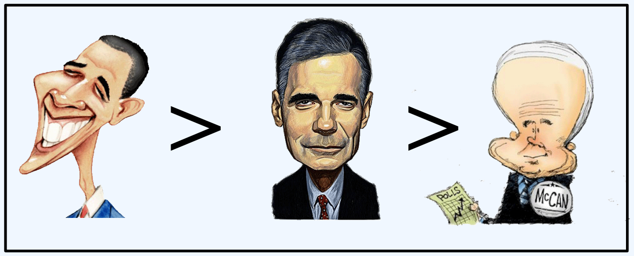

$\times 4$, 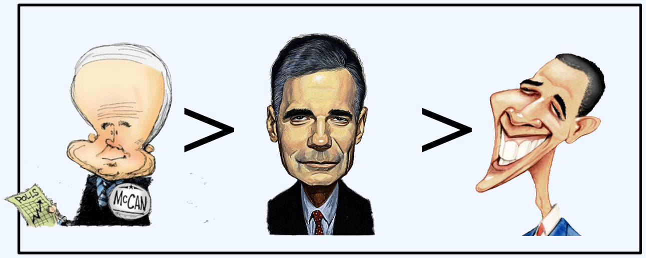

$\times 3$

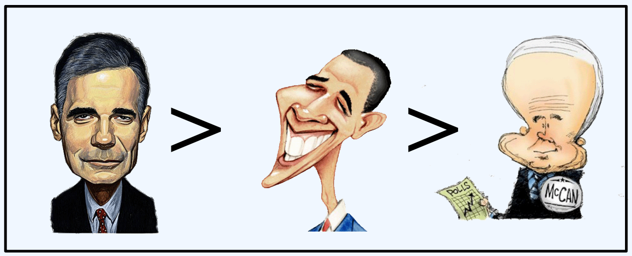

$\times 2$, 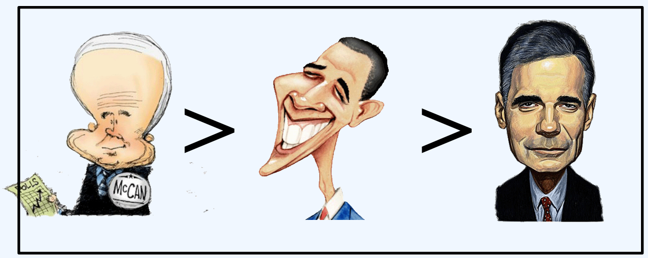

$\times 2$ }

Copeland score:

 : 2

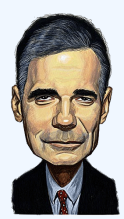 : 1

 : 0

The maximin rule

- A.k.a. Simpson or minimax
- The maximin score of an alternative a is

$$\mathrm{MS}_P(a)=\mathrm{min}_b\ \#\{a > b \text{ in } P\}$$

  - the smallest pairwise defeats
- The winner is the alternative with the highest maximin score
- WMG-based

Example: maximin

P= { 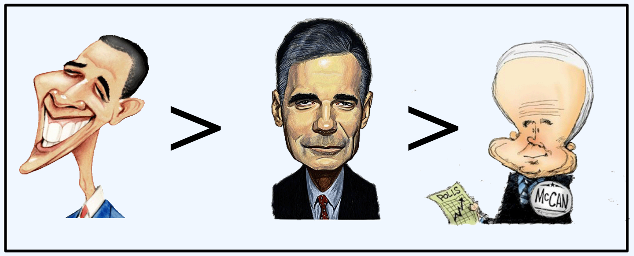 ×4, 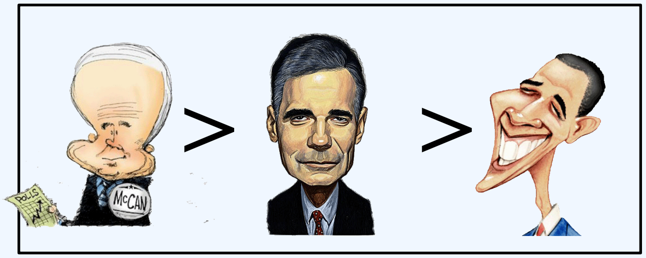 ×3

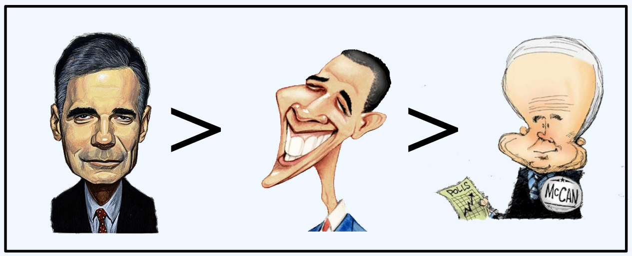 ×2, 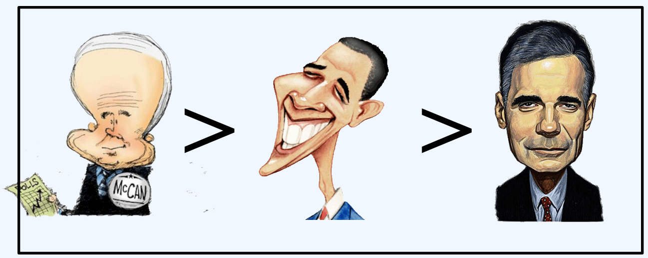 ×2 }

Maximin score:

 : 6

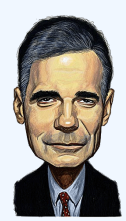 : 5

 : 5

Ranked pairs

- Given the WMG
- Starting with an empty graph G, adding edges to G in multiple rounds
  - In each round, choose the remaining edge with the highest weight
  - Add it to G if this does not introduce cycles
  - Otherwise discard it
- The alternative at the top of G is the winner

Example: ranked pairs

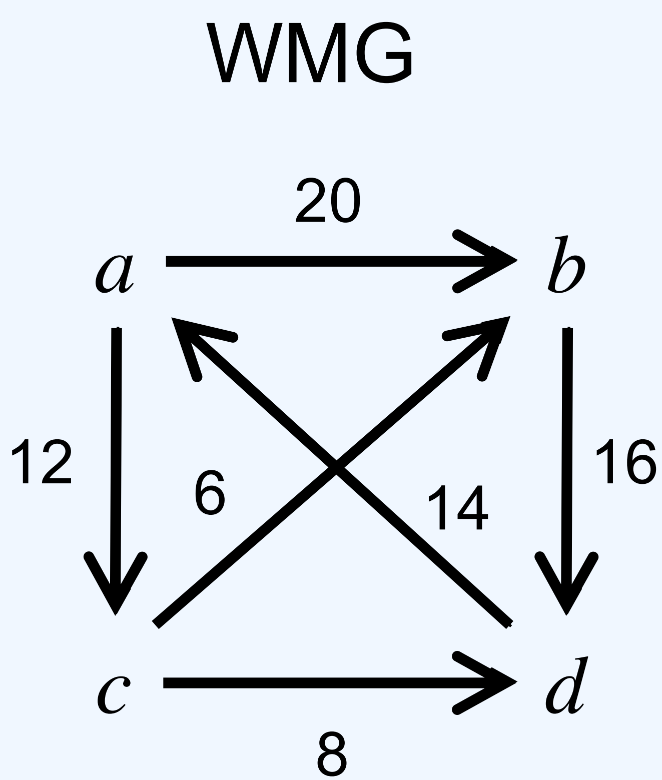

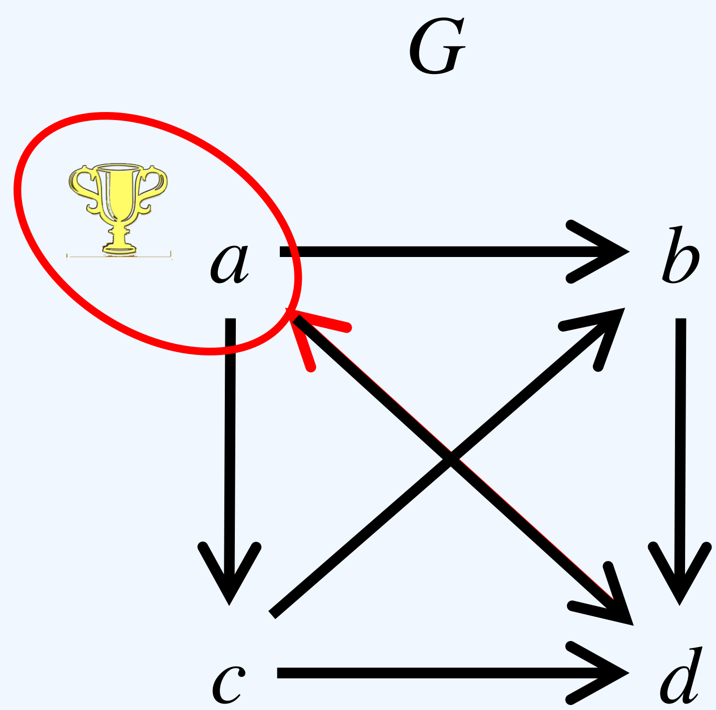

Q1: Is there always an alternative at the "top" of G?

Q2: Does it suffice to only consider positive edges?

Kemeny's rule

- Kendall tau distance
  - $\mathrm{K}(R,W)$= # {different pairwise comparisons}

$$\mathrm{K}(\ b \succ c \succ a\ ,\ a \succ b \succ c\ ) = ?$$

- Kemeny(D)=argmin$_W$ $\mathrm{K}(D,W)$=argmin$_W$ $\Sigma_{R\in D}\mathrm{K}(R,W)$
- For single winner, choose the top-ranked alternative in Kemeny(D)
- [reveals the truth]

Popular criteria for voting rules

(a.k.a. what people have done in the past 60 years)

How to evaluate and compare voting rules?

- No single numerical criteria
  - Utilitarian: the joint decision should maximize the total happiness of the agents
  - Egalitarian: the joint decision should maximize the worst agent's happiness
- Axioms: properties that a "good" voting rules should satisfy
  - measures various aspects of preference aggregation

Fairness axioms

- Anonymity: names of the voters do not matter
  - Fairness for the voters
- Non-dictatorship: there is no dictator, whose top-ranked alternative is always the winner, no matter what the other votes are
  - Fairness for the voters
- Neutrality: names of the alternatives do not matter
  - Fairness for the alternatives

A truth-revealing axiom

- Condorcet consistency: Given a profile, if there exists a Condorcet winner, then it must win
  - The Condorcet winner beats all other alternatives in pairwise comparisons
  - The Condorcet winner only has positive outgoing edges in the WMG
- Why this is truth-revealing?
  - why Condorcet winner is the truth?

The Condorcet Jury theorem [Condorcet 1785]

- Given
  - two alternatives $\{a,b\}$. a: liable, b: not liable
  - $0.5<p<1$,
- Suppose
  - given the ground truth (a or b), each voter's preference is generated i.i.d., such that
    - w/p p, the same as the ground truth
    - w/p $1-p$, different from the ground truth
- Then, as $n\rightarrow\infty$, the probability for the majority of agents' preferences is the ground truth goes to 1

Condorcet's model

[Condorcet 1785]

- Given a "ground truth" ranking W and p>1/2, generate each pairwise comparison in R independently as follows (suppose $c \succ d$ in W)

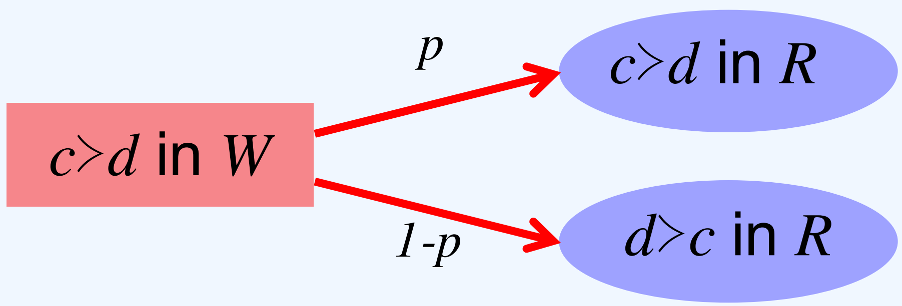

$$\Pr(b \succ c \succ a \mid a \succ b \succ c) = ?\ p(1-p)^2$$

- Its MLE is Kemeny's rule [Young JEP-95]

Truth revealing

Extended Condorcet Jury theorem

- Given
  - A ground truth ranking W
  - $0.5<p<1$,
- Suppose
  - each agent's preferences are generated i.i.d. according to Condorcet's model
- Then, as $n\rightarrow\infty$, with probability that $\rightarrow 1$
  - the randomly generated profile has a Condorcet winner
  - The Condorcet winner is ranked at the top of W
- If r satisfies Condorcet criterion, then as $n\rightarrow\infty$, r will reveal the "correct" winner with probability that $\rightarrow 1$.

Other axioms

- Pareto optimality: For any profile D, there is no alternative c such that every voter prefers c to $r(D)$
- Consistency: For any profiles D_1 and D_2, if $r(D_1)=r(D_2)$, then $r(D_1 \cup D_2)=r(D_1)$
- Monotonicity: For any profile D_1,
  - if we obtain D_2 by only raising the position of $r(D_1)$ in one vote,
  - then $r(D_1)=r(D_2)$
  - In other words, raising the position of the winner won't hurt it

Which axiom is more important?

	Condorcet criterion	Consistency	Anonymity/neutrality, non-dictatorship, monotonicity
Plurality	N	Y	Y
STV (alternative vote)	Y	N	Y

- Some axioms are not compatible with others
- Which rule do you prefer?

An easy fact

- Theorem. For voting rules that selects a single winner, anonymity is not compatible with neutrality
  - proof:

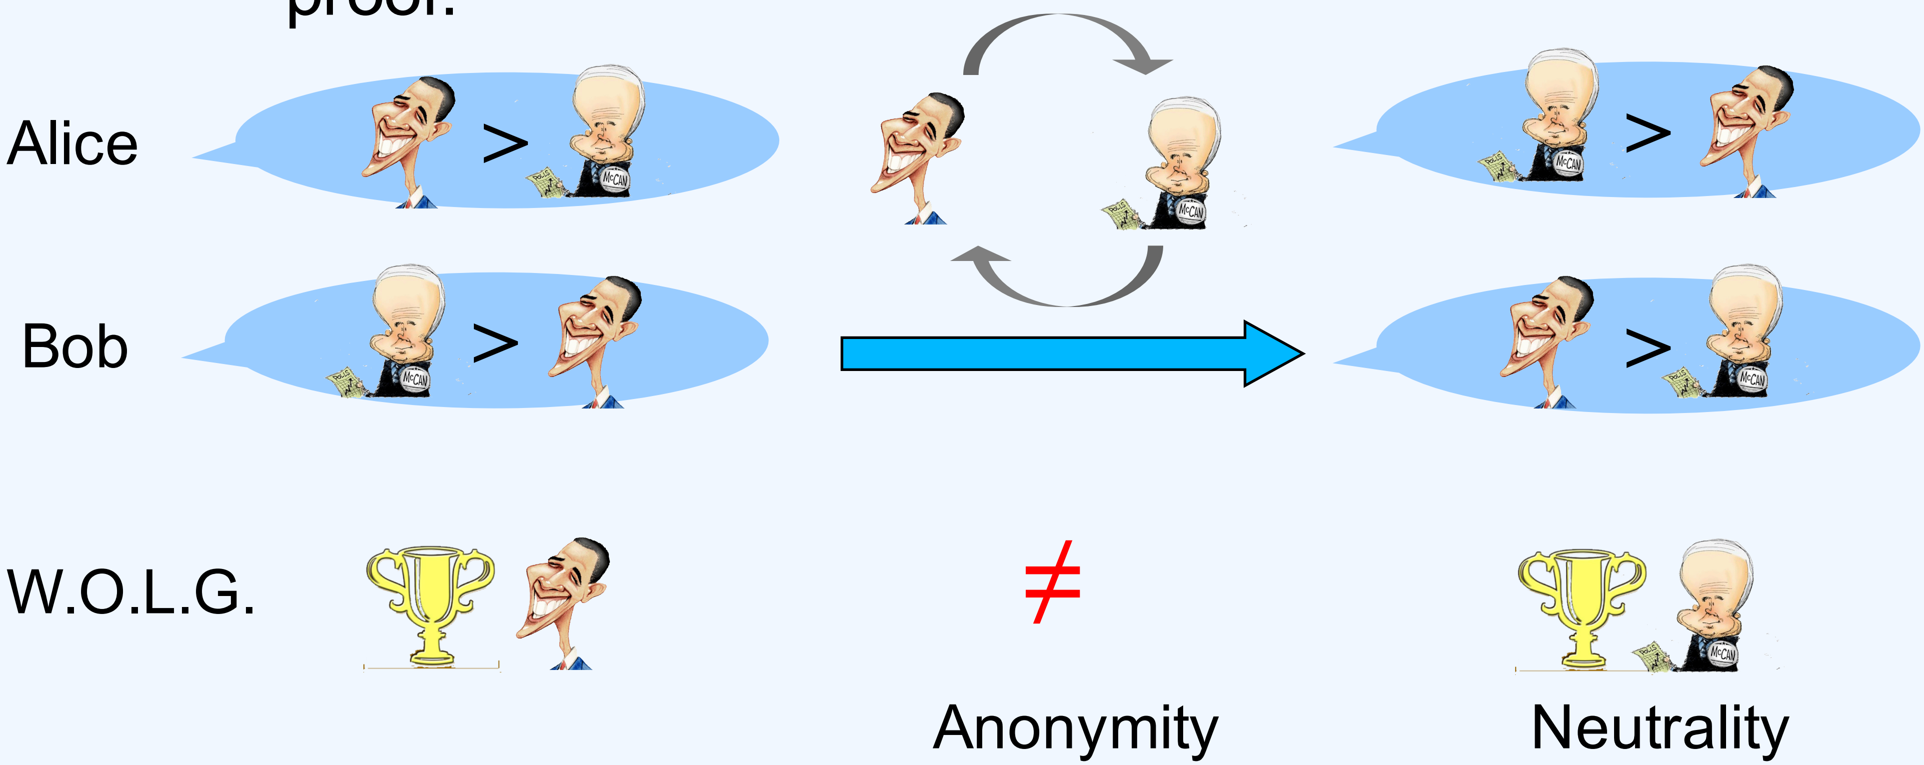

Another easy fact [Fishburn APSR-74]

- Theorem. No positional scoring rule satisfies Condorcet criterion:
  - suppose $s_1 > s_2 > s_3$

3 Voters

2 Voters

1 Voter

1 Voter

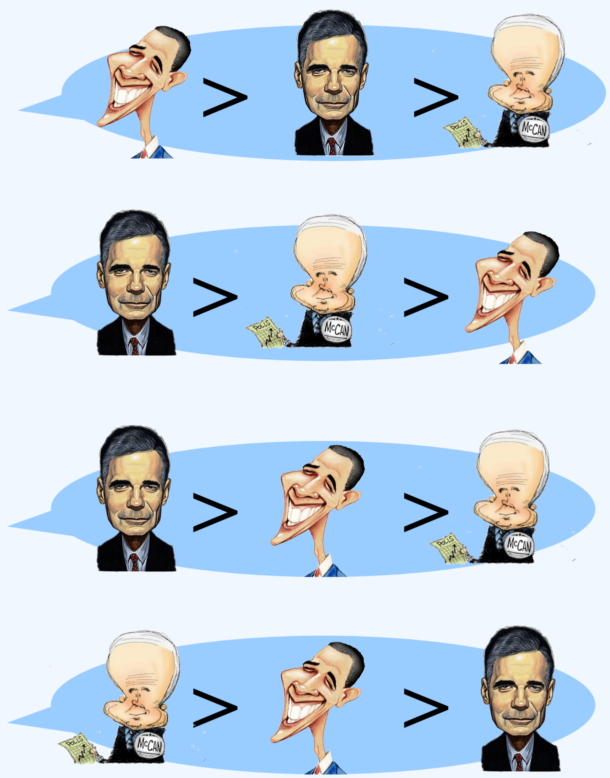

 is the Condorcet winner

CONTRADICTION

 : $3s_1 + 2s_2 + 2s_3$

$\wedge$

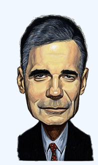 : $3s_1 + 3s_2 + 1s_3$

Arrow's impossibility theorem

- Recall: a social welfare function outputs a ranking over alternatives
- Arrow's impossibility theorem. No social welfare function satisfies the following four axioms
  - Non-dictatorship
  - Universal domain: agents can report any ranking
  - Unanimity: if a>b in all votes in D, then a>b in $r(D)$
  - Independence of irrelevant alternatives (IIA): for two profiles D_1=($R_1,\ldots,R_n$) and D_2=($R_1',\ldots,R_n'$) and any pair of alternatives a and b
    - if for all voter j, the pairwise comparison between a and b in R_j is the same as that in R_j'
    - then the pairwise comparison between a and b are the same in $r(D_1)$ as in $r(D_2)$

Other Not-So-Easy facts

- Gibbard-Satterthwaite theorem
  - Later in the “hard to manipulate” class
- Axiomatic characterization
  - Template: A voting rule satisfies axioms A1, A2, A2 ⇔ if it is rule X
  - If you believe in A1 A2 A3 are the most desirable properties then X is optimal
  - (unrestricted domain+unanimity+IIA) ⇔ dictatorships [Arrow]
  - (anonymity+neutrality+consistency+continuity) ⇔ positional scoring rules [Young SIAMAM-75]
  - (neutrality+consistency+Condorcet consistency) ⇔ Kemeny [Young&Levenglick SIAMAM-78]

Remembered all of these?

- Impressive! Now try a slightly larger tip of the iceberg at wiki

Change the world: 2011 UK Referendum

- The second nationwide referendum in UK history
  - The first was in 1975
- Member of Parliament election:

  Plurality rule ➔ Alternative vote rule
- 68% No vs. 32% Yes
- Why people want to change?
- Why it was not successful?
- Which voting rule is the best?

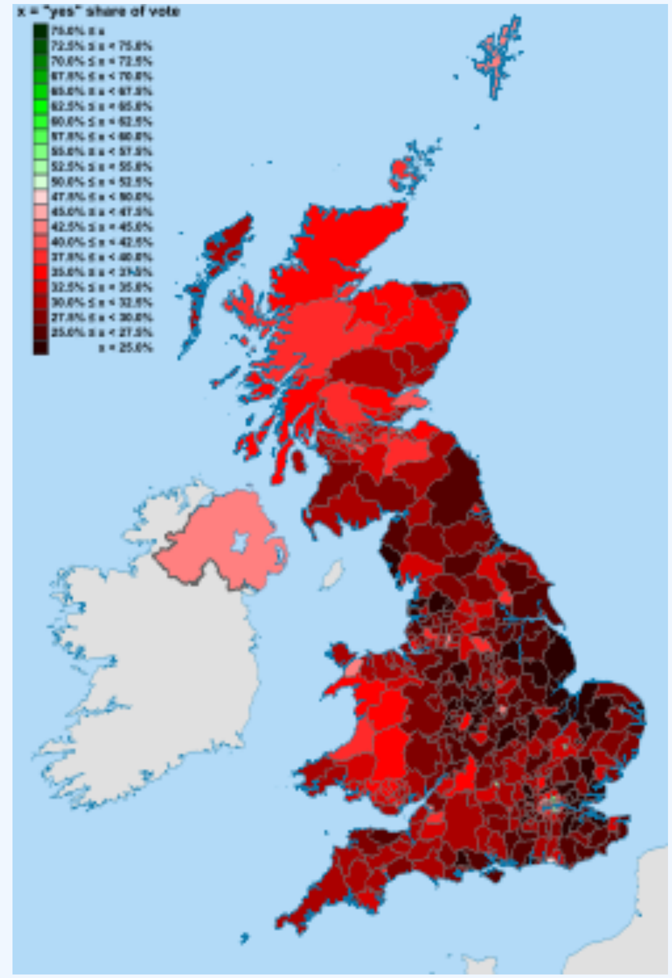

Wrap up

- Voting rules
  - positional scoring rules
  - multi-round elimination rules
  - WMG-based rules
  - A Ground-truth revealing rule (Kemeny's rule)
- Criteria (axioms) for "good" rules
  - Fairness axioms
  - A ground-truth-revealing axiom (Condorcet consistency)
  - Other axioms
- Evaluation
  - impossibility theorems
  - Axiomatic characterization

The reading questions

- What is the problem?
  - social choice
- Why we want to study this problem? How general it is?
  - It is very general and important
- How was problem addressed?
  - by designing voting rules for aggregation and axioms for evaluation and comparisons
- Appreciate the work: what makes the paper nontrivial?
  - No single numerical criterion for evaluation
- Critical thinking: anything you are not very satisfied with?
  - evaluation of axioms, computation, incentives

Looking forward

- How to apply these rules?
  - never use without justification: democracy or truth?
- Preview of future classes
  - Strategic behavior of the voters
    - Game theory and mechanism design
  - Computational social choice
    - Basics of computation
    - Easy-to-compute axiom
    - Hard-to-manipulate axiom
- You can start to work on the first homework!