CSC304 Lecture 16

Voting 2: Gibbard-Satterthwaite Theorem

Recap

- We introduced a plethora of voting rules
  - Plurality
  - Borda
  - Veto
  - k-Approval
  - STV
  - Plurality with runoff
  - Kemeny
  - Copeland
  - Maximin
- All these rules do something reasonable on a given preference profile
  - Only makes sense if preferences are truthfully reported

Recap

- Set of voters $N = \{1, \dots, n\}$
- Set of alternatives A, $|A| = m$
- Voter i has a preference ranking $\succ_i$ over the alternatives

1	2	3
a	c	b
b	a	a
c	b	c

- Preference profile $\vec{\succ}$ = collection of all voter rankings
- Voting rule (social choice function) f
  - Takes as input a preference profile $\vec{\succ}$
  - Returns an alternative $a \in A$

Strategyproofness

- Would any of these rules incentivize voters to report their preferences truthfully?

- A voting rule f is strategyproof if for every
  - preference profiles $\overrightarrow{\succ}$,
  - voter i, and
  - preference profile $\overrightarrow{\succ}'$ such that $\succ_j' = \succ_j$ for all $j \neq i$

  - it is not the case that $f\left(\overrightarrow{\succ}'\right) \succ_i f\left(\overrightarrow{\succ}\right)$

Strategyproofness

- None of the rules we saw are strategyproof!
- Example: Borda Count
  - In the true profile, b wins
  - Voter 3 can make a win by pushing b to the end

Winner
b

1	2	3
b	b	a
a	a	b
c	c	c
d	d	d

1	2	3
b	b	a
a	a	c
c	c	d
d	d	b

Winner
a

Borda's Response to Critics

My scheme is intended only for honest men!

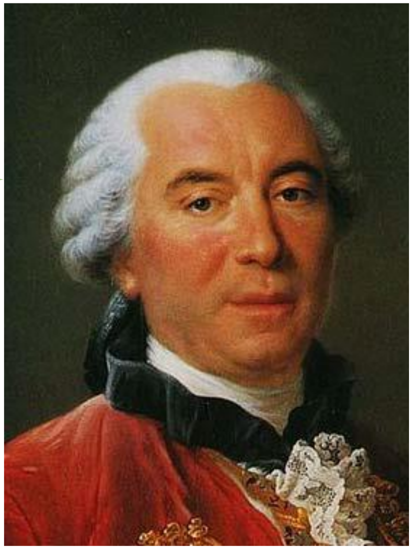

Random 18th century French dude

Strategyproofness

- Are there any strategyproof rules?
  - Sure
- Dictatorial voting rule
  - The winner is always the most preferred alternative of voter i
- Constant voting rule
  - The winner is always the same
- Not satisfactory (for most cases)

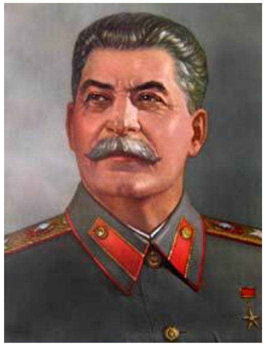

Dictatorship

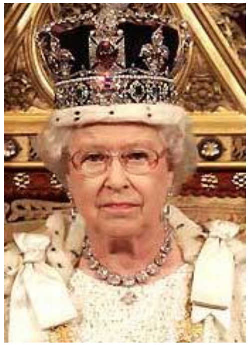

Constant function

Three Requirements

- Strategyproof: Already defined. No voter has an incentive to misreport.
- Onto: Every alternative can win under some preference profile.
- Nondictatorial: There is no voter i such that $f(\overrightarrow{\succ})$ is always the top alternative for voter i.

Gibbard-Satterthwaite

- Theorem: For $m \geq 3$, no deterministic social choice function can be strategyproof, onto, and nondictatorial simultaneously ☹
- Proof: We will prove this for $n = 2$ voters.
  - Step 1: Show that SP is equivalent to "strong monotonicity" [HW 3?]
  - Strong Monotonicity (SM): If $f(\overrightarrow{\succ}) = a$, and $\overrightarrow{\succ}'$ is such that $\forall i \in N, x \in A$: $a \succ_i x \Rightarrow a \succ_i' x$, then $f(\overrightarrow{\succ}') = a$.
    - If a is winning, and the votes change so that in each vote, a still defeats each alternative it defeated before, then a should still win.

Gibbard-Satterthwaite

- Theorem: For $m \geq 3$, no deterministic social choice function can be strategyproof, onto, and nondictatorial simultaneously ☹

- Proof: We will prove this for $n = 2$ voters.
  - Step 2: Show that SP+onto implies "Pareto optimality" [HW 3?]
  - Pareto Optimality (PO): If $a \succ_i b$ for all $i \in N$, then $f(\overrightarrow{\succ}) \neq b$.
    - If there is a different alternative that *everyone* prefers, your choice is not Pareto optimal (PO).

Gibbard-Satterthwaite

- Proof for n=2: Consider a problem instance $I(a, b)$

$I(a, b)$

$\succ_1$	$\succ_2$
a	b
b	a
Arbitrary	Arbitrary

Say $f(\succ_1, \succ_2) = a$

- PO: $f(\succ_1, \succ_2) \in \{a, b\}$

$\succ_1$	$\succ_2'$
a	b
b	Same as before
Same as before	a

$f(\succ_1, \succ_2') = a$

- PO: $f(\succ_1, \succ_2') \in \{\mathrm{a}, \mathrm{b}\}$
- SP: $f(\succ_1, \succ_2') \neq b$

$\succ_1''$	$\succ_2''$
a	A N Y
A N Y	

$f(\succ'') = a$

- Due to strong monotonicity

Gibbard-Satterthwaite

- Proof for n=2:
  - If f outputs a on instance $I(a, b)$, voter 1 can get a elected whenever she puts a first.
    - In other words, voter 1 becomes dictatorial for a.
    - Denote this by $D(1, a)$.
  - If f outputs b on $I(a, b)$
    - Voter 2 becomes dictatorial for b, i.e., we have $D(2, b)$.

- For every pair of alternatives (a, b), at least one of $D(1, a)$ and $D(2, b)$ holds.

Gibbard-Satterthwaite

- Proof for n=2:
  - Take a pair (a^*, b^*)
  - Suppose wlog that $D(1, a^*)$ holds
  - Then, we show that voter 1 is a dictator, i.e., $D(1, x)$ holds for every other x as well
  - Take $x \neq a^*$
  - Because $|A| \geq 3$, there exists $y \in A \backslash \{a^*, x\}$.
  - For (x, y), at least one of $D(1, x)$ and $D(2, y)$ holds
  - But $D(2, y)$ is incompatible with $D(1, a^*)$
    - Who wins if voter 1 puts a^* first and voter 2 puts y first?
  - Thus, we have $D(1, x)$, as required. ■

Circumventing G-S

- Randomization
  - Gibbard characterized all randomized strategyproof rules
  - Somewhat better, but still too restrictive
- Restricted preferences
  - Median for facility location (more generally, for single-peaked preferences on a line)
  - Will see other such settings later
- Money
  - E.g., VCG is nondictatorial, onto, and strategyproof, but charges payments to agents

Circumventing G-S

- Equilibrium analysis
  - Maybe good alternatives still win under Nash equilibria?
- Lack of information
  - Maybe voters don't know how other voters will vote?

Circumventing G-S

- Computational complexity (Bartholdi et al.)
  - Maybe the rule is manipulable, but it is NP-hard to find a successful manipulation?
  - Groundbreaking idea! NP-hardness can be good!!

- Not NP-hard: plurality, Borda, veto, Copeland, maximin, …
- NP-hard: Copeland with a peculiar tie-breaking, STV, ranked pairs, …

Circumventing G-S

- Computational complexity
  - Unfortunately, NP-hardness just says it is hard for *some worst-case instances.*
  - What if it is actually easy for most practical instances?
  - Many rules admit polynomial time manipulation algorithms for fixed #alternatives (m)
  - Many rules admit polynomial time algorithms that find a successful manipulation on almost all profiles (the fraction of profiles converges to 1)

- Interesting open problems regarding the design of voting rules that are hard to manipulate on average

Social Choice

- Let’s forget incentives for now.
- Even if voters reveal their preferences truthfully, we do not have a “right” way to choose the winner.

- Who is the right winner?
  - On profiles where the prominent voting rules have different outputs, all answers seem reasonable [HW3].

Axiomatic Approach

- Define axiomatic properties we may want from a voting rule

- We already defined some:
  - Majority consistency
  - Condorcet consistency
  - Ontoness
  - Strategyproofness
  - Strong monotonicity (equivalent to SP)
  - Pareto optimality

Axiomatic Approach

- We will see four more:
  - Unanimity
  - Weak monotonicity
  - Consistency (!)
  - Independence of irrelevant alternatives (IIA)
- Problem?
  - Cannot satisfy many interesting combinations of properties
  - Arrow’s impossibility result
  - Other similar impossibility results

Other Approaches?

- Statistical
  - There exists an objectively true answer
    - E.g., say the question is: “Sort the candidates by the #votes they will receive in an upcoming election.”
  - Every voter is trying to estimate the true ranking
  - Goal is to find the most likely ground truth given votes

- Utilitarian
  - Back to “numerical” welfare maximization, but we still ask voters to only report ranked preferences
    - $a \succ_i b \succ_i c$ simply means $v_i(a) \geq v_i(b) \geq v_i(c)$
  - How well can we maximize welfare subject to such partial information?