CSC2556

Lecture 4

Voting IV

Approaches to Voting

Approaches to Voting

- What does an approach give us?
  - A way to compare voting rules
  - Hopefully find a uniquely “optimal” voting rule

- Various approaches in the literature
  - Axiomatic approach
  - Distance rationalizability approach
  - Statistical approach
  - Implicit utilitarian approach
  - …

Axiomatic Approach

Axiomatic Approach

- Axiom:
  - A requirement that the voting rule must behave in a certain way
- Goal:
  - Define a set of reasonable axioms, and search for voting rules that satisfy them together
  - Ultimate hope: a unique voting rule satisfies the set of axioms simultaneously!
  - What often happens: no voting rule satisfies the axioms together ☹

We have already seen axioms!

- Condorcet consistency
- Majority consistency
- Strategyproofness
- Ontoness
- Non-dictatorship
- Strong monotonicity
- Pareto optimality

Axiomatic Approach

- Some axioms are weak and satisfied by all natural rules
  - Unanimity:
    - If all voters have the same top choice, that alternative is the winner.

$$\left(top(\succ_i) = a\ \forall i \in N\right) \Rightarrow f\left(\vec{\succ}\right) = a$$

  - Q: How does this compare to Pareto optimality?
  - Pareto optimality is weak but still violated by natural voting methods like voting trees

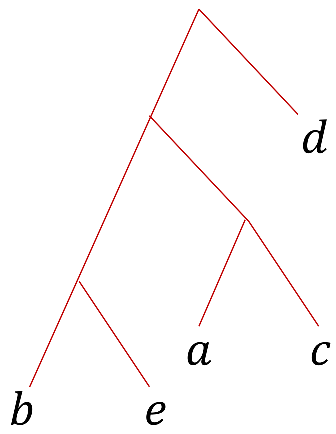

Axiomatic Approach

- Anonymity:
  - Permuting the votes does not change the winner
  - In other words, voter identities don't matter
  - Example: these two profiles must have the same winner:
    {voter 1: $a \succ b \succ c$, voter 2: $b \succ c \succ a$}
    {voter 1: $b \succ c \succ a$, voter 2: $a \succ b \succ c$}

- Neutrality:
  - Permuting alternative names just permutes the winner accordingly
  - Example:
    - Say a wins on {voter 1: $a \succ b \succ c$, voter 2: $b \succ c \succ a$}
    - We permute all names: $a \to b$, $b \to c$, and $c \to a$
    - New profile: {voter 1: $b \succ c \succ a$, voter 2: $c \succ a \succ b$}
    - Then, the new winner must be b

Axiomatic Approach

- Neutrality is tricky for deterministic rules
  - Incompatible with anonymity
    - Consider the profile {voter 1: $a \succ b$, voter 2: $b \succ a$}
    - Without loss of generality, say a wins
    - Imagine a different profile: {voter 1: $b \succ a$, voter 2: $a \succ b$}
      - Neutrality ⇒ we exchanged $a \leftrightarrow b$, so winner must be b
      - Anonymity ⇒ we exchanged the votes, so winner must be a

- We usually only require neutrality for…
  - Randomized rules: E.g., a rule could satisfy both by choosing a and b as the winner with probability ½ each, on both profiles
  - Deterministic rules that return a set of tied winners: E.g., a rule could return $\{a, b\}$ as tied winners on both profiles.

Axiomatic Approach

- Consistency: If a is the winner on two profiles, it must be the winner on their union.

$$f\left(\overrightarrow{\succ}_1\right) = a \wedge f\left(\overrightarrow{\succ}_2\right) = a \Rightarrow f\left(\overrightarrow{\succ}_1 + \overrightarrow{\succ}_2\right) = a$$

  - Example: $\overrightarrow{\succ}_1 = \{a \succ b \succ c\}$, $\overrightarrow{\succ}_2 = \{a \succ c \succ b, b \succ c \succ a\}$
  - Then, $\overrightarrow{\succ}_1 + \overrightarrow{\succ}_2 = \{a \succ b \succ c, a \succ c \succ b, b \succ c \succ a\}$

- Theorem [Young '75]:
  - Subject to mild requirements, a voting rule is consistent if and only if it is a positional scoring rule!

Axiomatic Approach

- **Weak monotonicity:** If a is the winner, and a is "pushed up" in some votes, a remains the winner.
  - $f(\vec{\succ}) = a \Rightarrow f(\vec{\succ}') = a$, where
    - $b \succ_i c \Leftrightarrow b \succ_i' c, \forall i \in N,\ b, c \in A \backslash \{a\}$ (Order of others preserved)
    - $a \succ_i b \Rightarrow a \succ_i' b, \forall i \in N,\ b \in A \backslash \{a\}$ (a only improves)

- Contrast with strong monotonicity
  - SM requires $f(\vec{\succ}') = a$ even if $\vec{\succ}'$ only satisfies the 2nd condition
  - Too strong; only satisfied by dictatorial or non-onto rules [GS Theorem]

Axiomatic Approach

- Weak monotonicity is satisfied by most voting rules
  - Popular exceptions: STV, plurality with runoff

- But violation of weak monotonicity helps STV be hard to manipulate
  - Theorem [Conitzer-Sandholm ‘06]:
    “Every weakly monotonic voting rule is easy to manipulate on average.”

Axiomatic Approach

- STV violates weak monotonicity

7 voters	5 voters	2 voters	6 voters
a	b	b	c
b	c	c	a
c	a	a	b

7 voters	5 voters	2 voters	6 voters
a	b	a	c
b	c	b	a
c	a	c	b

- First c, then b eliminated
- Winner: a

- First b, then a eliminated
- Winner: c

Axiomatic Approach

- Arrow's Impossibility Theorem
  - Applies to social welfare functions (profile → ranking)
  - Independence of Irrelevant Alternatives (IIA): If the preferences of all voters between a and b are unchanged, the social preference between a and b should not change
  - Pareto optimality: If all prefer a to b, then the social preference should be $a \succ b$
  - Theorem: IIA + Pareto optimality ⇒ dictatorship
- Interestingly, automated theorem provers can also prove Arrow's and GS impossibilities!

Axiomatic Approach

- Polynomial-time computability
  - Can be thought of as a desirable axiom
  - Two popular rules which attempt to make the pairwise comparison graph acyclic by inverting edges are NP-hard to compute:
    - Kemeny's rule: invert edges with minimum total weight
    - Slater's rule: invert minimum number of edges
  - Both rules can be implemented by straightforward integer linear programs
    - For small instances (say, up to 20 alternatives), NP-hardness isn't a practical concern.

Distortion Approach

Distortion Approach

- A quantitative approach to voting

- Three key steps:
  1. Assume that voters' *ranked* preferences are induced by their *underlying numerical utilities/costs* for the alternatives
  2. Set the goal (e.g. choose the alternative maximizing the sum of voters' utilities a.k.a. the social welfare)
  3. Select an alternative that *approximately optimizes* the goal as best as possible (the approximation ratio is called *distortion*)

- Increasingly popular in recent years

Distortion Approach

- Pros:
  - Uses minimal subjective assumptions
    - Need to assume underlying cardinal utilities/costs
    - Need to set a goal, which can be a subjective choice
  - Yields a uniquely optimal voting rule
- Cons:
  - Optimal rule often doesn't have an intuitive formula that humans can comprehend
  - Optimal rule can sometimes be difficult to compute

Utilitarian Framework

- Underlying utility profile $\vec{u} = (u_1, \dots, u_n)$
  - $u_i(a)$ = utility of voter i for alternative a
  - Normalization: $\sum_a u_i(a) = 1$ for all voters i

- Social welfare $sw(a, \vec{u}) = \sum_i u_i(a)$

- Ideal goal: choose $a^* \in \text{argmax}_a \, sw(a, \vec{u})$
  - If we observe $\vec{u}$, then we can compute a^* easily
  - However, we do not get to observe $\vec{u}$ directly

Utilitarian Framework

- Observed preference profile $\vec{\succ} = (\succ_1, \dots, \succ_n)$
  - Each voter i reports $\succ_i$ consistent with u_i
    - $u_i(a) > u_i(b) \Rightarrow a \succ_i b$
    - The voter can break ties arbitrarily

- Realistic goal: approximately maximize social welfare
  - Distortion of voting rule f

$$\text{dist}(f) = \sup_{\vec{u}} \frac{\max_a \text{sw}(a, \vec{u})}{\text{sw}\left(f(\vec{\succ}), \vec{u}\right)}$$

    - Implicit max over all possible $\vec{\succ}$ that can be induced from $\vec{u}$
    - If f is randomized, we need $E\left[\text{sw}\left(f(\vec{\succ}), \vec{u}\right)\right]$

Deterministic Rules

- Theorem [Caragiannis et al. '16]:
  Given ranked preferences, the optimal deterministic voting rule has $\Theta(m^2)$ distortion.

- Proof (lower bound):
  - High-level approach:
    - Take an arbitrary voting rule f
    - Construct a preference profile $\vec{\succ}$
    - Let f choose a winner a on $\vec{\succ}$
    - Reveal a utility profile $\vec{u}$ which could have induced $\vec{\succ}$ but on which a is $\Omega(m^2)$ times worse than the optimal alternative

Deterministic Rules

- Proof (lower bound):
  - Let f be any voting rule
  - Consider the preference profile $\vec{\succ}$ given on the right

${}^{n}/_{(m-1)}$ voters per column			
a_1	a_2	…	a_{m-1}
a_m	a_m	…	a_m
$\vdots$	$\vdots$	$\vdots$	$\vdots$

  - Case 1: $f(\vec{\succ}) = a_m$:
    - Infinite distortion. WHY?
  - Case 2: $f(\vec{\succ}) = a_i$ for some $i < m$:
    - Bad utility profile $\vec{u}$ consistent with $\vec{\succ}$:
      - Voters in column i have utility $1/m$ for every alternative
      - All other voters have utility $1/2$ for their top two alternatives
    - $\text{sw}(a_i, \vec{u}) = \frac{n}{m-1} \cdot \frac{1}{m}$, $\text{sw}(a_m, \vec{u}) \geq \frac{n-n/(m-1)}{2}$
    - Distortion = $\Omega(m^2)$

Deterministic Rules

- Proof (upper bound):
  - Claim: Plurality achieves $O(m^2)$ distortion
  - Suppose plurality winner is a.
    - At least n/m voters have a as their top choice
    - A voter has utility at least $1/m$ for their top choice
  - $sw(a, \vec{u}) \geq n/m^2$
  - $sw(a^*, \vec{u}) \leq n$ for every alternative a^*
  - $O(m^2)$ distortion

Randomized Rules

- Theorem [Boutilier et al. '12]:
  Given ranked preferences, the optimal randomized voting rule has distortion $O(\sqrt{m} \cdot \log^* m)$, $\Omega(\sqrt{m})$.

- Proof (lower bound):
  - Same high-level approach:
    - Take an arbitrary *randomized* voting rule f
    - Construct a preference profile $\vec{\succ}$
    - Let f choose a distribution over alternatives p
    - Reveal a utility profile $\vec{u}$ which could have induced $\vec{\succ}$ but on which the expected social welfare under p is $\Omega(\sqrt{m})$ times worse than the optimal social welfare

Randomized Rules

- Proof (lower bound):
  - Let f be an arbitrary rule
  - Consider $\vec{\succ}$ on the right:
    - $\sqrt{m}$ special alternatives
  - f must choose at least one special alternative (say a^*) w.p. at most $1/\sqrt{m}$

${}^{n}/_{\sqrt{m}}$ voters per column			
a_1	a_2	…	$a_{\sqrt{m}}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$

  - Bad utility profile $\vec{u}$ consistent with :
    - All voters ranking a^* first give utility 1 to a^*
    - All other voters give utility $1/m$ to each alternative
    - $\frac{n}{\sqrt{m}} \leq \text{sw}(a^*, \vec{u}) \leq \frac{2n}{\sqrt{m}}$
    - $sw(a, \vec{u}) \leq n/m$ for every other a
    - Distortion lower bound: $\sqrt{m}/3$ (proof on the board!)

Randomized Rules

- Proof (upper bound):
  - Given preference profile $\vec{\succ}$, define harmonic scores $\text{sc}(a, \vec{\succ})$:
    - Each voter gives $1/k$ points to her k^{th} most preferred alternative
    - Take the sum of points across voters
  - How does the harmonic score relate to social welfare?
    - It is an upper bound on social welfare
      - $\text{sw}(a, \vec{u}) \le \text{sc}(a, \vec{\succ})$ (WHY?)
    - On average, it is a relatively tight upper bound
      - $\sum_a sc(a, \vec{\succ}) = n \cdot \sum_{k=1}^{m} 1/k = n\, H_m \le n \cdot (\ln m + 1)$
      - $\sum_a sw\left(a, \vec{\succ}\right) = n$

Randomized Rules

- Proof (upper bound):
  - Golden rule f:
    - With probability ½:
      - Choose every a with probability proportional to $sc(a, \vec{\succ})$
    - With the remaining probability ½:
      - Choose every a with probability $1/m$ (uniformly at random)
  - $\text{dist}(f) \le 2\sqrt{m \cdot (\ln m + 1)}$ (proof on the board!)

Some Thoughts

- How do we interpret the distortion number?
  - Sometimes distortion can be large
    - E.g. $\Theta(m^2)$ for deterministic rules
  - But if all alternatives have bad worst-case approximation ratio, the alternative that minimizes it is still, in a sense, better than the others
    - The best we can do given partial information

- Optimal vs asymptotically optimal
  - Plurality and "golden rule" are (almost) asymptotically optimal
  - But one can also write an optimization program that chooses the exact alternative minimizing distortion on each input $\vec{\succ}$
  - Polytime for both deterministic (via a direct formula) and randomized (via a non-trivial LP) cases

Some Thoughts

- Elicitation-distortion tradeoff
  - What about other types of partial information?
    - There is work on considering less information than rankings as well as more information than rankings
    - One can analyze a tradeoff between eliciting less information and achieving low distortion
- Extensions
  - Selecting a subset of k alternatives or a ranking of alternatives
  - Participatory budgeting
  - Graph matching

Deployed @ ROBOVOTE

Metric Distortion

- Instead of utilities, voters have costs for alternatives
- Underlying metric d
  - Voters and alternatives are in an underlying *metric space* with distance function d, which satisfies the triangle inequality
    - $\forall x, y, z: d(x, y) + d(y, z) \geq d(x, z)$
  - Social cost $sc(a, d) = \sum_i d(i, a)$
  - **Ideal goal:** Choose $a^* \in \mathrm{argmin}_a\, sc(a, d)$
- Preference profile $\vec{\succ} = (\succ_1, \ldots, \succ_n)$
  - Voter i ranks the alternatives according to their distance from her
    - $d(i, a) < d(i, b) \Rightarrow a \succ_i b$
    - As before, the voter can break ties arbitrarily

Metric Distortion

- Metric distortion of a voting rule f

$$\text{dist}(f) = \sup_d \frac{\text{sc}\left(f\left(\vec{\succ}\right), d\right)}{\min_a sc(a, d)}$$

  - Implicit max over all possible $\vec{\succ}$ that can be induced from d
  - If f is randomized, we need $E\left[sc\left(f\left(\vec{\succ}\right), d\right)\right]$

➢ Once again, we can consider both deterministic and randomized rules

Deterministic Rules

- A simple lower bound of 3 with just two candidates

$^{n}/_{2}$ voters $: a \succ b$

$^{n}/_{2}$ voters $: b \succ a$

Rule selects wlog → **Winner** $\boldsymbol{a}$ → Bad instance

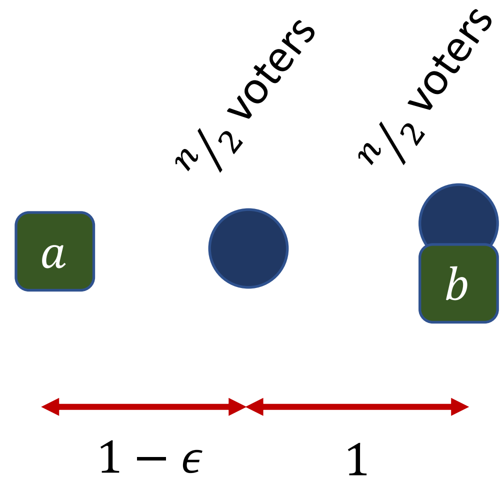

What about upper bounds?

Deterministic Rules

Distortion	Rule	Citation
Unbounded	k-approval ($k > 2$)	[Anshelevich et al., 2015]
$\Theta(m)$	Plurality, Borda count	[Anshelevich et al., 2015]
$\Theta(\sqrt{m})$	Ranked pairs, Schulze	[Kempe 2020]
$O(\log m)$, $\Omega\left(\sqrt{\log m}\right)$	STV	[Skowron and Elkind, 2017]
5	Copeland's rule	[Anshelevich et al., 2015]
$2 + \sqrt{5} \approx 4.236$	A new rule	[Munagala and Wang, 2019]
3	PluralityMatching	[Gkatzelis et al., 2020]

Randomized Rules

Distortion	Rule	Citation
$3 - {}^{2}/_{n}$	Random Dictatorship	[Anshelevich and Postl, 2017]
$3 - {}^{2}/_{m}$	Smart Dictatorship	[Kempe 2020, Gkatzelis et al. 2020]
≥ 2	Lower bound	Same example as before

- Major open question:
  - Does there exist a randomized voting rule with metric distortion 2?