Fair and Efficient Social Decision-Making

CSCI 699

Voting III: Distortion

Evi Micha

Credit for the slides: Nisarg Shah

Utilitarian Approach

Voting with Ranked Ballots

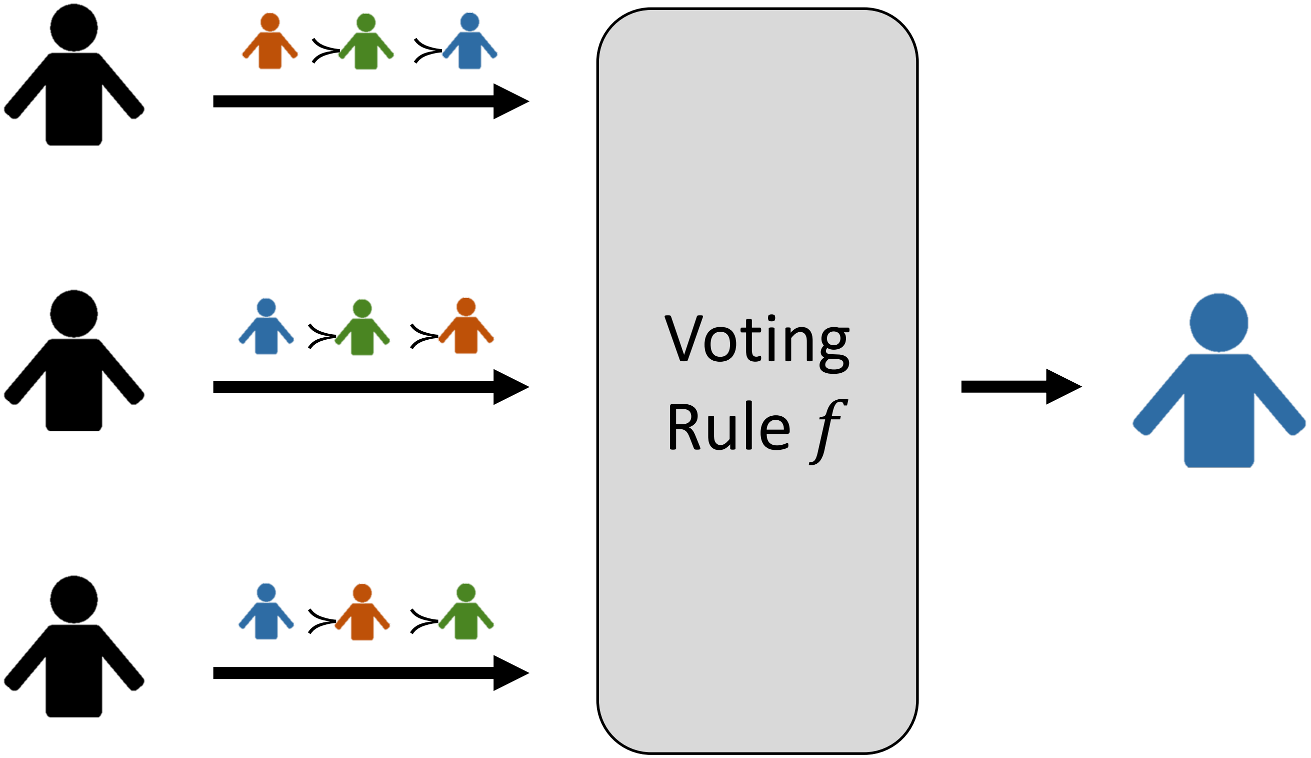

Utilitarian Voting with Ranked Ballots

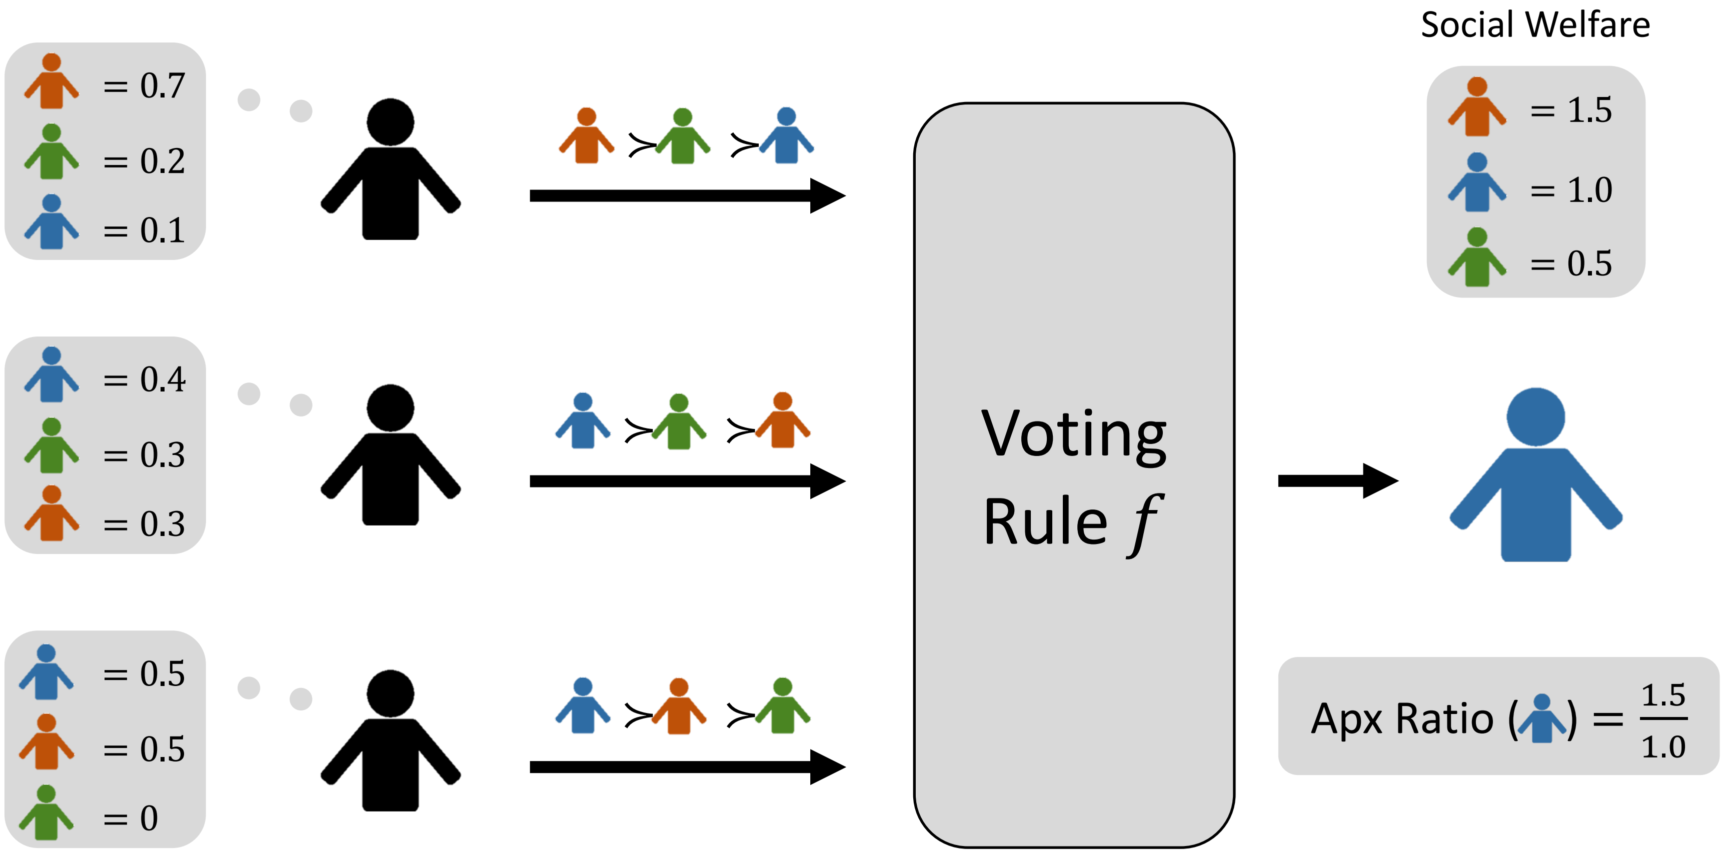

Optimal Voting Rules with Ranked Ballots

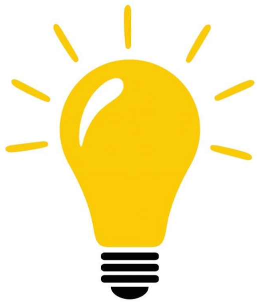

Minimize distortion
(Worst-case approximation ratio for utilitarian social welfare)

Voting with Ranked Ballots

- $N =$ set of n voters
- $A =$ set of m alternatives
  - $\Delta(A)$ = set of distributions over A
- $\overrightarrow{\succ} =$ observed ranked preference profile
  - $\succ_i$ = preference ranking of voter i
  - $a \succ_i b$ means the voter ranks a higher than b
- (Randomized) Voting rule f
  - Maps every preference profile $\overrightarrow{\succ}$ to a distribution over alternatives $f(\overrightarrow{\succ}) = x \in \Delta(A)$
  - We say that f is deterministic if $f(\overrightarrow{\succ})$ has singleton support for every $\overrightarrow{\succ}$

Utilitarian Distortion

1. There exists an underlying utility profile $\vec{u}$ such that for each $i \in N$:
   - Consistency (denoted $u_i \triangleright \succ_i$): $\forall a, b : a \succ_i b \Rightarrow u_i(a) \geq u_i(b)$
   - Unit-sum: $\sum_a u_i(a) = 1$
     - [Aziz 2019] provides seven justifications!
   - Risk-neutrality: For $x \in \Delta(A)$, $u_i(x) = \sum_a u_i(a) \cdot x(a)$

2. If we knew the utilities, we would want to maximize the (utilitarian) social welfare
   - $sw(x, \vec{u}) = \sum_{i \in N} u_i(x)$

3. Because this is impossible given the limited ranked information, we want to best approximate the social welfare in the worst case.

Utilitarian Distortion

- Distortion

$$\mathrm{dist}(x, \vec{\succ}) = \sup_{\vec{u} \rhd \vec{\succ}} \frac{\max_{a \in A} sw(a, \vec{u})}{sw(x, \vec{u})}$$

- Given voting rule f

$$dist(f) = \max_{\vec{\succ}} \mathrm{dist}\left(f(\vec{\succ}), \vec{\succ}\right)$$

What is the lowest possible $dist(f)$? Which voting rule achieves it?

Example

Voters

Alternatives

$1 : a \succ b \succ c$

$\quad {}^{1}/_{2} \quad {}^{1}/_{2} \quad 0$

$2 : b \succ a \succ c$

$\quad {}^{2}/_{3} \quad {}^{1}/_{4} \quad {}^{1}/_{12}$

$3 : a \succ c \succ b$

$\quad {}^{1}/_{3} \quad {}^{1}/_{3} \quad {}^{1}/_{3}$

- Suppose we choose a:
  - How much better is b?

$$\frac{sw(b,\vec{u})}{sw(a,\vec{u})} = \frac{{}^{1}/_{2} + {}^{2}/_{3} + {}^{1}/_{3}}{{}^{1}/_{2} + {}^{1}/_{4} + {}^{1}/_{3}} = \frac{18}{13}$$

  - How much better is c?

$$\frac{sw(c,\vec{u})}{sw(a,\vec{u})} = \frac{0 + {}^{1}/_{12} + {}^{1}/_{3}}{{}^{1}/_{2} + {}^{1}/_{4} + {}^{1}/_{3}} = \frac{5}{13}$$

Example

Voters

Alternatives

$1 : a \succ b \succ c$

$^1/_3 \quad ^1/_3 \quad ^1/_3$

$2 : b \succ a \succ c$

$1 \quad 0 \quad 0$

$3 : a \succ c \succ b$

$^1/_3 \quad ^1/_3 \quad ^1/_3$

- Suppose we choose a:
  - How much better can b be?

$$\max_{\vec{u} \rhd \vec{\succ}} \frac{sw(b, \vec{u})}{sw(a, \vec{u})} = \frac{^1/_3 + 1 + {^1/_3}}{^1/_3 + 0 + {^1/_3}} = \frac{5}{2}$$

  - How much better can c be?

$$\max_{\vec{u} \rhd \vec{\succ}} \frac{sw(c, \vec{u})}{sw(a, \vec{u})} = \frac{^1/_3 + 0 + {^1/_3}}{^1/_3 + 0 + {^1/_3}} = 1$$

  - Hence, $dist(a, \vec{\succ}) = \frac{5}{2}$

Example

Voters

Alternatives

$1 : a \succ b \succ c$

$1 \quad 0 \quad 0$

$2 : b \succ a \succ c$

${}^{1}/_{3} \quad {}^{1}/_{3} \quad {}^{1}/_{3}$

$3 : a \succ c \succ b$

$1 \quad 0 \quad 0$

- Suppose we choose b:
  - How much better can a be?

$$\max_{\vec{u} \triangleright \vec{\succ}} \frac{sw(a, \vec{u})}{sw(b, \vec{u})} = \frac{1 + {}^{1}/_{3} + 1}{0 + {}^{1}/_{3} + 0} = 7$$

Example

Voters

Alternatives

$1 : a \succ b \succ c$

$1 \quad 0 \quad 0$

$2 : b \succ a \succ c$

$1/2 \quad 1/2 \quad 0$

$3 : a \succ c \succ b$

$1 \quad 0 \quad 0$

- Suppose we choose c:
  - How much better can a be?

$$\max_{\vec{u} \triangleright \vec{\succ}} \frac{sw(a, \vec{u})}{sw(c, \vec{u})} = \frac{1 + 1/2 + 1}{0 + 0 + 0} = inf$$

Optimal Deterministic Distortion

- Theorem [Caragiannis, Procaccia, 2011; Caragiannis, Nath, Procaccia, Shah, 2017]
  - For deterministic aggregation of ranked ballots, the optimal distortion is $\Theta(m^2)$
- 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 bad utility profile $\vec{u}$ consistent with $\vec{\succ}$ in which a is $\Omega(m^2)$ factor worse than the optimal alternative

Deterministic Rules

- Proof (lower bound):
  - Let f be any deterministic voting rule
  - Consider $\vec{\succ}$ on the right
  - 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
    - $\text{sw}(a_i, \vec{u}) = \frac{n}{m-1} \cdot \frac{1}{m}$, $\text{sw}(a_m, \vec{u}) \geq \frac{n - n/(m-1)}{2} = \Omega(n)$
    - Distortion = $\Omega(m^2)$

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

Deterministic Rules

- Proof (upper bound):
  - Plurality rule: Select an alternative a that is the top choice of the most voters
  - For this plurality winner:
    - At least n/m voters have a as their top choice (pigeonhole principle)
    - Every voter has utility at least $1/m$ for their top choice (pigeonhole principle)
  - Hence, for every consistent utility profile $\vec{u}$:
    - $sw(a, \vec{u}) \geq n/m^2$
    - $sw(a^*, \vec{u}) \leq n$ for every alternative a^*
  - $dist(a, \vec{\succ}) = O(m^2)$

Optimal Randomized Distortion

- Theorem [Boutilier, Caragiannis, Haber, Lu, Procaccia, and Sheffet, 2015]
  - For randomized aggregation of ranked ballots, the optimal distortion is $O(\sqrt{m} \cdot \log^* m)$ but $\Omega(\sqrt{m})$
- 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 x over alternatives
    - Reveal a bad utility profile $\vec{u}$ consistent with $\vec{\succ}$ in which the expected social welfare under x is $\Omega(\sqrt{m})$ factor worse than the optimal social welfare

Randomized Rules

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

- Proof (lower bound):
  - Let f be an arbitrary rule
  - Consider $\vec{\succ}$ on the right with $\sqrt{m}$ special alternatives
  - f returns distribution x in which at least one special alternative (say a_i) must be chosen w.p. at most ${}^{1}/_{\sqrt{m}}$
  - Bad utility profile $\vec{u}$ consistent with $\vec{\succ}$:
    - All voters ranking a_i first have utility 1 for a_i
    - All other voters have utility ${}^{1}/_{m}$ for every alternative
    - $\text{sw}(a_i, \vec{u}) = \Theta\left({}^{n}/_{\sqrt{m}}\right)$ but $sw(a, \vec{u}) \le {}^{n}/_{m}$ for every other alternative a
    - $sw(x, \vec{u}) \le \left({}^{1}/_{\sqrt{m}}\right) \cdot \Theta\left({}^{n}/_{\sqrt{m}}\right) + \left(1 - {}^{1}/_{\sqrt{m}}\right) \cdot \left({}^{n}/_{m}\right) = O\left({}^{n}/_{m}\right)$
    - Hence, $dist(x, \vec{u}) = \Omega(\sqrt{m})$

Optimal Randomized Distortion

- Harmonic Rule
  - The rule that achieves $O(\sqrt{m} \cdot \log^* m)$ distortion is complicated, but they propose a simpler harmonic rule that achieves $O\left(\sqrt{m \cdot \log m}\right)$ distortion

Harmonic Rule

- Each voter i awards ${}^{1}/_{r}$ points to her r^{th} ranked alternative for every $r \in \{1, \dots m\}$
- Harmonic score of alternative a, denoted $hsc(a, \vec{\succ})$, is the total point awarded to a
- W.p. ½ , choose each $a \in A$ with probability proportional to $hsc(a, \vec{\succ})$
- W.p. ½, choose each $a \in A$ uniformly at random

- Key proof idea:
  - $hsc(a, \vec{\succ}) \geq sw(a, \vec{u})$ for every a, while $\sum_a hsc(a, \vec{\succ}) = O(\log m) \cdot \sum_a sw(a, \vec{u})$

Optimal Randomized Distortion

- Theorem [Ebadian, Kahng, Peters, Shah, 2022]
  - For randomized aggregation of ranked ballots, the optimal distortion is $\Theta(\sqrt{m})$.

Metric Distortion

[Anshelevich, Bhardwaj, Elkind, Postl, Skowron, 2018]

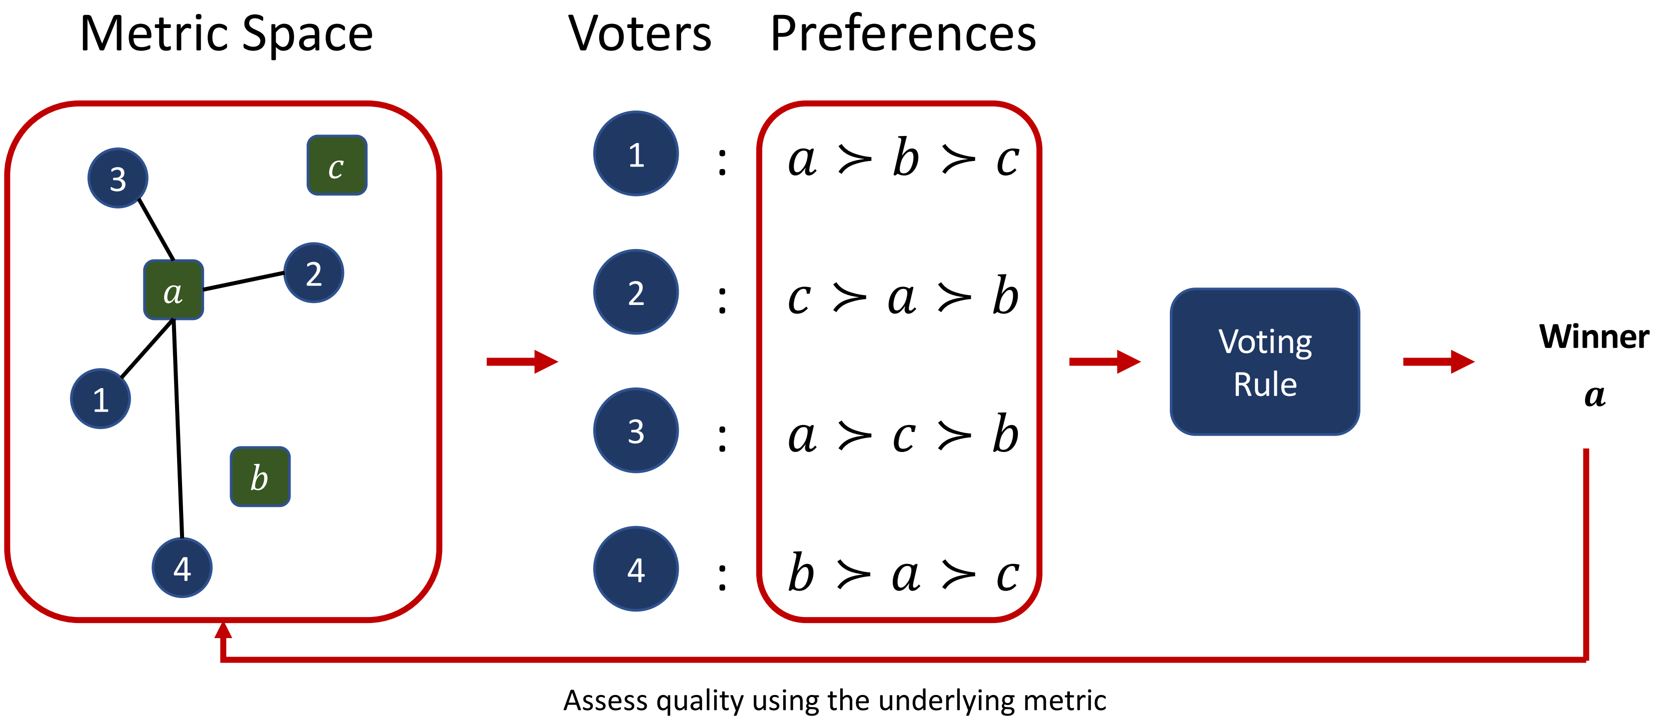

Why The Metric?

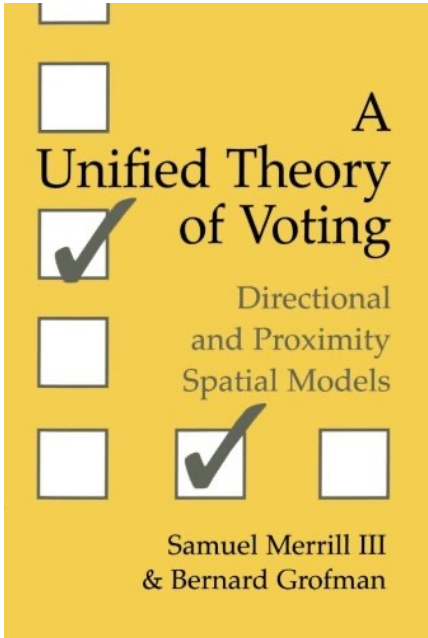

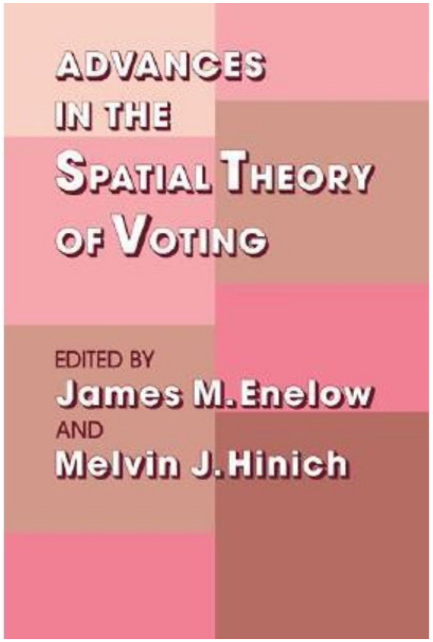

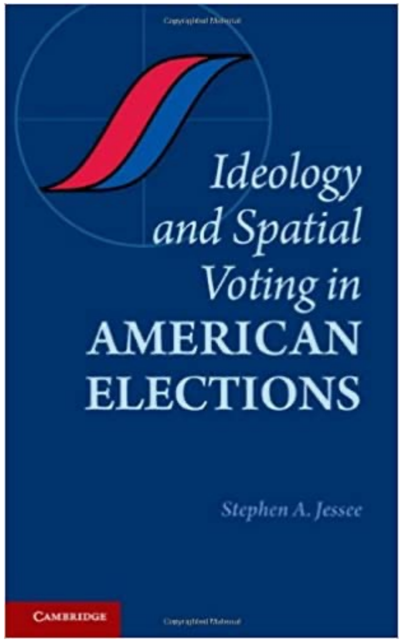

Why The Metric?

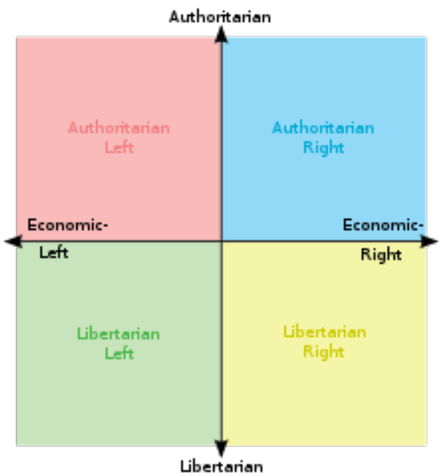

2D Models

3D Models

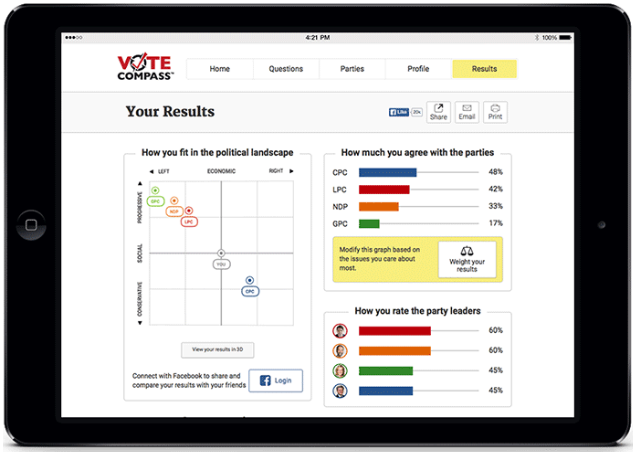

Popular Tools

Metric Distortion

1. There exists an underlying metric d over voters and alternatives such that:
   - Consistency (denoted $d \rhd \overrightarrow{\succ}$) : $\forall a, b : a \succ_i b \Rightarrow d(i,a) \leq d(i,b)$
   - Triangle inequality: $\forall x, y, z, d(x,y) + d(y,z) \geq d(x,z)$
   - Risk-neutrality: For $x \in \Delta(A)$, $c_i(x) = \sum_a d(i,a) \cdot x(a)$

2. If we knew the costs, we would minimize the social cost
   - $sc(x,d) = \sum_{i \in N} d(i,x)$

3. Because this is impossible given the limited ranked information, we want to best approximate the social cost in the worst case.

Metric Distortion

- Distortion

$$\text{dist}(x, \vec{\succ}) = \sup_{d \,\rhd\, \vec{\succ}} \frac{sc(x,d)}{\min_{a \in A} sc(a,d)}$$

- Given voting rule f

$$dist(f) = \max_{\vec{\succ}} \text{dist}\left(f(\vec{\succ}), \vec{\succ}\right)$$

What is the lowest possible distortion of deterministic and randomized rules? Which voting rules achieves it?

Lower Bound

[Anshelevich, Bhardwaj, Elkind, Postl, Skowron, 2018]

- A simple lower bound of 3 (deterministic rules) with just two candidates

$n/2$ voters $: a \succ b$

$n/2$ voters $: b \succ a$

Rule selects → (w.l.o.g.) **Winner** $\boldsymbol{a}$ → Bad instance

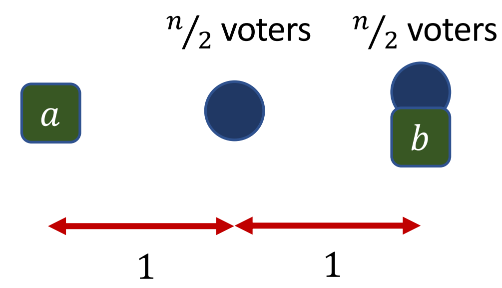

$$sc(a,d) = \frac{3n}{2},\ sc(b,d) = \frac{n}{2} \Rightarrow \text{distortion} \geq \frac{sc(a,d)}{sc(b,d)} \geq 3$$

Can a deterministic rule achieve distortion 3?

Deterministic Rules

- Question: What is the distortion of veto?
- Unbounded!

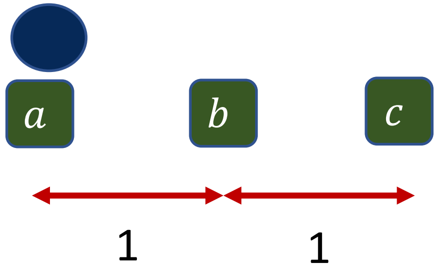

Deterministic Rules

- Theorem [Anshelevich, Bhardwaj, Elkind, Postl, Skowron, 2018]:

Rule	Distortion
k-approval ($k > 2$)	Unbounded
Plurality, Borda count	$\Theta(m)$
Harmonic rule*	$O\left(\frac{m}{\sqrt{\log m}}\right), \Omega\left(\frac{m}{\log m}\right)$
Best positional scoring rule	$\Omega(\sqrt{\log m})$
STV	$O(\log m),\ \Omega(\sqrt{\log m})$
Copeland's rule	5
Best deterministic rule	≥ 3

*Deterministic version of the harmonic rule,
which simply picks an alternative with the largest harmonic score

- The instance-optimal deterministic rule can be computed in polynomial time by solving a number of linear programs.

- Open question: What is the best distortion achievable by any positional scoring rule?

Deterministic Rules

- Theorem [Munagala, Wang, 2019]:
  - There exists a deterministic voting rule with distortion $2 + \sqrt{5} \approx 4.236$.

- Lemma [Munagala, Wang, 2019]: If f is a voting rule such that for every election, the domination graph of $f(\overrightarrow{\succ})$ has a perfect matching, then f has distortion equal to 3.

Domination Graph of Candidate a

Edge (i, j) exists when, in i's vote, a weakly defeats the top choice of j

$a \succ b \succ c$

$c \succ a \succ b$

$a \succ c \succ b$

$b \succ a \succ c$

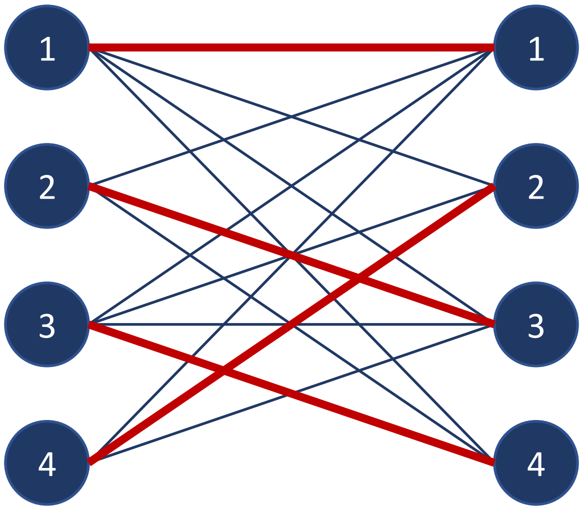

$a \succ b \succ c$

$c \succ a \succ b$

$a \succ c \succ b$

$b \succ a \succ c$

Perfect Matching

Main Lemma

- Lemma [Munagala, Wang, 2019]: If f is a voting rule such that for every election, the domination graph of $f(\overrightarrow{\succ})$ has a perfect matching, then f has distortion equal to 3.
- Proof
- Let a be the optimal alternative

$sc(a) = \sum_{i \in N} d(i, a)$

$\leq \sum_{i \in N} d(i, top(M(i))$ ($a \succcurlyeq_i$ top(M(i)) from the definition of the domination graph)

$\leq \sum_{i \in N} d(i, b) + d(b, top(M(i))$ (triangle inequality)

$\leq \sum_{i \in N} d(i, b) + d(b, top(i)$ (M is a perfect mathing)

$\leq \sum_{i \in N} d(i, b) + d(b, i) + d(i, top(i))$ (triangle inequality)

$\leq \sum_{i \in N} d(i, b) + d(b, i) + d(i, b)$

$\leq 3 \cdot sc(b)$

Optimal Distortion

- Theorem [Gkatzelis, Halpern, Shah, 2020]:
  - There always exists an alternative whose domination graph admits a perfect matching, and PluralityMatching outputs any such alternative.

- Theorem [Kizilkaya, Kempe, 2022]:
  - There always exists an alternative whose domination graph admits a perfect matching, and Plurality Veto outputs any such alternative.

Randomized Rules

- Theorem [Anshelevich, Bhardwaj, Elkind, Postl, Skowron, 2018]:
  - No randomized rule has distortion better than 2.
  - RandomDictatorship has distortion $3 - {}^{2}/_{n}$.

- Theorem [Kempe 2020a]:
  - There is a randomized voting rule with access to only plurality votes with distortion $3 - {}^{2}/_{m}$.

- Theorem [Charikar, Ramakrishnan, 2022; Pulyassary, Swamy, 2021]:
  - No randomized rule has distortion better than 2.112 for all m.
    - Weaker lower bounds for fixed, finite m

- Theorem [Charikar, Ramakrishnan, Wang, Wu, 2024]:
  - There is a randomized voting rules with distortion less than 2.753.

- Open question: What is the optimal metric distortion of randomized rules?