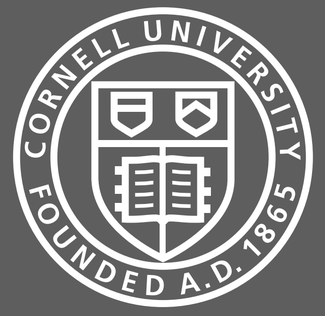
Implicit Utilitarian Voting

Paul Gölz

Lecture 16, ORIE 4360/5360, Fall ’26



Recap: Cardinal Utilities

Each voter i has a utility $u_i(x) \in \mathbb{R}$ for each alternative x, which measures the amount of happiness/satisfaction they get if x wins.

A utility profile $\boldsymbol{u} = (u_1, \ldots, u_n) \in (M \to \mathbb{R})^n$ is a vector with a utility function for each voter.

The utilitarian welfare of alternative x is $\mathrm{SW}(x, \boldsymbol{u}) := \sum_{i \in N} u_i(x)$.

Recap: Cardinal Utilities

Last time, we saw that cardinal utilities are much more expressive than rankings/approvals, for example because they can express intensity.

But we also discussed reasons not to directly ask voters for cardinal utilities: difficulty of elicitation and big opportunities for manipulation.

Now: Analyze ranked-choice elections through the possible cardinal utilities behind the rankings.

Back to our Basic Voting Model

Voters $N = \{1,2,\ldots,n\}$.

Alternatives $M = \{a, b, c, \ldots\}, |M| = m$.

Each voter i has a ranking $\sigma_i \in \mathscr{L}$ over the alternatives; $a \succ_{\sigma_i} b$ means that i prefers a to b.

A preference profile $\boldsymbol{\sigma} \in \mathscr{L}^n$ is a collection of all voters' rankings.

A social choice function is a function $f : \mathscr{L}^n \to M$.

Consistent Rankings

Let $u_i : M \to \mathbb{R}$ be a utility function, and $\sigma_i \in \mathscr{L}$ a ranking of alternatives. σ_i is consistent with u_i ("$u_i \rhd \sigma_i$") if, for all $x, y \in M$ with $x \succ_{\sigma_i} y$, it holds that $u_i(x) \geq u_i(y)$.

Let $u_i(a) = 0.2$, $u_i(b) = 0.6$, $u_i(c) = 0.2$.
Which rankings are consistent with u_i?

pollev.com/goelz

Distortion of a Preference Profile

For a given preference profile $\boldsymbol{\sigma} \in \mathscr{L}^n$ and social choice function f, define f's distortion on $\boldsymbol{\sigma}$ as:

$$\mathrm{dist}(f, \boldsymbol{\sigma}) := \sup_{\boldsymbol{u} \rhd \boldsymbol{\sigma}} \frac{\max_{x \in M} \mathrm{SW}(x, \boldsymbol{u})}{\mathrm{SW}(f(\boldsymbol{\sigma}), \boldsymbol{u})},$$

where "$\boldsymbol{u} \rhd \boldsymbol{\sigma}$" ranges over all utility profiles $\boldsymbol{u} = (u_1, \ldots, u_n) \in (M \to \mathbb{R}_{\geq 0})^n$ such that $u_i \rhd \sigma_i$ and $\sum_{x \in M} u_i(x) = 1$ for all $i \in N$.

Implicit Utilitarian Voting

Conceptually: want to maximize utilitarian welfare, but only know the rankings not the utilities ("implicit utilitarian voting").

$$\mathrm{dist}(f, \boldsymbol{\sigma}) := \sup_{\boldsymbol{u} \rhd \boldsymbol{\sigma}} \frac{\max_{x \in M} \mathrm{SW}(x, \boldsymbol{u})}{\mathrm{SW}(f(\boldsymbol{\sigma}), \boldsymbol{u})}.$$

Supremum means we take worst-case utilities, which make ratio as high as possible. But definition restricted to nonnegative utilities, that add up to 1 for each voter.

Measures how much better we could have done than $f(\boldsymbol{\sigma})$. Ratio at least 1, could be ∞ if denominator is zero.

Compute Borda Distortion

1	2	3	4
a	c	d	e
b	b	b	b
c	d	e	a
d	e	a	c
e	a	c	d

$$\mathrm{dist}(f, \boldsymbol{\sigma}) := \sup_{\boldsymbol{u} \rhd \boldsymbol{\sigma}} \frac{\max_{x \in M} \mathrm{SW}(x, \boldsymbol{u})}{\mathrm{SW}(f(\boldsymbol{\sigma}), \boldsymbol{u})},$$

where "$\boldsymbol{u} \rhd \boldsymbol{\sigma}$" ranges over all utility profiles $\boldsymbol{u} = (u_1, \ldots, u_n) \in (M \to \mathbb{R}_{\geq 0})^n$ such that $u_i \rhd \sigma_i$ and $\sum_{x \in M} u_i(x) = 1$ for all $i \in N$.

What is the distortion of Borda count on the profile above?

pollev.com/goelz

Distortion of a Voting Rule

We can also define the distortion of a voting rule without referring to any fixed preference profile:

$$\mathrm{dist}(f) := \sup_{\boldsymbol{\sigma} \in \mathscr{L}^n} \mathrm{dist}(f, \boldsymbol{\sigma}).$$

I.e., consider the worst-case preference profile.

For example, last poll question shows that the distortion of Borda can be infinite.

Distortion of Plurality

Theorem [Caragiannis & Procaccia 2011]**.**
$\text{dist}(f_{\text{plurality}}) \leq m^2 - 1.$

Proof. If x is ranked first in σ_i, then $u_i(x) \geq 1/m$ for all $u_i \rhd \sigma_i$ (assuming utilities nonneg., sum 1).

The plurality winner must be ranked first by at least n/m voters. So, in any consistent $\boldsymbol{u}$,
$\text{SW}(f_{\text{plurality}}, \boldsymbol{u}) \geq n/m^2$.

$\sum_{y \in M} \text{SW}(y, \boldsymbol{u}) = n$, so no other alternative can have more than $m^2 - 1$ times higher welfare.



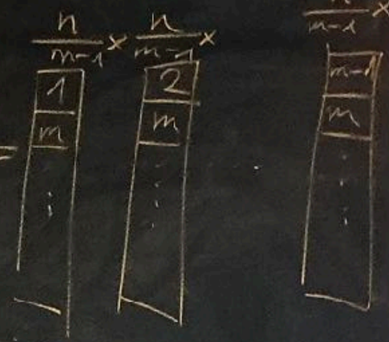
Lower Bound on Distortion

Theorem [Caragiannis & Procaccia 2011]**.** Each (deterministic) social choice function has distortion $\Omega(m^2)$ (i.e., at least on the order m^2).

[Proof on Blackboard.]

Proof

Proof: For exposition, assume that $m-1 \mid n$. Fix any social choice function f. Claim: $\text{dist}(f, \sigma) = \Omega(m^2)$.



Case $f(\sigma) = m$. If every voter has utility one for top alternative and zero otherwise, distortion is ∞. ✓

Case: $1 \le f(\sigma) \le m-1$. Say that the bloc of voters ranking $f(\sigma)$ first has utility $1/m$ for all alternatives. And say that all other blocs have utility $1/2$ for their top two and zero otherwise.

$$SW(f(\sigma), u) = \frac{n}{m-1} \cdot \frac{1}{m}.$$

$$SW(m, u) = \frac{n}{m-1} \cdot \frac{1}{m} + \left(n - \frac{n}{m-1}\right) \frac{1}{2} = \frac{n}{2} - \frac{n}{2(m-1)} + \frac{n}{m(m-1)}$$

$$\text{dist}(f) \ge \text{dist}(f, \sigma) \ge \frac{\frac{n}{2} - \frac{n}{2(m-1)} + \frac{n}{m(m-1)}}{\frac{n}{m(m-1)}} = m(m-1)\left(\frac{1}{2} - \frac{1}{2(m-1)} + \frac{1}{m(m-1)}\right)$$

$$= \Omega(m^2).$$

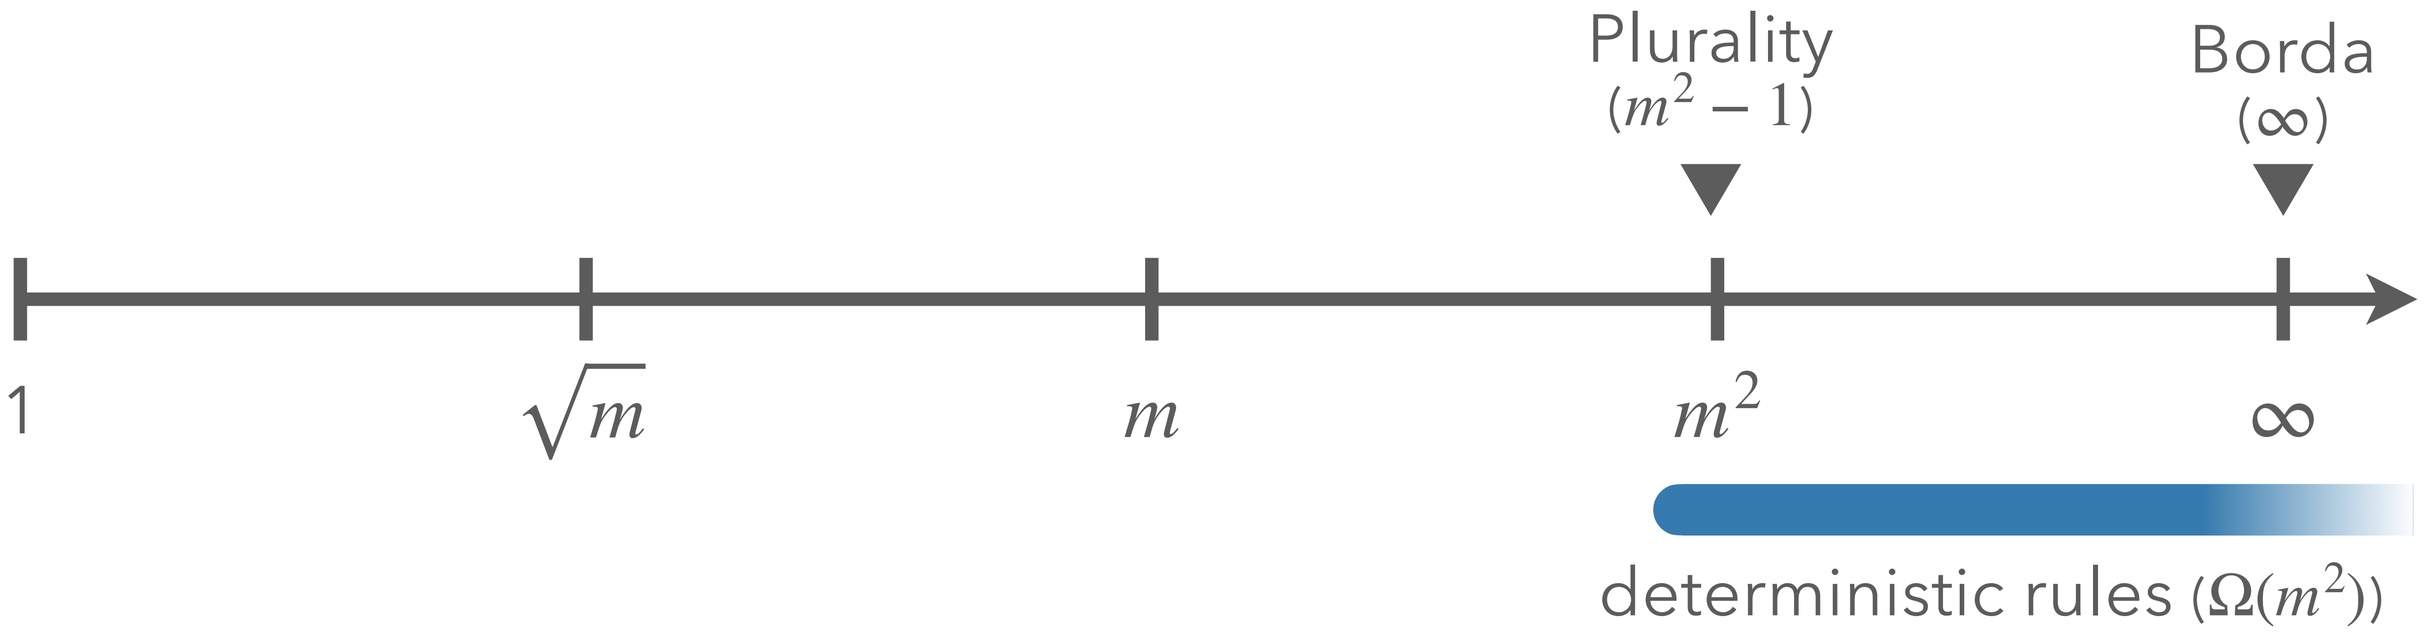
The Revenge of Plurality

Theorem [Caragiannis & Procaccia 2011]**.** $\text{dist}(f_{\text{plurality}}) \leq m^2 - 1.$

Theorem [Caragiannis & Procaccia 2011]**.** Each (deterministic) social choice function has distortion $\Omega(m^2)$ (i.e., at least on the order m^2).

For once, plurality comes out on top!

Distortion Landscape



Instance-Optimality

Plurality has optimal worst-case bound (up to constant factors). Can we find a voting rule that has the best distortion on every $\boldsymbol{\sigma}$?

$f^* = \text{argmin}_{x \in M} \sup_{\boldsymbol{u} \rhd \boldsymbol{\sigma}} \frac{\max_{x' \in M} \text{SW}(x', \boldsymbol{u})}{\text{SW}(x, \boldsymbol{u})}$ has the best distortion, by definition.

Computing Instance-Optimal Rule

Steps:

- For each $a \neq b$, calculate $d_{a,b} := \sup_{\boldsymbol{u} \triangleright \boldsymbol{\sigma}} \frac{\mathrm{SW}(b, \boldsymbol{u})}{\mathrm{SW}(a, \boldsymbol{u})}$.
- Return $\operatorname{argmin}_{a \in M} \max_{b \in M} d_{a,b}$.

Question

1	2	3	4	5
b	*c*	*d*	*e*	*a*
a	*a*	*b*	*d*	*d*
c	*e*	*e*	*a*	*b*
d	*b*	*c*	*b*	*e*
e	*d*	*a*	*c*	*c*

What is $d_{b,a} := \sup_{\boldsymbol{u} \rhd \boldsymbol{\sigma}} \frac{\mathrm{SW}(a, \boldsymbol{u})}{\mathrm{SW}(b, \boldsymbol{u})}$?

Computing Instance-Optimal Rule

How to calculate $d_{a,b} := \sup_{\boldsymbol{u} \rhd \boldsymbol{\sigma}} \frac{\mathrm{SW}(b, \boldsymbol{u})}{\mathrm{SW}(a, \boldsymbol{u})}$?

Can compute $\boldsymbol{u}$ that makes b as good as possible while making a as bad as possible. For each i:

- If $b \succ_{\sigma_i} a$, equally divide utility between b and all better-ranked alternatives, zero for a.
- If $a \succ_{\sigma_i} b$ and a not ranked first, give all utility to first-ranked.
- Else, give all alternatives utility $1/m$.

Bibliography

Procaccia & Rosenschein: *The distortion of cardinal preferences in voting*. Cooperative Information Agents (2006).

Boutilier et al.: *Optimal social choice functions*. EC (2012).

Ebadian et al.: *Optimized Distortion and Proportional Fairness in Voting*. TEAC (2024).