Frugal Preference Aggregation: The Ex-Ante Condorcet Approach

Klaus Nehring and Clemens Puppe

COMSOC Summer School 2023

July 18, 2023

Agenda

Motivation

- Majoritarian choice under 'ambiguous' beliefs about voters' prefs.
- **Frugal aggregation model:**
  - Information about individual top choices.
  - Ambiguous beliefs about underlying complete preferences.
- **Ex-ante Condorcet approach:**
  - Every pair of alternatives induces possible expected majorities...
  - ... and thus yields ex-ante net majority tournament.
  - Complete ignorance about ex-post preferences.
  - Maximal elements: 'Ex-ante Condorcet winners.'
- Applications: budget allocation ('participatory budgeting'), spatial voting, collective choice in space of characteristics.
- Different assumptions on the epistemic state of social evaluator:
  - Plain convexity
  - Symmetry

Motivation

- Majoritarian choice under ‘ambiguous’ beliefs about voters’ prefs.
- **Frugal aggregation model:**
  - Information about individual top choices.
  - Ambiguous beliefs about underlying complete preferences.
- **Ex-ante Condorcet approach:**
  - Every pair of alternatives induces possible expected majorities...
  - ... and thus yields ex-ante net majority tournament.
  - Complete ignorance about ex-post preferences.
  - Maximal elements: ‘Ex-ante Condorcet winners.’
- Applications: budget allocation (‘participatory budgeting’), spatial voting, collective choice in space of characteristics.
- Different assumptions on the epistemic state of social evaluator:
  - Plain convexity → generic plurality winner.
  - Symmetry → (strict) Tukey median.

Relation to Literature

- Frugal aggregation as opposed to
  - (i) standard *Arrovian preference aggregation* on economic domains (Le Breton & Weymark, 2004), and
  - (ii) standard *spatial voting* (Austen-Smith & Banks, 1999) (which is a degenerate special case).
- (Non-Bayesian) preference aggregation under incomplete information (Boutilier & Rosenschein, 2016; Lang, 2020).
- Decision making under complete ignorance (Luce & Raiffa, 1957; Nehring, 2000; 2009).
- Participatory budgeting (Aziz & Shah, 2020, Goel et al., 2019; Freeman et al., 2021).

Background: Pref. Aggregation on Economic Domains

- Consider a *common* restricted domain $\mathcal{D}_i = \mathcal{D} \subseteq \mathcal{R}$ on some set of alternatives X, where $\mathcal{R}$ is the set of all weak orders.
- $\mathcal{D}$ has the **free triple** property if the restriction of $\mathcal{D}$ to any triple of distinct alternatives is unrestricted.
- **Proposition** If $\mathcal{D}$ has the free triple property, then every social welfare function $F : \mathcal{D}^n \longrightarrow \mathcal{R}$ satisfying IIA and WP is dictatorial.

Background: Pref. Aggregation on Economic Domains

- Consider a *common* restricted domain $\mathcal{D}_i = \mathcal{D} \subseteq \mathcal{R}$ on some set of alternatives X, where $\mathcal{R}$ is the set of all weak orders.
- $\mathcal{D}$ has the **free triple** property if the restriction of $\mathcal{D}$ to any triple of distinct alternatives is unrestricted.
- **Proposition** If $\mathcal{D}$ has the free triple property, then every social welfare function $F : \mathcal{D}^n \longrightarrow \mathcal{R}$ satisfying IIA and WP is dictatorial.
- **Example:** Let $X \subseteq \mathbb{R}^L$ be convex, and consider the space of all convex preferences $\mathcal{D}_{\mathrm{co}}$ on X. This domain does **not** have the free triple property.

Background: Pref. Aggregation on Economic Domains

- Consider a *common* restricted domain $\mathcal{D}_i = \mathcal{D} \subseteq \mathcal{R}$ on some set of alternatives X, where $\mathcal{R}$ is the set of all weak orders.
- $\mathcal{D}$ has the **free triple** property if the restriction of $\mathcal{D}$ to any triple of distinct alternatives is unrestricted.
- **Proposition** If $\mathcal{D}$ has the free triple property, then every social welfare function $F : \mathcal{D}^n \longrightarrow \mathcal{R}$ satisfying IIA and WP is dictatorial.
- **Example:** Let $X \subseteq \mathbb{R}^L$ be convex, and consider the space of all convex preferences $\mathcal{D}_{\mathrm{co}}$ on X. This domain does **not** have the free triple property.
- **Proposition** Suppose $X \subseteq \mathbb{R}^L$, and for all distinct $a, b \in X$, there is $c \in X$ such that $\{a, b, c\}$ is not collinear (note: $\Rightarrow L > 1$); then, every social welfare function $F : \mathcal{D}_{\mathrm{co}}^n \longrightarrow \mathcal{R}$ satisfying IIA and WP is dictatorial (follows from Kalai, Muller & Satterthwaite, 1979).

Background: (No) Condorcet Winners in Spatial Voting

- Standard spatial voting assumes that $X \subseteq \mathbb{R}^L$ and that voters have **Euclidean preferences**, i.e. a voter with top θ_i has utility function

$$u_{\theta_i}(x) \;=\; -(||x-\theta_i||_2)^2 \;=\; -(x-\theta_i)^T \cdot (x-\theta_i).$$

Geometrically, voters have *circles* as indifference curves.

Background: (No) Condorcet Winners in Spatial Voting

- Standard spatial voting assumes that $X \subseteq \mathbb{R}^L$ and that voters have **Euclidean preferences**, i.e. a voter with top θ_i has utility function

$$u_{\theta_i}(x) \;=\; -(||x-\theta_i||_2)^2 \;=\; -(x-\theta_i)^T \cdot (x-\theta_i).$$

Geometrically, voters have *circles* as indifference curves.

- Condorcet winners?

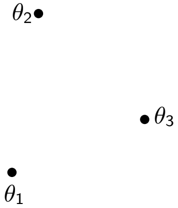

Background: (No) Condorcet Winners in Spatial Voting

- Standard spatial voting assumes that $X \subseteq \mathbb{R}^L$ and that voters have **Euclidean preferences**, i.e. a voter with top θ_i has utility function

$$u_{\theta_i}(x) \;=\; -(||x - \theta_i||_2)^2 \;=\; -(x - \theta_i)^T \cdot (x - \theta_i).$$

Geometrically, voters have *circles* as indifference curves.

- Condorcet winners?

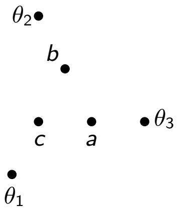

Background: (No) Condorcet Winners in Spatial Voting

- Standard spatial voting assumes that $X \subseteq \mathbb{R}^L$ and that voters have **Euclidean preferences**, i.e. a voter with top θ_i has utility function

$$u_{\theta_i}(x) = -(||x - \theta_i||_2)^2 = -(x - \theta_i)^T \cdot (x - \theta_i).$$

Geometrically, voters have *circles* as indifference curves.

- Condorcet winners?

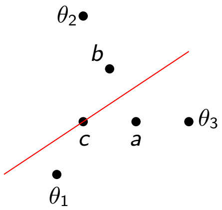

Background: (No) Condorcet Winners in Spatial Voting

- Standard spatial voting assumes that $X \subseteq \mathbb{R}^L$ and that voters have **Euclidean preferences**, i.e. a voter with top θ_i has utility function

$$u_{\theta_i}(x) \;=\; -(||x-\theta_i||_2)^2 \;=\; -(x-\theta_i)^T \cdot (x-\theta_i).$$

Geometrically, voters have *circles* as indifference curves.

- Condorcet winners?

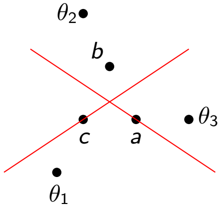

Background: (No) Condorcet Winners in Spatial Voting

- Standard spatial voting assumes that $X \subseteq \mathbb{R}^L$ and that voters have **Euclidean preferences**, i.e. a voter with top θ_i has utility function

$$u_{\theta_i}(x) \ = \ -(||x - \theta_i||_2)^2 \ = \ -(x - \theta_i)^T \cdot (x - \theta_i).$$

  Geometrically, voters have *circles* as indifference curves.

- Condorcet winners?

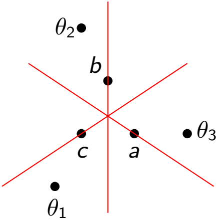

Background: (No) Condorcet Winners in Spatial Voting

- Standard spatial voting assumes that $X \subseteq \mathbb{R}^L$ and that voters have **Euclidean preferences**, i.e. a voter with top θ_i has utility function

$$u_{\theta_i}(x) \; = \; -(||x - \theta_i||_2)^2 \; = \; -(x - \theta_i)^T \cdot (x - \theta_i).$$

  Geometrically, voters have *circles* as indifference curves.

- Condorcet winners?

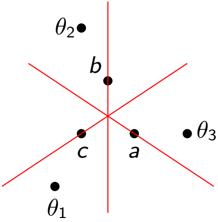

Thus, $a \succ_{\mathrm{Maj}} b$, $b \succ_{\mathrm{Maj}} c$, $c \succ_{\mathrm{Maj}} a$!

Background: (No) Condorcet Winners in Spatial Voting

- Hence, in the above situation, **every** element of $\{a, b, c\}$ beats **any other** element of $\{a, b, c\}$ via a *majority path*.
- **Proposition (McKelvey, 1976)** Generically, every $x \in X$ beats every $y \in X$ along some majority path in the standard spatial voting model.
- This negative result has been significantly generalized to the case of general continuous preferences and voting procedures beyond pairwise majority comparisons by McKelvey (1979).
- **Upshot:** In spatial voting, Condorcet winners fail to exist almost always. And even worse, generically an agenda setter can induce **every** alternative as the winner of a sequential majority vote by choosing an appropriate sequence of intermediate comparisons ('McKelvey-Schofield Chaos Theorem').

Agenda

The Ex-Ante Condorcet Approach

- $X = \{x, y, ...\}$ abstract space of alternatives.
- $\theta = (\theta_1, ..., \theta_n) \in X^n$ **known** profile of (unique) top choices.
- π prior belief over underlying profile $(\succcurlyeq_1, ..., \succcurlyeq_n)$ of complete ('ex-post') preferences on X.
- Π convex set of priors: social evaluator's **ambiguous beliefs.**
  - E.g. $X \subseteq \mathbb{R}^L$, Π_{co} all priors over profiles of convex preferences.
- (θ, Π) epistemic state of social evaluator.
- Every π induces expected majority count for all $x, y \in X$:

$$m_{(\theta,\pi)}(x, y) \;:=\; E_\pi[\#\{i : x \succ_i y\}].$$

- Every ambiguous belief Π thus induces interval of possible expected majority margins $\left[m^-_{(\theta,\Pi)}(x, y)\,,\; m^+_{(\theta,\Pi)}(x, y)\right]$.

Ex-Ante Condorcet Winners

- The **ex-ante majority relation** is defined by

$$xR_{(\theta,\Pi)}y \quad :\Longleftrightarrow \quad m^{-}_{(\theta,\Pi)}(x,y) \geq m^{-}_{(\theta,\Pi)}(y,x)$$
$$\Longleftrightarrow \quad m^{+}_{(\theta,\Pi)}(x,y) \geq m^{+}_{(\theta,\Pi)}(y,x).$$

- Majority intervals may overlap but are unambiguously ordered.
  - Choice based on $\alpha \min + (1-\alpha)\max$ ('Hurwicz pessimism – optimism index,' Luce & Raiffa) independent of α.
  - Independent of ambiguity attitude.
- **Ex-ante Condorcet winners**:

$$\mathrm{CW}(\theta,\Pi) \;:=\; \{x \in X \mid xR_{(\theta,\Pi)}y \text{ for all } y \in X\}.$$

- **Main finding**: in interesting cases, ex-ante CWs exist even if ex-post CWs fail to exist (e.g. in standard spatial voting).

Simple Examples

- Unrestricted beliefs: Π_{univ}
  - $xP_{(\theta,\Pi)}y$ iff more agents have top at x than top at y.
  - $\mathrm{CW}(\theta,\Pi) =$ plurality winners.

Simple Examples

- Unrestricted beliefs: Π_{univ}
  - $xP_{(\theta,\Pi)}y$ iff more agents have top at x than top at y.
  - $\mathrm{CW}(\theta,\Pi) =$ plurality winners.
- Convexity ('single-peakedness') on the line: Π_{co} on $X = \mathbb{R}$
  - For $x < y$,

    $xP_{(\theta,\Pi)}y$ iff more agents have top in $(-\infty, x]$ than in $[y, +\infty)$.
  - $\mathrm{CW}(\theta,\Pi) =$ medians.

Agenda

The Multi-Dimensional Problem

Now, X convex subset of $\mathbb{R}^L$, e.g.

$$X_{\mathrm{res}} = \left\{ x = (x^1, \dots, x^L) \in \mathbb{R}^L \ : \ \sum_{\ell=1}^{L} x^\ell = M \right\}$$

The Multi-Dimensional Problem

Now, X convex subset of $\mathbb{R}^L$, e.g.

$$X_{\mathrm{res}} = \left\{ x = (x^1, ..., x^L) \in \mathbb{R}^L \ : \ \sum_{\ell=1}^{L} x^\ell = M \right\}$$

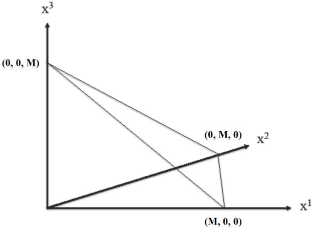

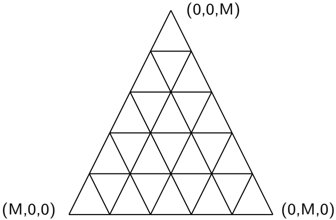

Convexity of Preferences

- Suppose Π_{co} consists of all possible priors on profiles of individually convex preferences (**'plain convex model'**).

Proposition

If individual tops $\{\theta_1, ..., \theta_n\}$ are in general position (no three collinear), then $\mathrm{CW}(\theta, \Pi_{\mathrm{co}})$ contains all plurality winners. Moreover, if there is a unique plurality winner θ_{i^*}, then $\mathrm{CW}(\theta, \Pi_{\mathrm{co}}) = \{\theta_{i^*}\}$.

Conundrum

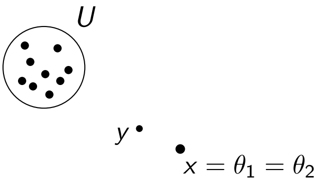

Conundrum

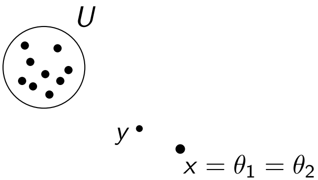

- Tops randomly drawn from U (in particular, in general position).
- Convexity gives no info about pref between y and x.
- With duplication, ex-ante Condorcet winner is x.
- In fact, possibly **all** voters with top in U prefer x to y.
- But alternative x's claim for winner is like 'grasping at straws.'
- Indeed, prefs must be very specifically and individually tailored.

Agenda

Separating Top from Preference Structure

- **Intuition:** tops contain no info about preference 'structure.'
- **Flexible workhorse:** quadratic preferences, i.e.,

$$u_{\theta_i}(x) = -(x - \theta_i)^T \cdot Q \cdot (x - \theta_i)$$

with positive definite Q.

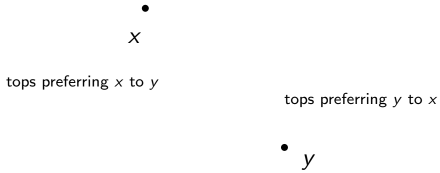

Separating Top from Preference Structure

- **Intuition:** tops contain no info about preference 'structure.'
- **Flexible workhorse:** quadratic preferences, i.e.,

$$u_{\theta_i}(x) = -(x - \theta_i)^T \cdot Q \cdot (x - \theta_i)$$

with positive definite Q.

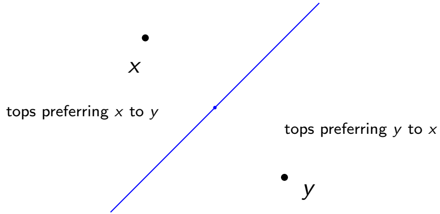

Separating Top from Preference Structure

- **Intuition:** tops contain no info about preference 'structure.'
- **Flexible workhorse:** quadratic preferences, i.e.,

$$u_{\theta_i}(x) = -(x - \theta_i)^T \cdot Q \cdot (x - \theta_i)$$

with positive definite Q.

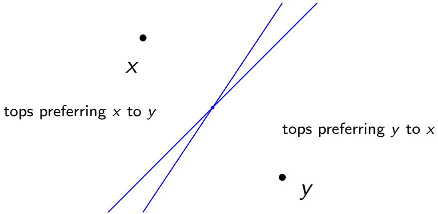

Separating Top from Preference Structure

- **Intuition:** tops contain no info about preference 'structure.'
- **Flexible workhorse:** quadratic preferences, i.e.,

$$u_{\theta_i}(x) = -(x-\theta_i)^T \cdot Q \cdot (x-\theta_i)$$

with positive definite Q.

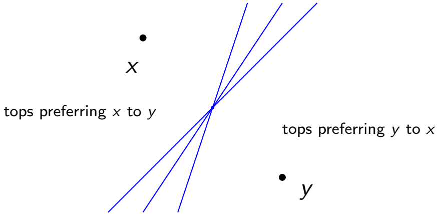

Separating Top from Preference Structure

- **Intuition:** tops contain no info about preference 'structure.'
- **Flexible workhorse:** quadratic preferences, i.e.,

$$u_{\theta_i}(x) = -(x - \theta_i)^T \cdot Q \cdot (x - \theta_i)$$

with positive definite Q.

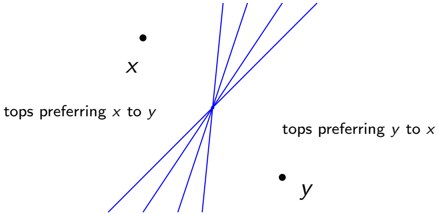

Separating Top from Preference Structure

- **Intuition:** tops contain no info about preference 'structure.'
- **Flexible workhorse:** quadratic preferences, i.e.,

$$u_{\theta_i}(x) = -(x-\theta_i)^T \cdot Q \cdot (x-\theta_i)$$

with positive definite Q.

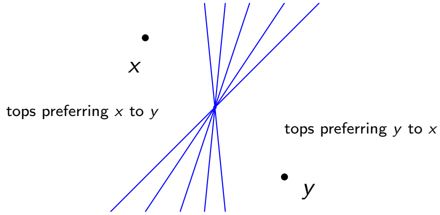

Separating Top from Preference Structure

- **Intuition:** tops contain no info about preference 'structure.'
- **Flexible workhorse:** quadratic preferences, i.e.,

$$u_{\theta_i}(x) = -(x - \theta_i)^T \cdot Q \cdot (x - \theta_i)$$

with positive definite Q.

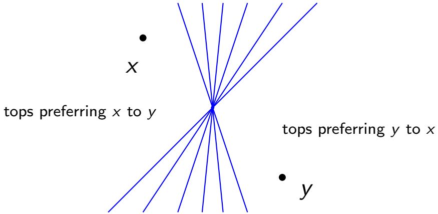

Quadratic Model as Approximation of Convex Model

Π_{quad} set of all priors on profiles of quadratic preferences.

Proposition

Π_{quad} *induces same ex-ante majority relation as* Π_{co}. *In particular,* $\mathrm{CW}(\theta, \Pi_{\text{quad}}) = \mathrm{CW}(\theta, \Pi_{\text{co}})$.

The Symmetric Quadratic Model

- Quadratic model allows one to formulate above intuition that "tops contain no info about preference structure" as follows:
- A prior π on profiles of quadratic preferences is **symmetric** if, for all i, j,

$$\pi_{\mathcal{Q}_i} = \pi_{\mathcal{Q}_j}.$$

- $\Pi_{\mathrm{quad}}^{sym}$ set of all **symmetric** priors on profiles of quadratic preferences.

Question

$$\mathrm{CW}(\theta, \Pi_{\mathrm{quad}}^{sym}) \ ?$$

The Tukey Median

- *Tukey depth* of x at θ: $\mathfrak{d}(x,\theta) := \min_{H \ni x} \#(\theta \cap H)$.
- Let $\mathfrak{d}(\theta) := \max_{x \in X} \mathfrak{d}(x,\theta)$.
- The **Tukey median** (Tukey, 1975) is defined by

$$T(\theta) := \{x \in X : \mathfrak{d}(x,\theta) = \mathfrak{d}(\theta)\}.$$

The Tukey Median

- *Tukey depth* of x at θ: $\mathfrak{d}(x,\theta) := \min_{H \ni x} \#(\theta \cap H)$.
- Let $\mathfrak{d}(\theta) := \max_{x \in X} \mathfrak{d}(x,\theta)$.
- The **Tukey median** (Tukey, 1975) is defined by

$$T(\theta) := \{x \in X : \mathfrak{d}(x,\theta) = \mathfrak{d}(\theta)\}.$$

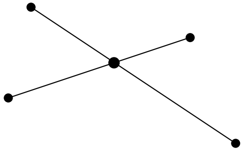

The Tukey Median

- *Tukey depth* of x at θ: $\mathfrak{d}(x,\theta) := \min_{H \ni x} \#(\theta \cap H)$.
- Let $\mathfrak{d}(\theta) := \max_{x \in X} \mathfrak{d}(x,\theta)$.
- The **Tukey median** (Tukey, 1975) is defined by

$$T(\theta) := \{x \in X : \mathfrak{d}(x,\theta) = \mathfrak{d}(\theta)\}.$$

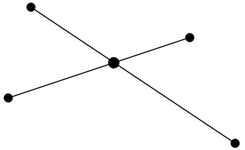

- Tukey median is one (affinely invariant) multi-dimensional median see, e.g., survey by Rousseeuw & Hubert (2017).

Example: The Pentagon

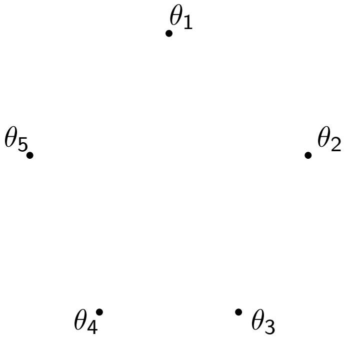

Example: The Pentagon

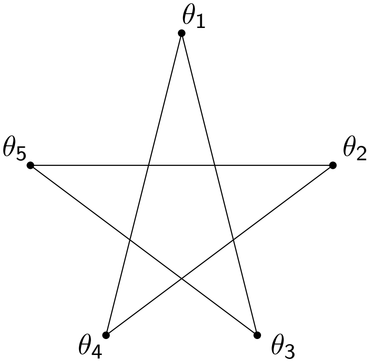

Example: The Pentagon

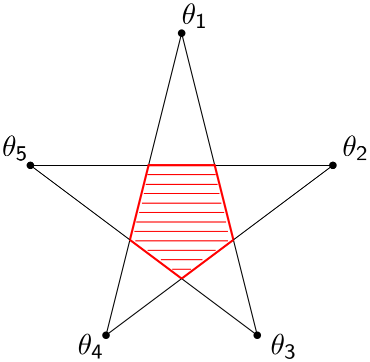

Strict Tukey Median and Main Result

- Denote by $\mathcal{H}_x^* := \{H \ni x : \#(\theta \cap H) = \mathfrak{d}(\theta)\}$ and let

$$T^*(\theta) \;:=\; \{x \in T(\theta) \mid \text{ for no } y \in T(\theta),\ \mathcal{H}_y^* \subsetneq \mathcal{H}_x^*\}.$$

- By construction, $T^*(\theta) \subseteq T(\theta)$.
- $T^*(\theta)$ **strict Tukey median**.

Strict Tukey Median and Main Result

- Denote by $\mathcal{H}_x^* := \{H \ni x : \#(\theta \cap H) = \mathfrak{d}(\theta)\}$ and let

$$T^*(\theta) \;:=\; \{x \in T(\theta) \mid \text{ for no } y \in T(\theta),\ \mathcal{H}_y^* \subsetneq \mathcal{H}_x^*\}.$$

- By construction, $T^*(\theta) \subseteq T(\theta)$.
- $T^*(\theta)$ **strict Tukey median**.

Theorem

For all profiles θ, the strict Tukey median is non-empty and

$$\mathrm{CW}(\theta, \Pi_{\mathrm{quad}}^{sym}) \;=\; T^*(\theta).$$

Sketch of Proof

Step 1

The symmetric quadratic model $\Pi^{sym}_{\mathrm{quad}}$ *induces the same ex-ante majority relation as the* **uniform** *quadratic model* $\Pi^{unif}_{\mathrm{quad}}$ *consisting of all priors that assign full mass to uniform profiles*

$$[(\theta_1, Q), (\theta_2, Q), ..., (\theta_n, Q)]$$

for some common quadratic from Q*.*

Sketch of Proof

Step 2

This implies that, under symmetry, the ex-ante majority relation coincides locally with **relative Tukey depth:**

$$xR_{(\theta,\mathrm{Tuk})}y \;:\Leftrightarrow\; \min_{H\ni x, H\not\ni y} \#(\theta\cap H) \geq \min_{H\ni y, H\not\ni x} \#(\theta\cap H).$$

Sketch of Proof

Step 2

This implies that, under symmetry, the ex-ante majority relation coincides locally with **relative Tukey depth:**

$$xR_{(\theta,\mathrm{Tuk})}y \;\; :\Leftrightarrow \;\; \min_{H\ni x, H\not\ni y} \#(\theta \cap H) \geq \min_{H\ni y, H\not\ni x} \#(\theta \cap H).$$

Observe that this solves above conundrum:

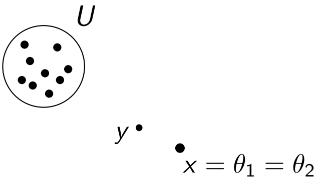

Sketch of Proof

Step 2

This implies that, under symmetry, the ex-ante majority relation coincides locally with **relative Tukey depth:**

$$xR_{(\theta,\mathrm{Tuk})}y \ :\Leftrightarrow \ \min_{H\ni x, H\not\ni y} \#(\theta\cap H) \geq \min_{H\ni y, H\not\ni x} \#(\theta\cap H).$$

Observe that this solves above conundrum:

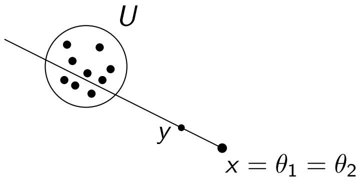

Sketch of Proof

Step 3

Now show that the strict Tukey median consists of the maximal elements of the relative Tukey depth:

$$T^*(\theta) \ = \ \left\{x \in X : xR_{(\theta,\mathrm{Tuk})}y \text{ for all } y \in X\right\}.$$

Step 4

Finally, show that $\{x \in X : xR_{(\theta,\mathrm{Tuk})}y \text{ for all } y \in X\}$ *is non-empty by appeal to Zorn's Lemma.*

Example: Relative vs.Absolute Tukey Depth

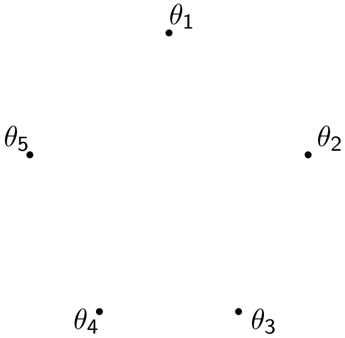

Example: Relative vs.Absolute Tukey Depth

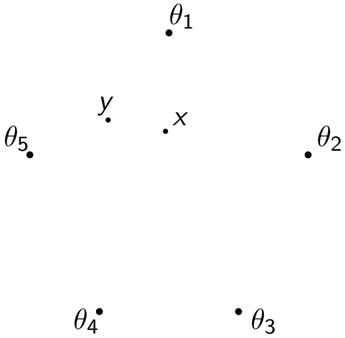

Example: Relative vs.Absolute Tukey Depth

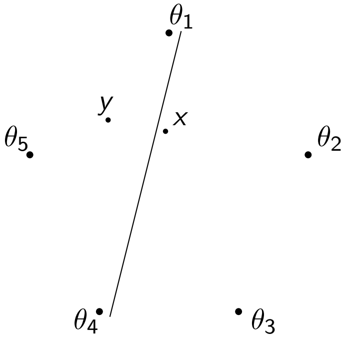

Minimal depth of x relative to y: $\min_{x \in H \not\ni y} \#(\theta \cap H) = 2$.

Example: Relative vs. Absolute Tukey Depth

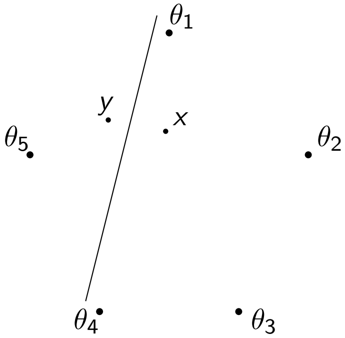

Minimal depth of y relative to x: $\min_{y \in H \not\ni x} \#(\theta \cap H) = 1$.

Example: Relative vs. Absolute Tukey Depth

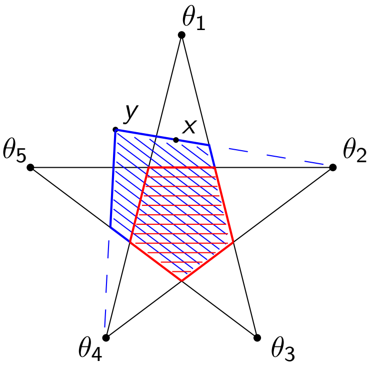

Example: Relative vs. Absolute Tukey Depth

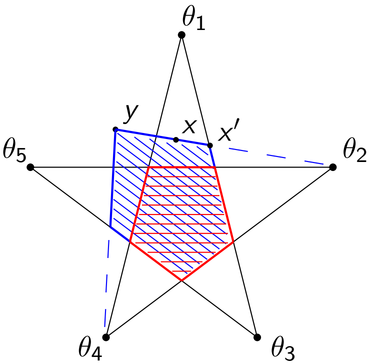

Upper contour set of y (in terms of relative depth) is not open!

Conclusion

Thanks for your attention!!

The Separably Convex Model

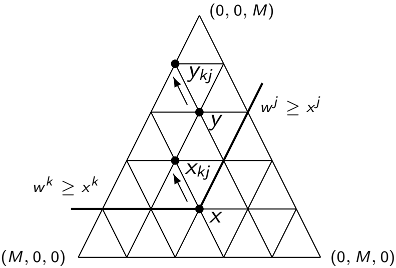

Example: All pref's representable by $u(x) = \sum_\ell u^\ell(x^\ell)$ with concave u^ℓ.

The Separably Convex Model

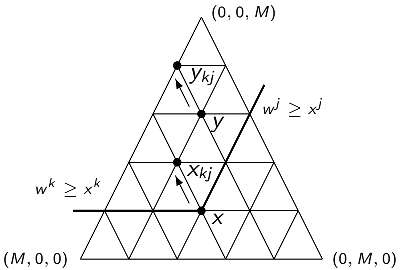

Example: All pref's representable by $u(x) = \sum_\ell u^\ell(x^\ell)$ with concave u^ℓ.

Observation

x is preferred to y by all voters with top θ_i iff x on shortest L_1-path connecting θ_i and y.

Global vs.Local Ex-Ante Condorcet Winners

The ex-ante majority relation under separably convex preferences may have cycles between 'distant' alternatives ...

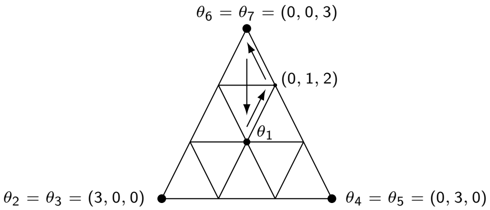

Global vs.Local Ex-Ante Condorcet Winners

The ex-ante majority relation under separably convex preferences may have cycles between 'distant' alternatives ...

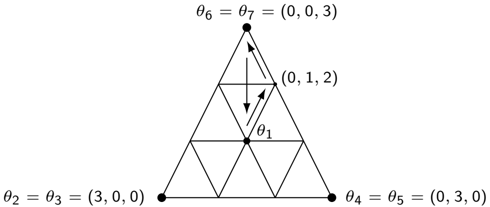

... but **no local cycles**. Indeed, we have the following result:

Separably Convex Model and the L_1-median

Definition

The L_p-**median** is the choice correspondence that selects, for all profiles θ,

$$\arg\min_{x \in X} \sum_{i=1}^{n} ||\theta_i - x||_p.$$

Theorem

For all θ,

$$\mathrm{CW}^{loc}(\theta, \Pi_{\mathrm{sepco}}) \; = \; L_1\text{-}median.$$