Optimizing Threshold-Based Fairness

John Hooker,* Peter Zhang
Carnegie Mellon University
Özgün Elçi
Amazon

**Presenter*

INFORMS 2025

Modeling Fairness

Some **fairness** applications:

- Health care resources.
- Facility location
    - Ambulance, fire stations
- Transportation.
    - Bus stops
- Telecommunications.
- Traffic signal timing
- Disaster recovery

Modeling Fairness

Balancing **efficiency** and **fairness**

- **Location:** average vs. worst-case service level
- **Telecom** or **traffic signals**: thruput vs. worst-case wait time
- **Disaster recovery**: average vs. worst-case response time.
    - For example, power restoration after storm

Modeling Fairness

Optimization models normally maximize **utility**.

- Utility = wealth, health, revenue, negative cost, service level
- This can lead to **very unfair** resource distribution**.**

Alternative: maximize a **social welfare function** (SWF)

- Subject to problem constraints.
- The SWF accounts for both **total utility** and **fairness**.
- May have a parameter that governs the **trade-off** between them.

Modeling Fairness

A widely used SWF is **alpha fairness**.

- Parameter α = degree of fairness.
- Larger $\alpha \Rightarrow$ greater fairness
- $\alpha = 0$: **utilitarian** (max total utility)
- $\alpha = \infty$: **maximin** (max smallest utility)
    - Inspired by **Difference Principle** of John Rawls.
- $\alpha = 1$: **proportional fairness** (Nash bargaining solution)

Above: John Rawls, 1921-2002

John Nash, 1928-2015

Modeling Fairness

Problem with alpha fairness.

- How to **interpret** α in practice?
- How to **choose an** α that suits the application?

We propose two **threshold SWFs** with more-or-less interpretable trade-off parameters.

- **Utility threshold** SWF.
  - Which has an integer programming formulation.
- **Equity threshold** SWF.
  - Which has a linear programming formulation.

Alpha Fairness

Larger $\alpha \geq 0$ corresponds to greater fairness

$$W_\alpha(\boldsymbol{u}) = \begin{cases} \dfrac{1}{1-\alpha} \displaystyle\sum_i u_i^{1-\alpha} & \text{for } \alpha \geq 0,\ \alpha \neq 1 \\ \displaystyle\sum_i \log(u_i) & \text{for } \alpha = 1 \end{cases}$$

Mo & Walrand 2000; Verloop, Ayesta & Borst 2010

- $\boldsymbol{u_i}$ = utility allocated to stakeholder *i*
- **Nonlinear** but **concave**.
- **Used** in engineering.
- Can be **derived** from certain axioms.

Lan & Chiang 2011

Alpha Fairness

Contours for 2 stakeholders

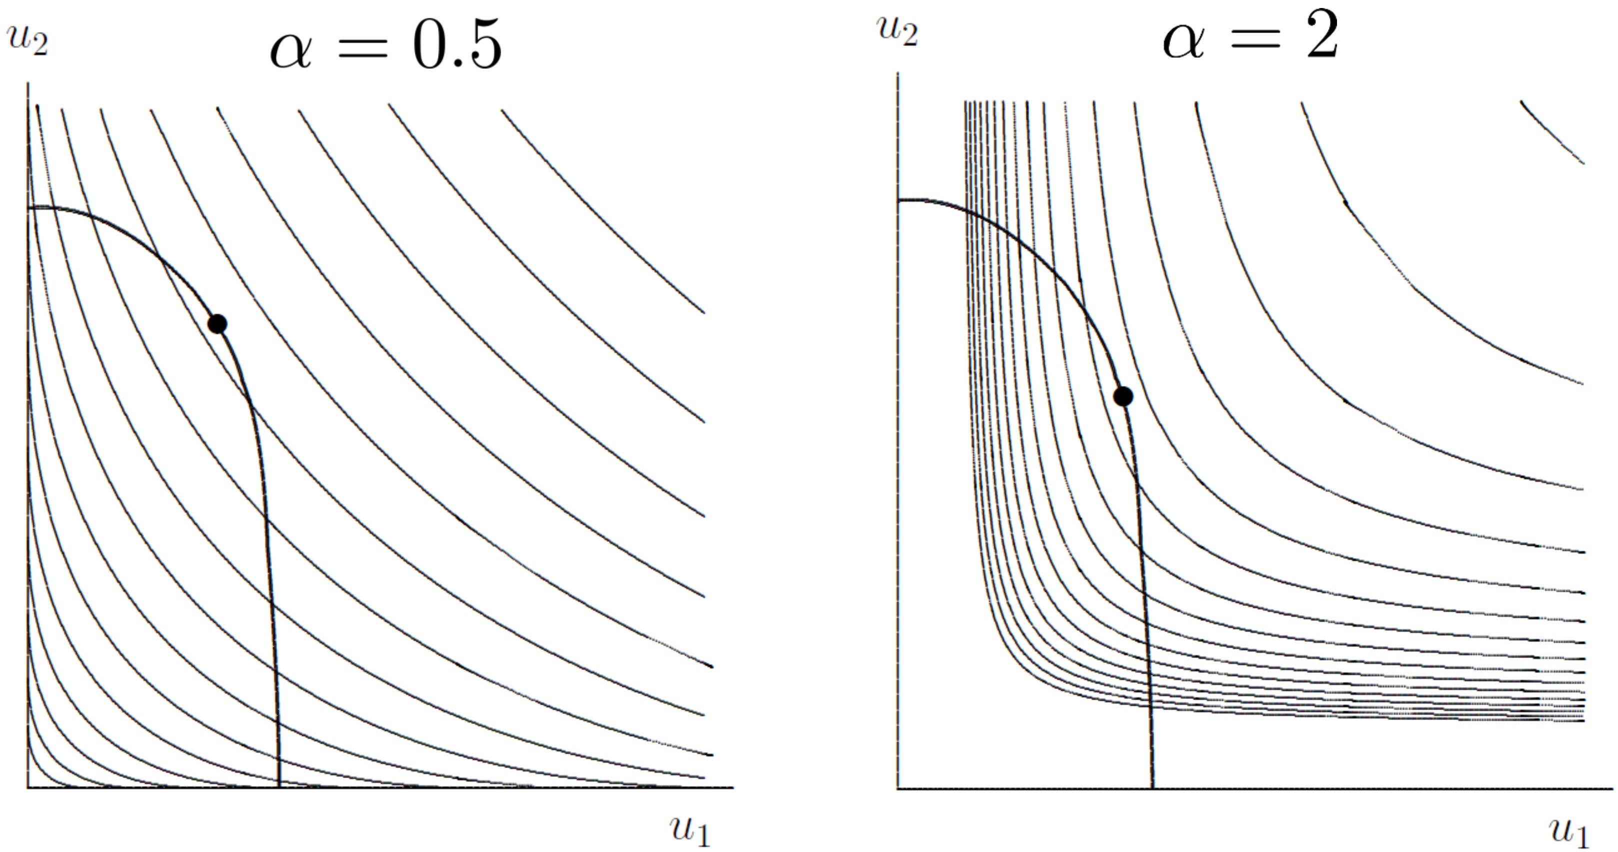

When α > 1, egalitarian distribution can have same social welfare as extreme inequality.

Alpha Fairness

We study properties of SWFs when maximized subject to a **budget constraint**.

$$\sum_i a_i u_i \leq B$$

a_i = unit utility cost of stakeholder *i*
$1/a_i$ = **conversion efficiency** of stakeholder *i*

Solution for alpha fairness has **closed form:**

$$u_i = \frac{B}{a_i^{1/\alpha} \sum_j a_j^{1-1/\alpha}}, \text{ all } i$$

Alpha Fairness

Example:

Maximum alpha fairness subject to budget constraint

$$5u_1 + 6u_2 + \cdots + 12u_8 \leq 100$$

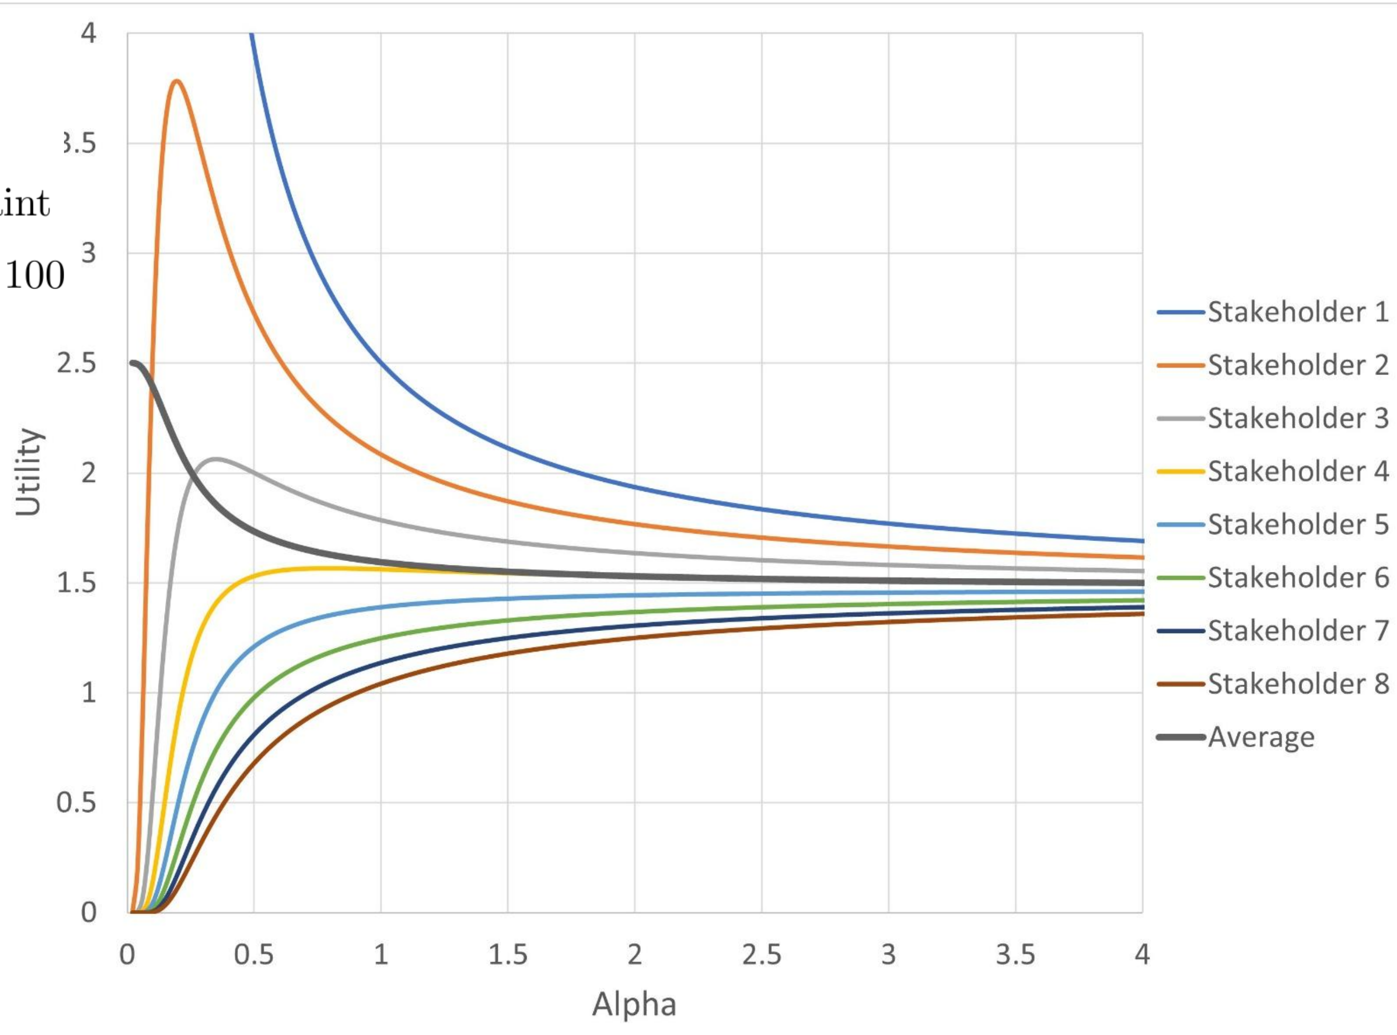

When $\alpha = 0$, most efficient stakeholder gets everything.

Unclear how to choose α in practice

Example of Alpha Fairness

Investment in electric generating capacity and transmission - Liberia

- Pure efficiency objective neglects the hinterland.
- Emphasis on fairness reduces total benefit.

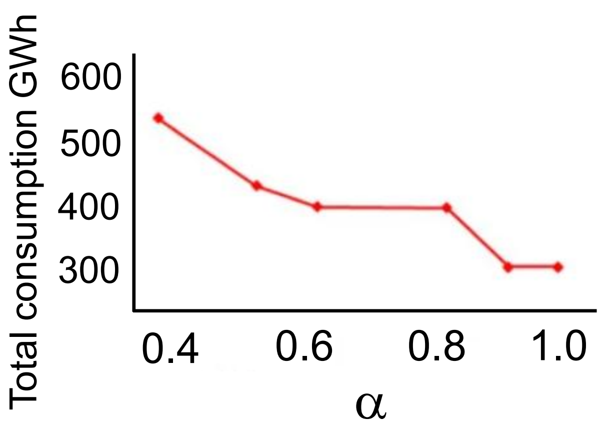

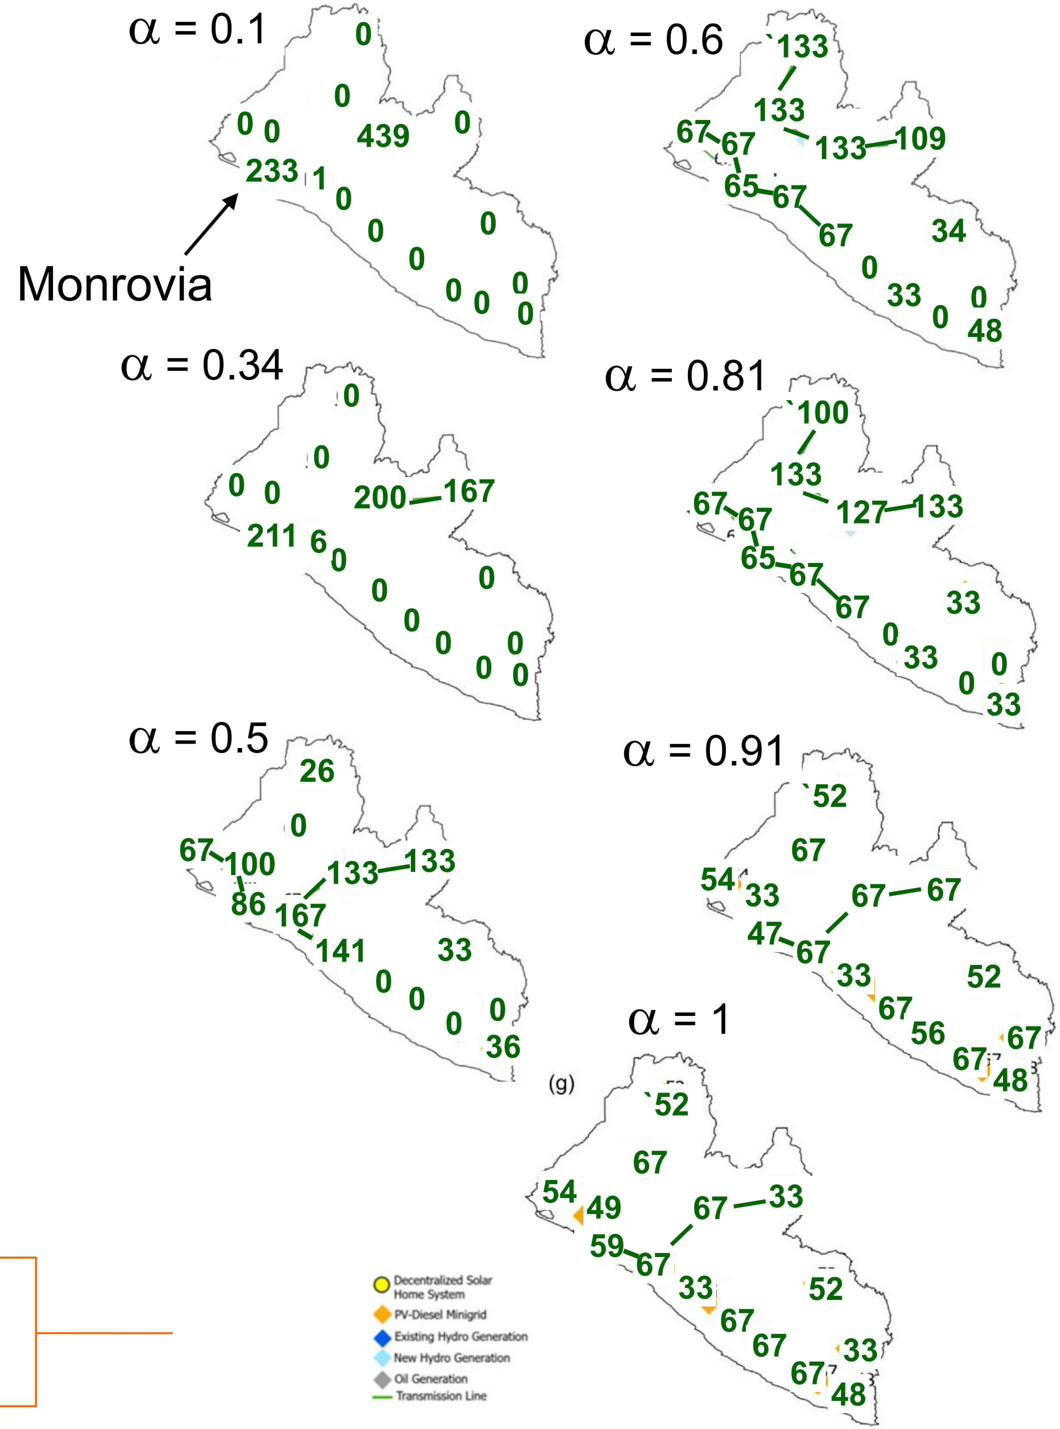

Sackey, Nock, Cao, Armanios, Davis 2023

Example of Alpha Fairness

Investment in electric generating capacity and transmission - Liberia

- Pure efficiency objective neglects the hinterland.
- Emphasis on fairness reduces total benefit.
- Elicited value of **α = 0.81**
    - Based on survey of engineering graduate students.

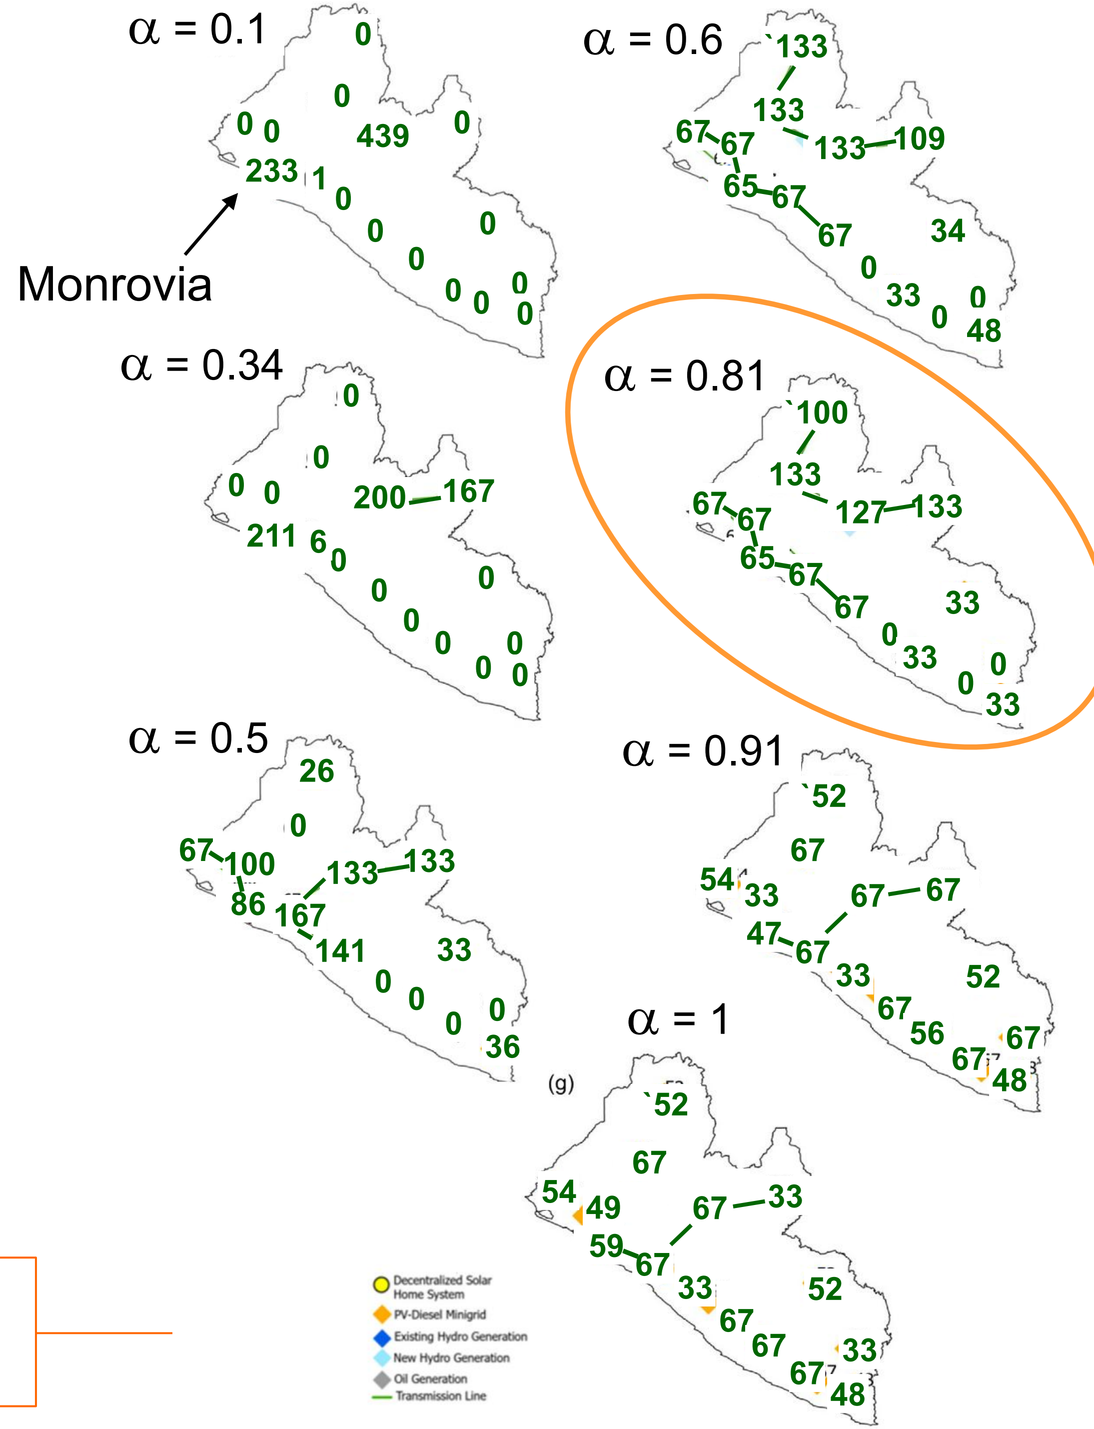

Sackey, Nock, Cao, Armanios, Davis 2023

Threshold Methods

Combining utility and maximin

Utility threshold: Use a maximin criterion until the **utility cost becomes too great**, then switch some stakeholders to utilitarian.

- Fairness is a primary concern, but without sacrificing too much utility.
- As in a medical context, emergency facility location, task assignment.

Equity threshold: Use a utilitarian criterion until the **inequity becomes too great**, then switch some stakeholders to maximin.

- Efficiency is the primary concern, but without excessive sacrifice by any individual.
- As in telecommunications, disaster recovery, traffic control..

Williams & Cookson 2000

Utility Threshold

Optimal solution

Maximin solution results in
too much utility sacrifice
for stakeholder 2

Feasible set

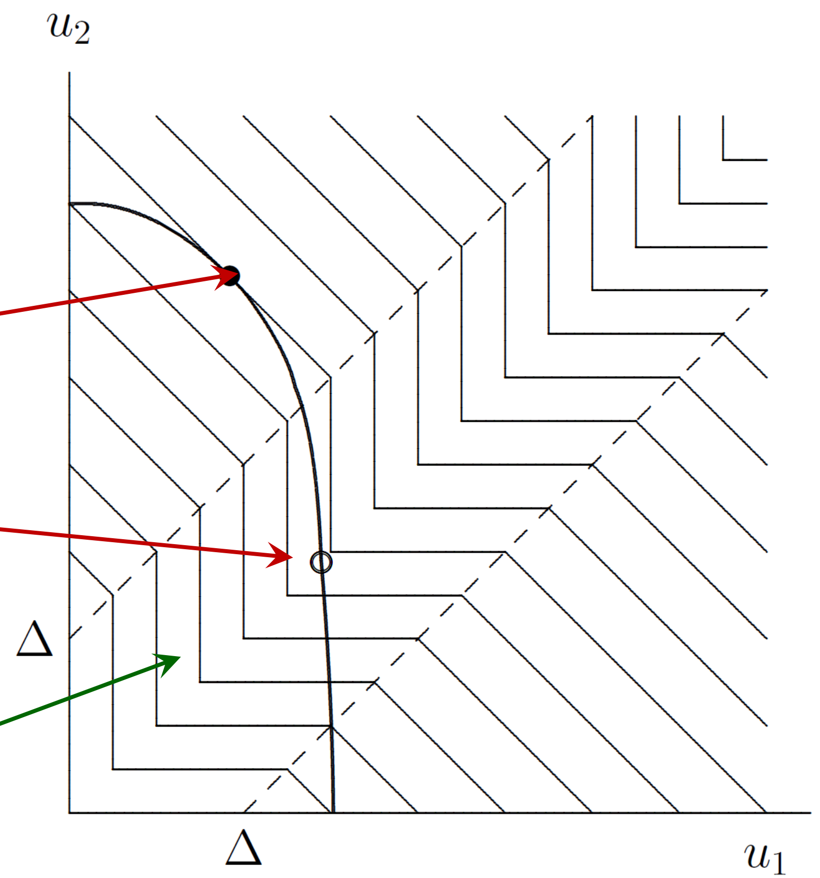

A. Williams & Cookson 2000

$$W(u_1, u_2) = \begin{cases} u_1 + u_2, & \text{if } |u_1 - u_2| \geq \Delta \\ 2\min\{u_1, u_2\} + \Delta, & \text{otherwise} \end{cases}$$

Utility Threshold

Generalization to *n* stakeholders

$$W(\boldsymbol{u}) = (n-1)\Delta + \sum_{i=1}^{n} \max\left\{u_i - \Delta, u_{\min}\right\}$$

$$\text{where } u_{\min} = \min_i\{u_i\}$$

JH & H.P. Williams 2012

- $\Delta = 0$ corresponds to utilitarian criterion, $\Delta = \infty$ to maximin.
- Δ is chosen so that individuals with utility **within Δ of smallest** are sufficiently deprived to **deserve priority**.

Utility Threshold

MILP model

$$\max\left\{ z \;\middle|\; \begin{array}{l} z \le (n-1)\Delta + \sum_i v_i \\ u_i - \Delta \le v_i \le u_i - \delta_i \Delta, \;\; \text{all } i \\ w \le v_i \le w + \delta_i (M - \Delta), \;\; \text{all } i \\ u_i \ge 0, \delta_i \in \{0,1\}, \;\; \text{all } i \\ \text{other problem constraints} \end{array} \right\}$$

The formulation is **sharp** (describes convex hull of hypograph).

JH & H.P. Williams 2012

Utility Threshold

Theorem. Solution subject to a budget constraint is purely **utilitarian** when

$$\Delta \leq B\left(\frac{1}{\min_i\{a_i\}} - \frac{n}{\sum_i a_i}\right)$$

and otherwise is purely **maximin**.

Elçi, JH, Zhang 2025

Utility Threshold

Purely utilitarian solution

Stakeholders have **very different** costs, or Δ is **small**.

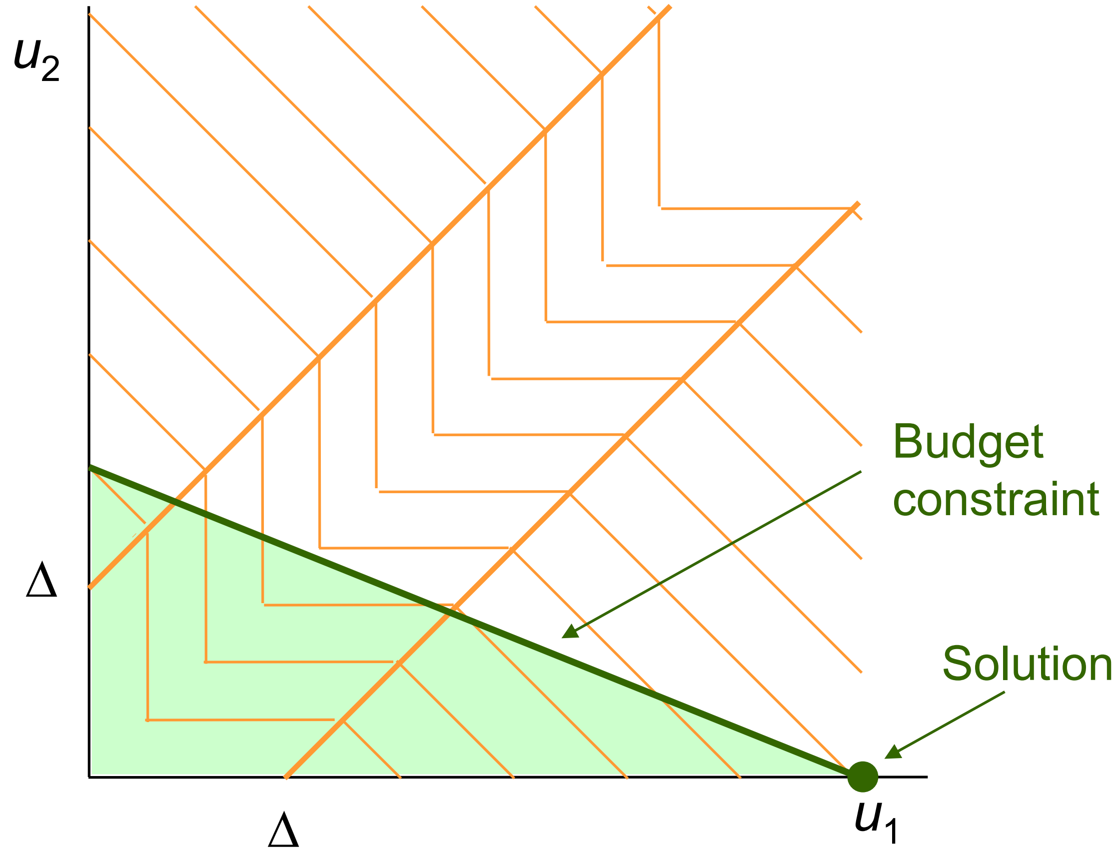

Utility Threshold

Purely maximin solution

Stakeholders have **similar** costs,
or Δ is **large**.

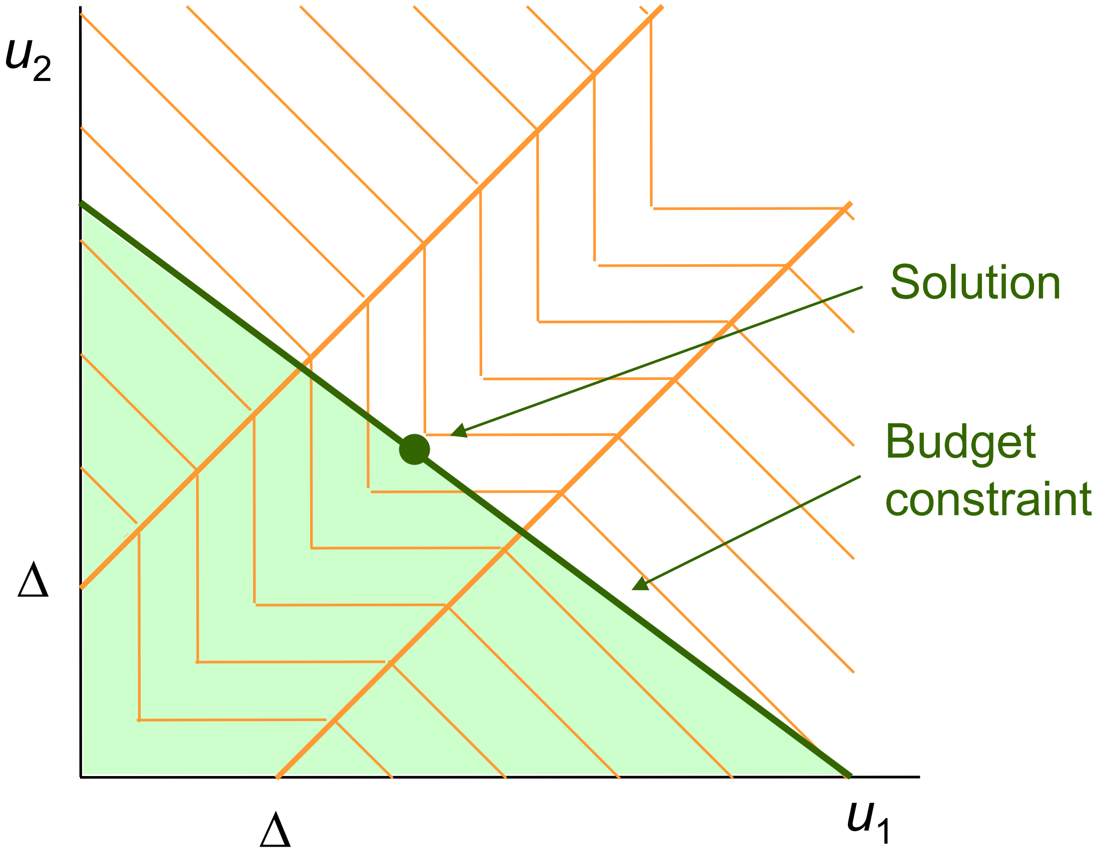

Utility Threshold

Theorem. When maximizing the SWF subject to a **budget constraint and upper bounds** d_i at most one utility is **strictly between** its upper bound and the smallest utility.

Here, **one** utility u_2 is **strictly between** upper bound d_2 and the smallest utility u_1.

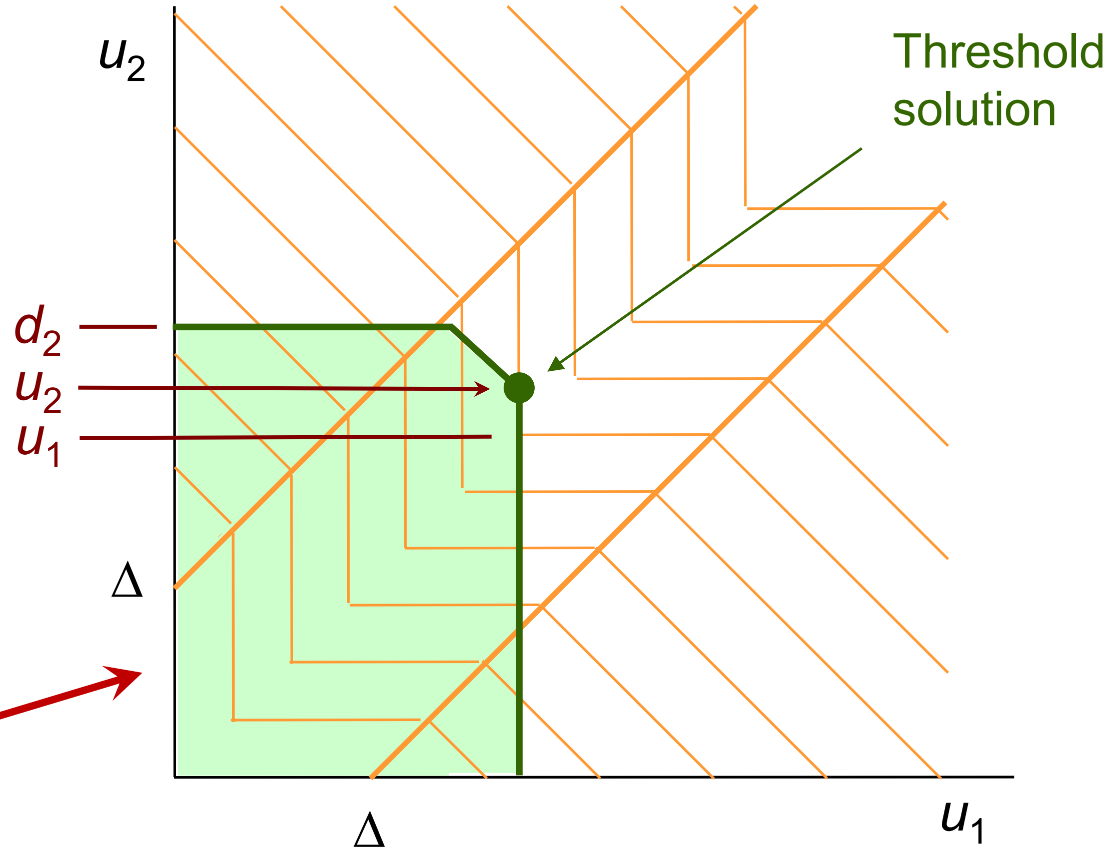

Equity Threshold

Utilitarian solution
leaves person 1
overly deprived

Optimal solution

Feasible set

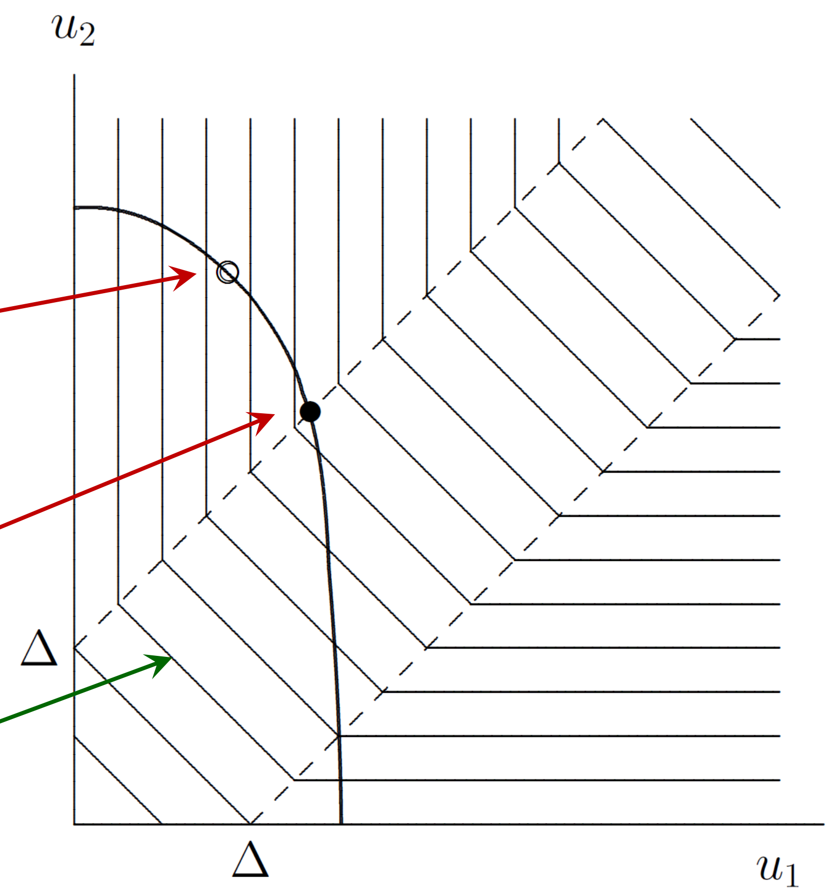

Williams & Cookson 2000

$$W(u_1, u_2) = \begin{cases} 2\min\{u_1, u_2\} + \Delta, & \text{if } |u_1 - u_2| \geq \Delta \\ u_1 + u_2, & \text{otherwise} \end{cases}$$

Equity Threshold

Can be viewed as a piecewise linear approximation of alpha fairness, with interpretable Δ.

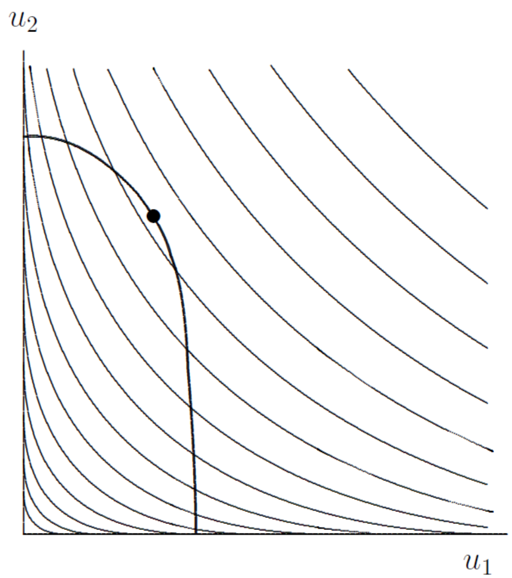

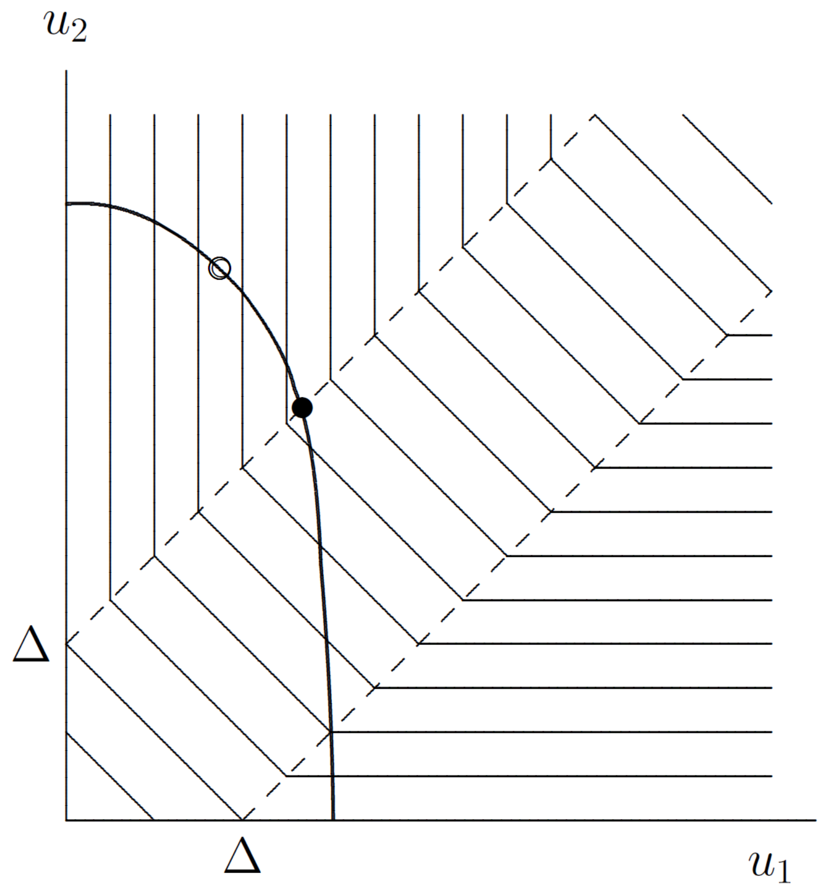

$$W(u_1, u_2) = \begin{cases} 2\min\{u_1, u_2\} + \Delta, & \text{if } |u_1 - u_2| \geq \Delta \\ u_1 + u_2, & \text{otherwise} \end{cases}$$

Equity Threshold

Generalization to *n* stakeholders

$$W(\boldsymbol{u}) = n\Delta + \sum_{i=1}^{n} \min\{u_i - \Delta, u_{min}\}$$

Elçi, JH, Zhang 2025

Interpretation of Δ:

- Stakeholders with utility that is Δ **greater than the worst-off** are considered **well-off**, and
- Raising the utility of **well-off** stakeholders improves social welfare only if **worst-off benefit an equal amount**.
- Values **reversed**: $\Delta = \infty$ corresponds to utilitarian, $\Delta = 0$ to maximin.

Elçi, JH, Zhang 2025

Equity Threshold

LP model

$$\max \left\{ n\Delta + \sum_i v_i \;\middle|\; \begin{array}{l} w \le u_i,\; v_i \le w,\; \text{all } i \\ v_i \le u_i - \Delta,\; \text{all } i \\ w \ge 0 \\ \text{other problem constraints} \end{array} \right\}$$

Elçi, JH, Zhang 2025

Equity Threshold

Example:

Maximum equity threshold SWF subject to budget constraint

$$5u_1 + 6u_2 + \cdots + 12u_8 \leq 100$$

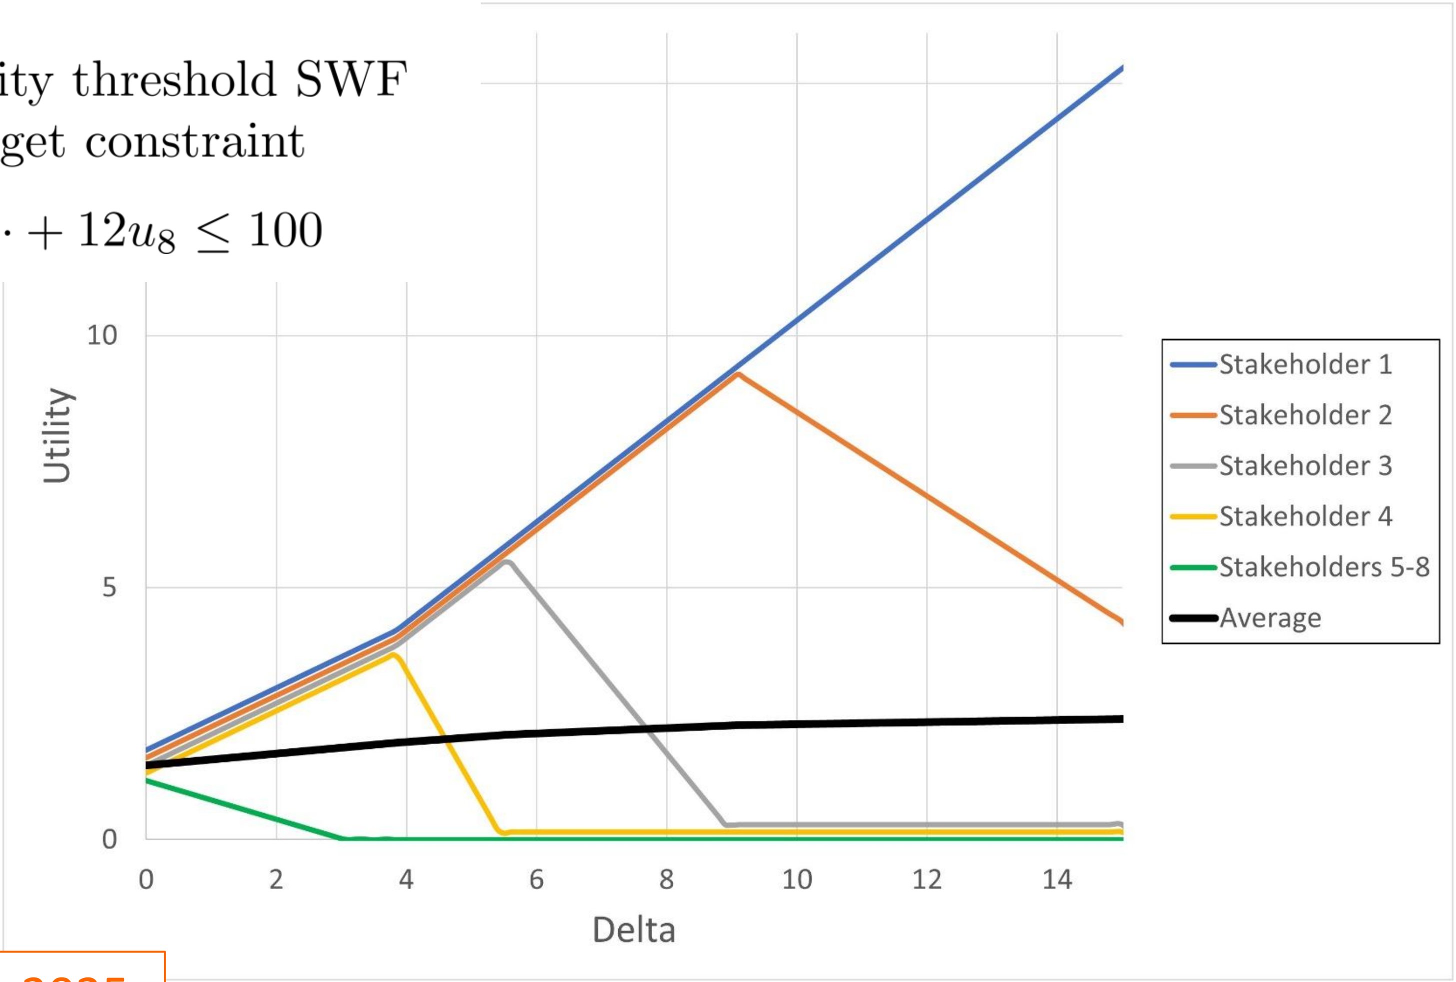

Elçi, JH, Zhang 2025

Utility Threshold

Theorem. Maximizing the equity threshold SWF with parameter Δ and subject to budget constraint is **equivalent** to maximizing **total utility** subject to budget constraint and a **max range of Δ**.

Elçi, JH, Zhang 2025

Utility Threshold with Leximax

Combines utility and **leximax** to provide **more sensitivity to equity**.

SWFs $W_1, \ldots, W_n$ are maximized sequentially, where W_1 is the utility threshold SWF defined earlier, and W_k for $k \geq 2$ is

$$W_k(\boldsymbol{u}) = \sum_{i=1}^{k-1}(n-i+1)u_{\langle i\rangle} + (n-k+1)\min\left\{u_{\langle 1\rangle} + \Delta, u_{\langle k\rangle}\right\} + \sum_{i=k}^{n}\max\left\{0,\ u_{\langle i\rangle} - u_{\langle 1\rangle} - \Delta\right\}$$

where $u_{\langle 1\rangle}, \ldots, u_{\langle n\rangle}$ are $u_1, \ldots, u_n$ in nondecreasing order.

Chen & JH 2021

Utility Threshold with Leximax

MILP model

for maximizing $W_k(\boldsymbol{u})$

$$\max \left\{ z \;\middle|\; \begin{array}{l} z \leq (n-k+1)\sigma + \sum_{i=I_k} v_i \\ 0 \leq v_i \leq M\delta_i, \quad i \in I_k \\ v_i \leq u_i - \bar{u}_{i_1} - \Delta + M(1-\delta_i), \quad i \in I_k \\ \sigma \leq \bar{u}_{i_1} + \Delta, \quad \sigma \leq w \\ w \leq u_i \leq w + M(1-\epsilon_i), \quad i \in I_k \\ \sum_{i \in I_k} \epsilon_i = 1, \quad w \geq \bar{u}_{i_{k-1}} \\ u_i - \bar{u}_{i_1} \leq M, \quad \delta_i, \epsilon_i \in \{0,1\}, \quad i \in I_k \\ \text{other problem constraints} \end{array} \right\}$$

The problem $\max\ W_k(\boldsymbol{u})$ determines value of smallest u_i in its solution. Also $\bar{u}_{i_j} = j$th u_i determined, $I_k = \{i \mid u_i \text{ not yet determined}\}$.

Chen & JH 2021

Utility Threshold with Leximax

Example:

Maximum utility threshold SWF with leximax subject to budget constraint

$$5u_1 + 6u_2 + \cdots + 12u_8 \leq 100$$

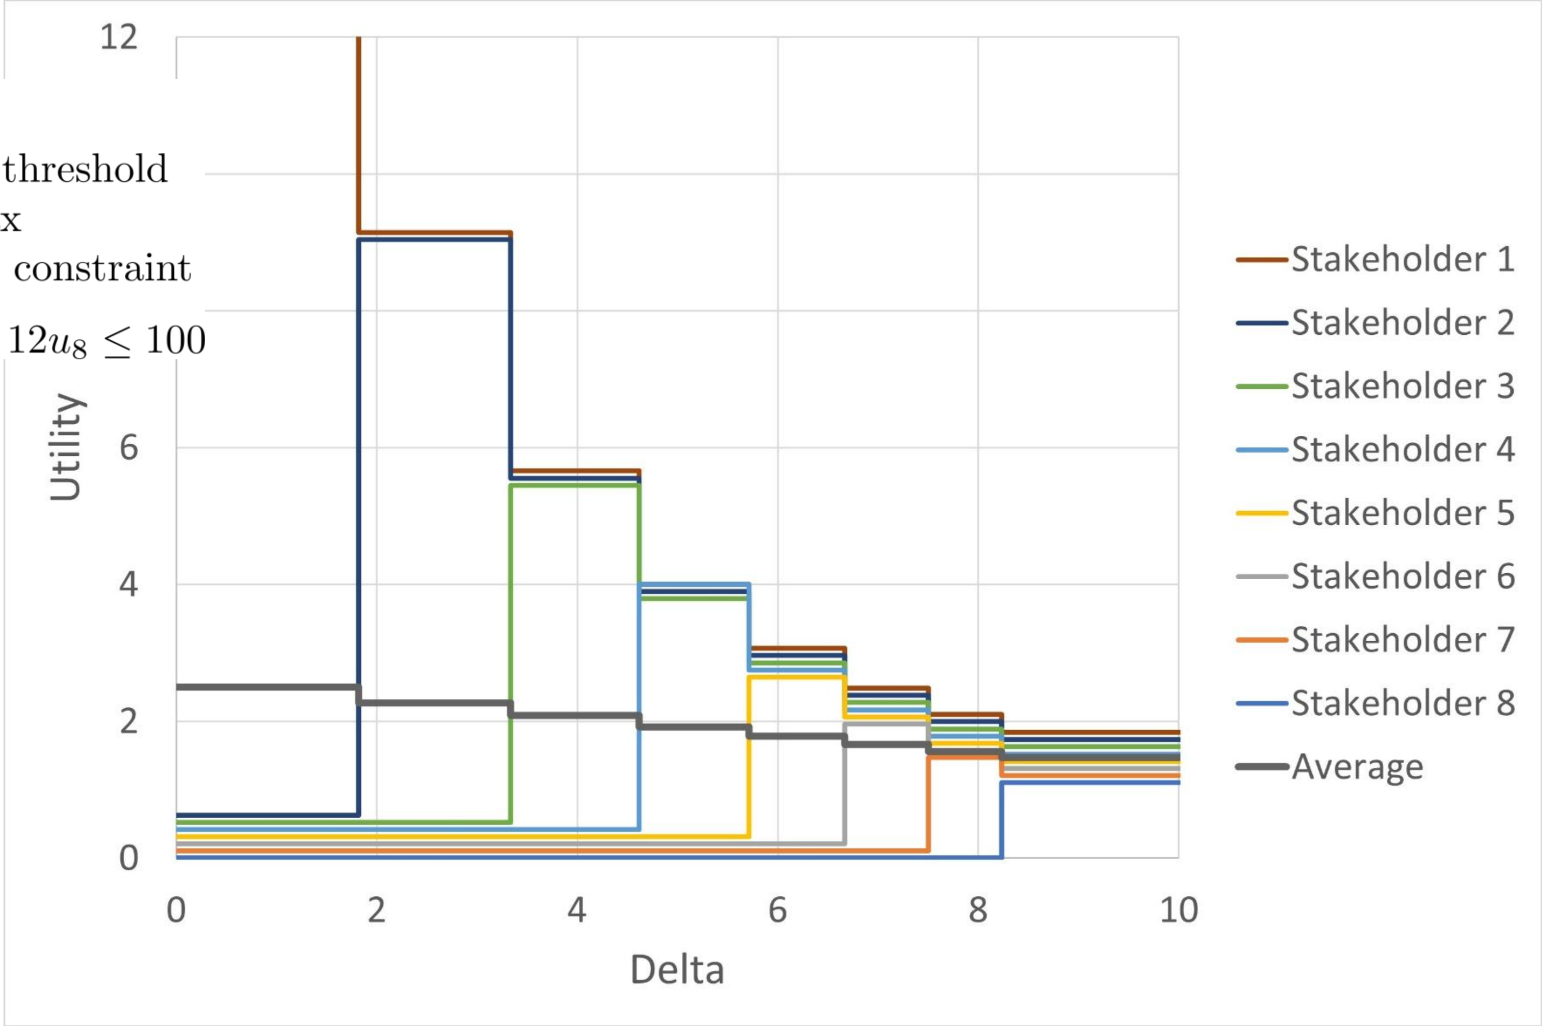

Example of Utility Threshold with Leximax

- Select **earthquake shelter** locations in Istanbul.
- Utility = negative **distance** of each neighborhood to **nearest shelter**.
- 50 neighborhoods, 50 potential shelter locations.
- Solution time = 1 to 18 seconds for each value of Δ.

Problem due to Mostajabdaveh, Gutjahr & Salman 2019

Threshold SWF

Utility + leximax

Horizontal lines correspond to neighborhoods

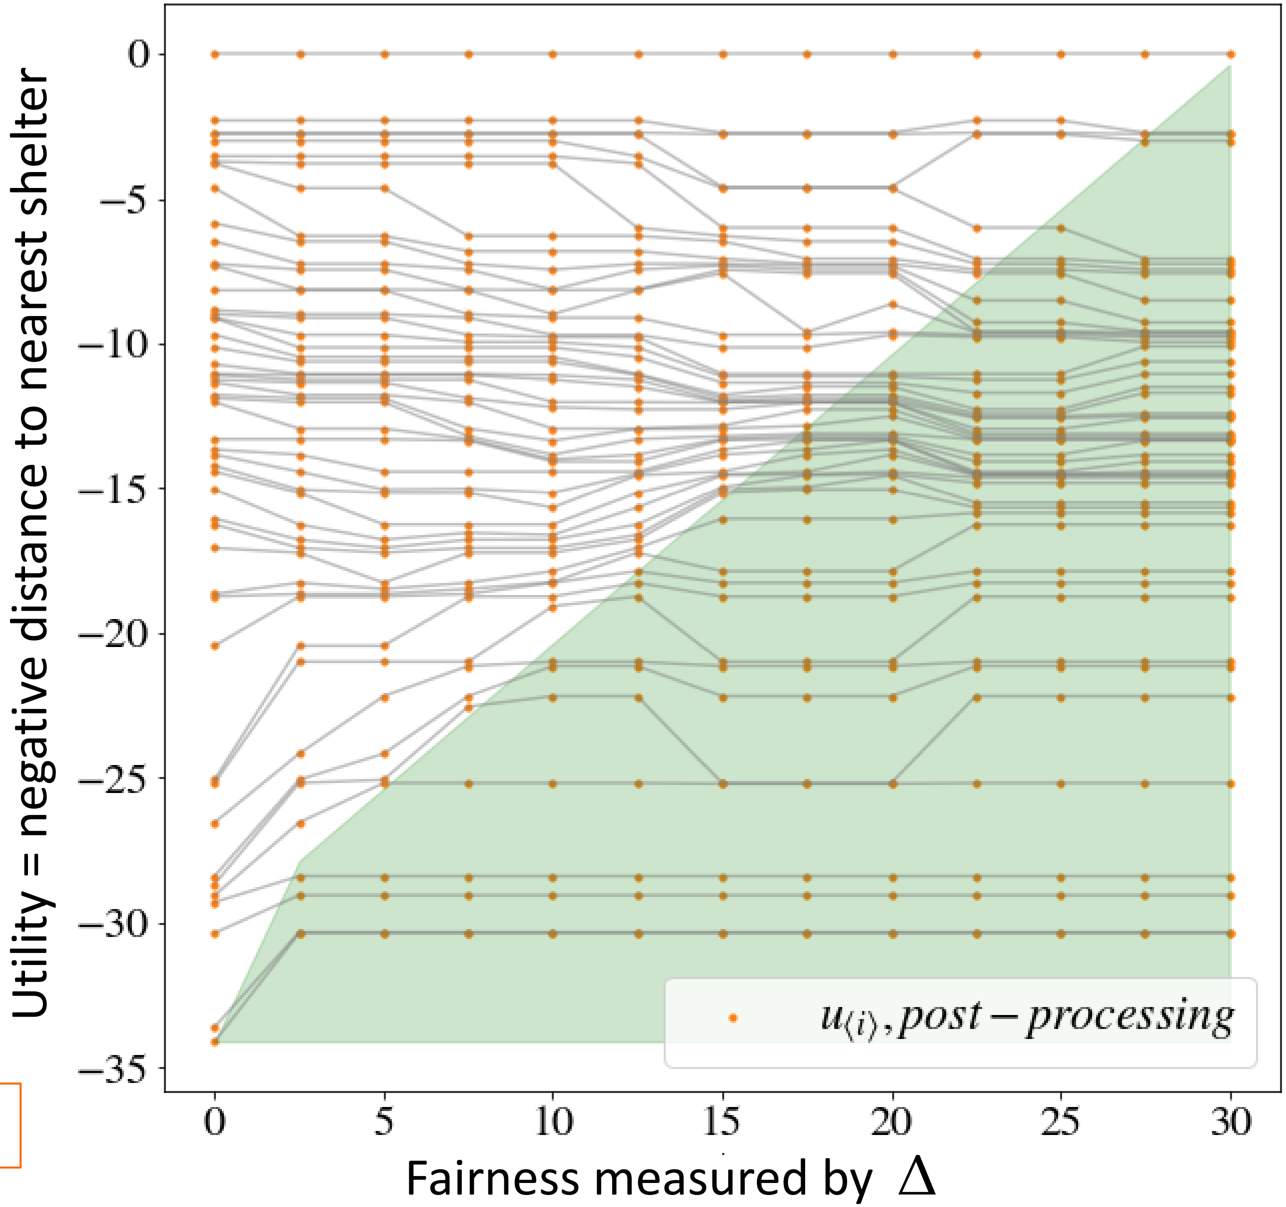

Chen & JH 2021

Research Question

How to combine **equity** threshold with **leximax**?

- Is the formulation linear?
- Does it require solving a sequence of LPs?
- What does the distribution look like?

Conclusions

Possible advantages of **threshold** SWFs over alpha fairness:

- **Interpretable** trade-off parameters
  - and practical **LP** and **MILP** formulations.
- **Equity** threshold SWF is a piecewise linear approximation of alpha fairness.
  - with **LP** rather than **nonlinear** formulation

Possible advantage of **alpha fairness**:

- Distribution varies **smoothly** with α

References

References and more details may be found in

V. Chen & J.N. Hooker, A guide to formulating fairness in an optimization model, *Annals of OR,* 2023.

Ö. Elçi, J.N. Hooker & P. Zhang, *The Structure of Fair Solutions*, Springer, 2025.