Equity through Social Welfare Optimization

John Hooker

Carnegie Mellon University

UCLA Instutute for Pure and Applied Mathematics

January 2024

Some results represent joint work with…

Violet (Xinying) Chen
Stevens Institute of Technology

Özgün Elçi
Amazon

H. Paul Williams
London School of Economics

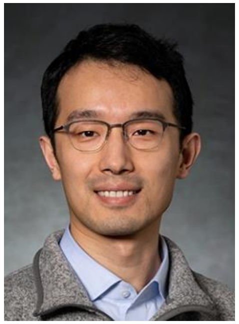

Peter Zhang
CMU

Modeling Fairness

- A growing interest in incorporating **fairness** into models
  - Health care resources.
  - Facility location (e.g., emergency services, infrastructure).
  - Telecommunications.
  - Traffic signal timing
  - Disaster recovery (e.g., power restoration)

Modeling Fairness

- Example: **Emergency facility location**
  - Locations in densely populated zone minimize **average response time,** but are unfair to those in outlying areas
  - Locations that minimize **worst-case response time** result in poor service for most of the population
- A more **equitable** solution
  - …would compromise between equity and efficiency.

Modeling Fairness

- Example: **Traffic signal timing**
  - **Throughput** is maximized by giving constant green light to the major street, red light to cross street.
  - Then motorists on the cross street **wait forever**.
- A more **equitable** solution would find a compromise.
  - For example, by using **proportional fairness** (Nash Bargaining solution), a special case of **alpha fairness.**

Modeling Fairness

- Similar example: **Telecommunications**
  - Must compromise between **maximizing total throughput** and **minimizing worst-case latency**
- An **early adopter** of fairness modeling.
  - Alpha fairness, Jain's index, QoE fairness, G's fairness index, Bossaert's fairness index
  - All but alpha fairness are pure inequality measures.

Modeling Fairness

- Example: **Disaster relief**
  - Power restoration can focus on **urban** areas first (**efficiency**).
  - This can leave rural areas without power for weeks/months.
  - This happened in Puerto Rico after Hurricane Maria (2017).

– A more **equitable** solution
  - …would give some priority to rural areas without overly sacrificing efficiency.

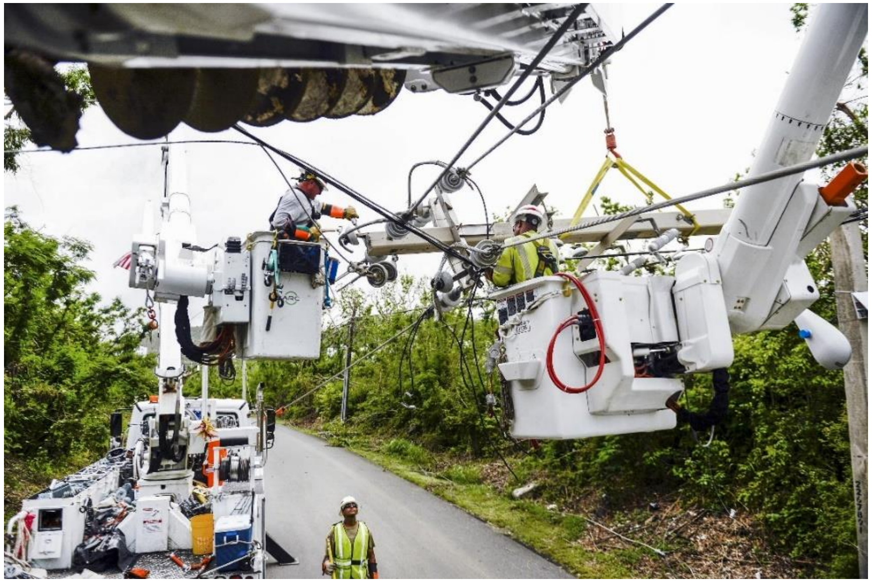

Modeling Fairness

- Optimization models are normally formulated to **maximize utility**.
    - where utility = wealth, health, negative cost, etc.
    - This can lead to **very unfair** resource distribution.

- For example...

Maximize Utility?

Utility maximizing
distribution
for 2 persons

Utility contours

$$u_1 + u_2$$

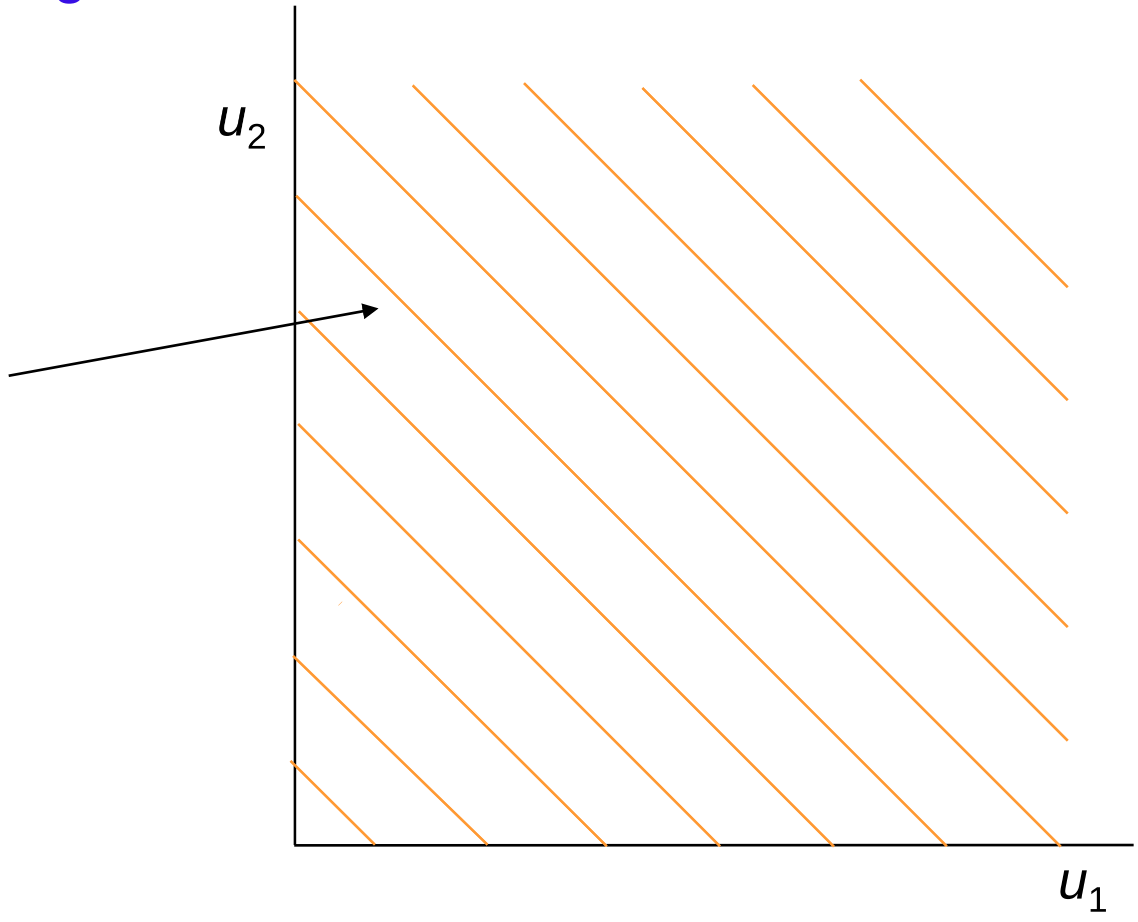

Maximize Utility?

Utility maximizing distribution for 2 persons, subject to budget constraint

Utility contours

$$u_1 + u_2$$

Feasible region

$$a_1 u_1 + a_2 u_2 \leq B$$

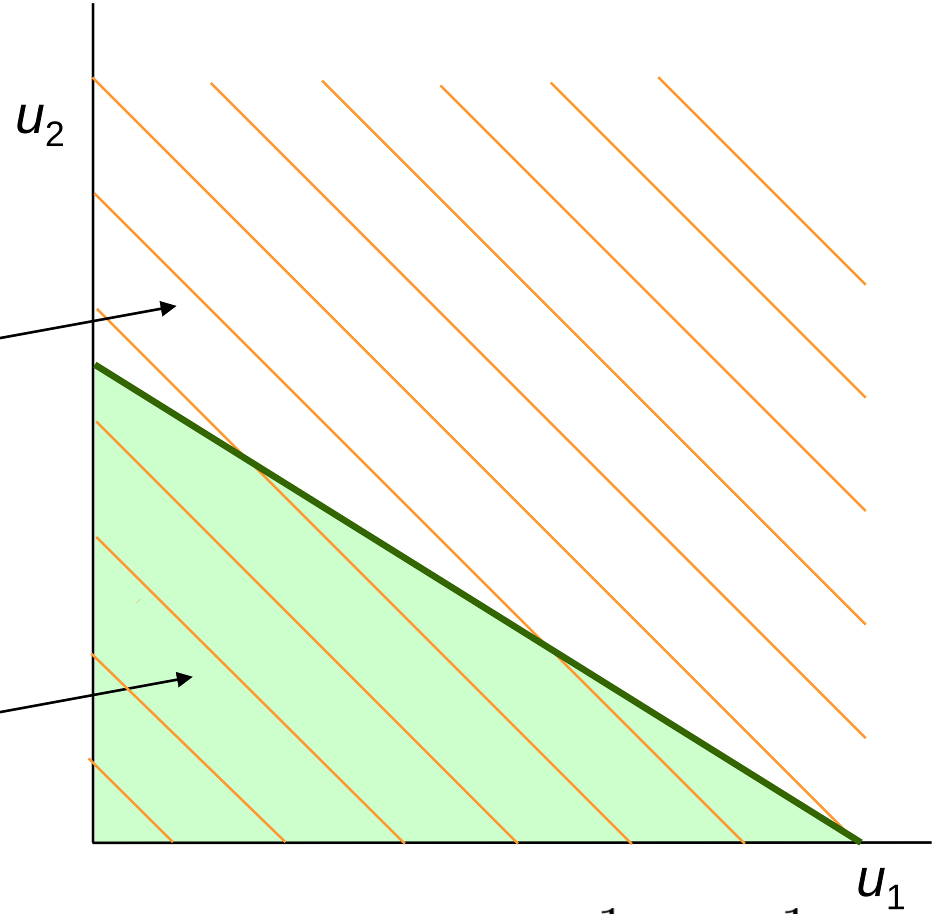

Person 1 has greater **conversion efficiency**: $\frac{1}{a_1} > \frac{1}{a_2}$

Maximize Utility?

Utility maximizing distribution for 2 persons, subject to budget constraint

Utility contours

$$u_1 + u_2$$

Feasible region

$$a_1 u_1 + a_2 u_2 \leq B$$

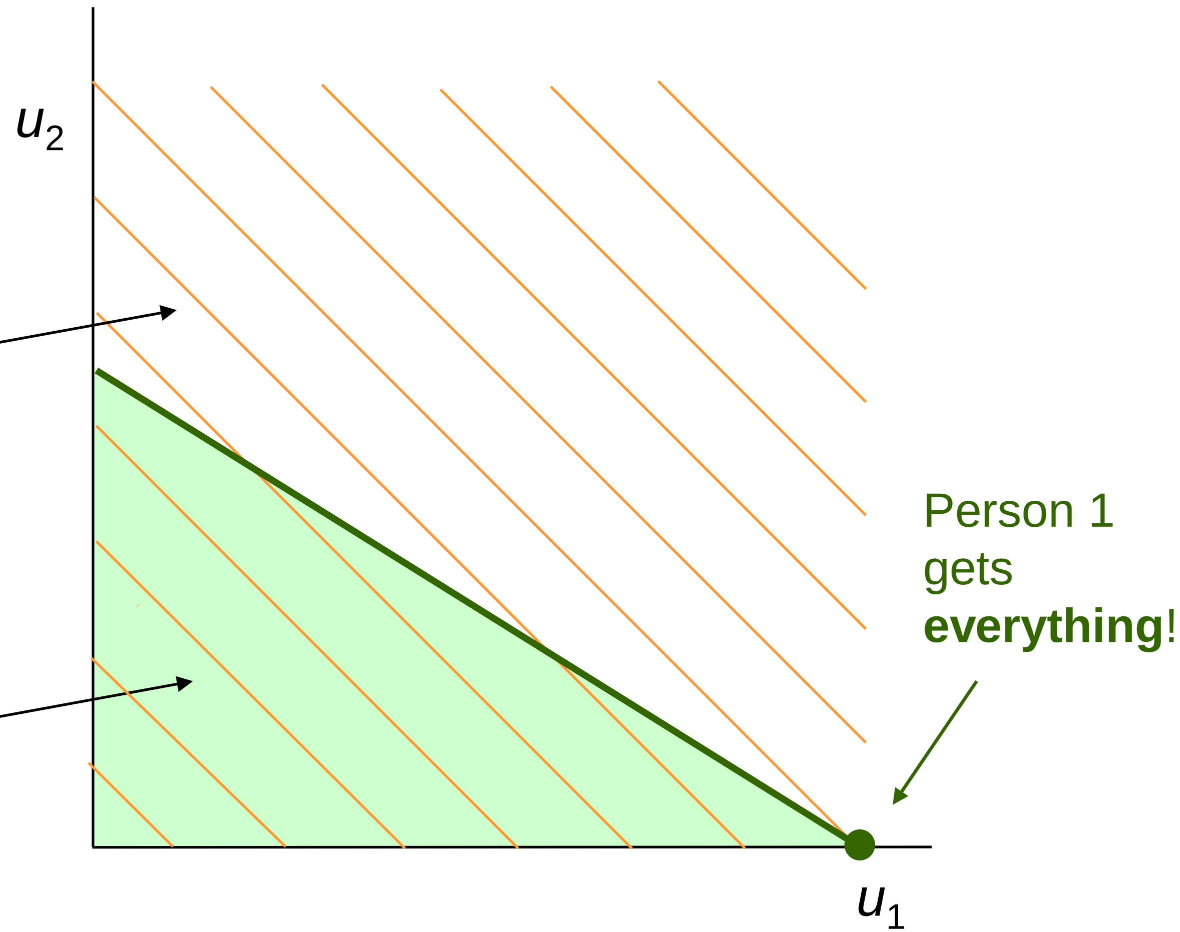

Person 2 makes **less efficient use of resources** (e.g., has a more serious disease)

Modeling Fairness

- There is **no one** concept of fairness.
  - The appropriate concept **depends on the context**.
- How to **choose the right one**?
- For each of several fairness models, we…
  - Describe the **optimal solutions** they deliver
  - Determine their implications for **hierarchical** distribution
  - Study how they incentivize **efficiency improvements** and **competition vs. cooperation**.
- We also take a brief excursion into **social choice theory**.

Modeling Fairness

- We focus on fairness models that **balance equity and efficiency** in some principled way.
  - Why not use an **inequality bound?**
    - This provides no guidance for the equity-efficiency **trade-off**.
  - Why not use a **convex combination**?
    - It is unclear how to interpret the **multipliers** assigned to equity and efficiency (which are typically expressed in incommensurable units).

Generic Model

- We formulate each fairness criterion as a **social welfare function (SWF)**.

Individual utilities

$$W(\boldsymbol{u}) = W(u_1, \ldots, u_n)$$

  - Measures desirability of the **magnitude and distribution of utilities** across individuals.
  - The **SWF** becomes the **objective function** of the optimization model.

Generic Model

The social welfare optimization problem

$$\max_{\boldsymbol{u},\boldsymbol{x}} \left\{ W(\boldsymbol{u}) \;\middle|\; \boldsymbol{u} = U(\boldsymbol{x}),\ \boldsymbol{x} \in S_{\boldsymbol{x}} \right\}$$

Vector of utilities enjoyed by individuals (or groups) → $\boldsymbol{u}$

Vector of resource allocations to individuals (or groups) → $\boldsymbol{x}$

Social welfare function → W

Vector of utility functions $U_1, \ldots, U_n$ → U

Problem constraints → $S_{\boldsymbol{x}}$

Generic Model

We state structural results for a linearly constrained model

Reciprocals of conversion efficiencies

Individual utilities

$$\max_{\boldsymbol{u}} \left\{ W(\boldsymbol{u}) \;\middle|\; \sum_i a_i u_i \leq B, \;\; \mathbf{0} \leq \boldsymbol{u} \leq \boldsymbol{d} \right\}$$

Social welfare function

Budget constraint

Utility bounds (upper bounds **optional**)

Conversion efficiency of individual *i* = 1/a_i

The **linear** budget constraint specifies conversion efficiencies while allowing **fairness properties** to be indicated **transparently** in the SWF.

References

- References and more details may be found in

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

Ö. Elçi, J.N. Hooker & P. Zhang, Structural properties of fair solutions, submitted 2023.

Fairness for the disadvantaged

Criterion	*Linear?*	*Contin?*
Maximin (Rawlsian)	yes	yes
Leximax (lexicographic)	yes	yes

Combining efficiency & fairness
Classical methods

Criterion	*Linear?*	*Contin?*
Alpha fairness	yes	yes
Proportional fairness (Nash bargaining)	yes	yes
Kalai-Smorodinsky bargaining	no	yes

Linear = fairness model introduces only **linear** expressions
Contin. = fairness model introduces only **continuous** variables

Combining efficiency & fairness
Threshold methods

Criterion	*Linear?*	*Contin?*
Utility + maximin – Utility threshold	yes	no
Utility + maximin – Equity threshold	yes	yes
Utility + leximax – Predefined priorities	yes	no
Utility + leximax – No predefined priorities	yes	no

Linear = fairness model introduces only **linear** expressions
Contin. = fairness model introduces only **continuous** variables

Hierarchical Distribution

Two-level hierarchy

- **National authority** allocates resources to **regions**.
- **Each region** combines these resources with its own resources and allocates to **subregions**.

Regional decomposability

- **Each region's** allocation to subregions is **the same** as in a **national solution** that uses the **same SWF**.
- Surprisingly, some SWFs are **not regionally decomposable**.

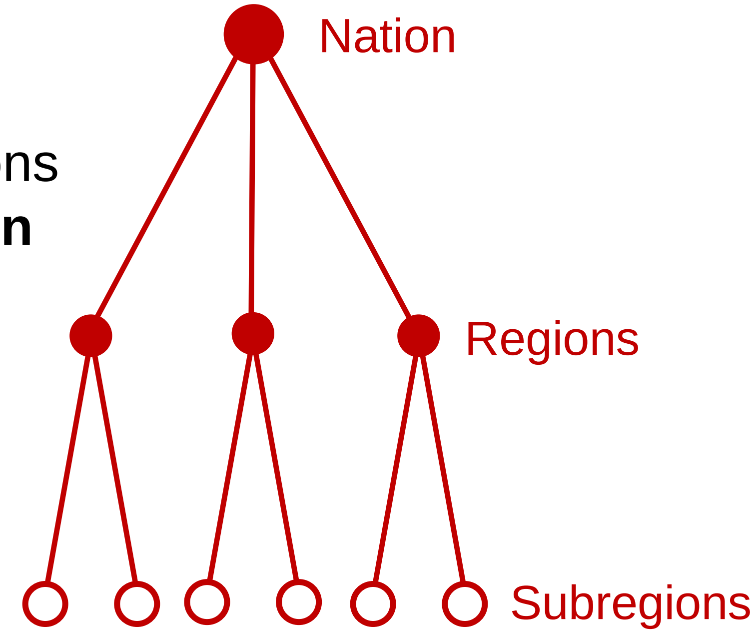

Hierarchical Distribution

Sufficient condition for regional decomposability

SWF $W(\boldsymbol{u})$ is *monotonically separable* when for any partition $\boldsymbol{u} = (\boldsymbol{u}^1, \boldsymbol{u}^2)$, $W(\bar{\boldsymbol{u}}^1) \geq W(\boldsymbol{u}^1)$ and $W(\bar{\boldsymbol{u}}^2) \geq W(\boldsymbol{u}^2)$ imply $W(\bar{\boldsymbol{u}}) \geq W(\boldsymbol{u})$.

In particular, a separable SWF is monotonically separable.

Theorem.

A monotonically separable SWF is regionally decomposable.

Incentives and Sharing

My **incentive rate** =

$$\frac{\text{\% increase in my optimal utility allotment}}{\text{\% increase in my conversion efficiency}}$$

A **positive** incentive rate indicates a reward for **improving** efficiency.

My **cross-subsidy rate** with respect to another individual =

$$\frac{\text{\% increase in the other individual's optimal utility allotment}}{\text{\% increase in my conversion efficiency}}$$

Positive cross-subsidy rates indicate **cooperation**.

Negative cross-subsidy rates indicate **competition**.

Utilitarian SWF

Maximize total utility: $$W(\boldsymbol{u}) = \sum_i u_i$$

Optimal solution subject to budget constraint:

- Most efficient person gets everything.

Regionally decomposable?

- Separable SWF → **yes**.

Incentive rate?

- **1** for most efficient person, **0** for others.

Cross-subsidy rates?

- All **zero**

Maximin

Maximize minimum utility: $W(\boldsymbol{u}) = \min_i \{u_i\}$

Suggested by social contract argument for **Difference Principle** of John Rawls, which applies only to design of social institutions and distribution of "primary goods."

Optimal solution subject to budget constraint:

- Everyone gets **equal** utility.

Maximin

2-person example with budget constraint

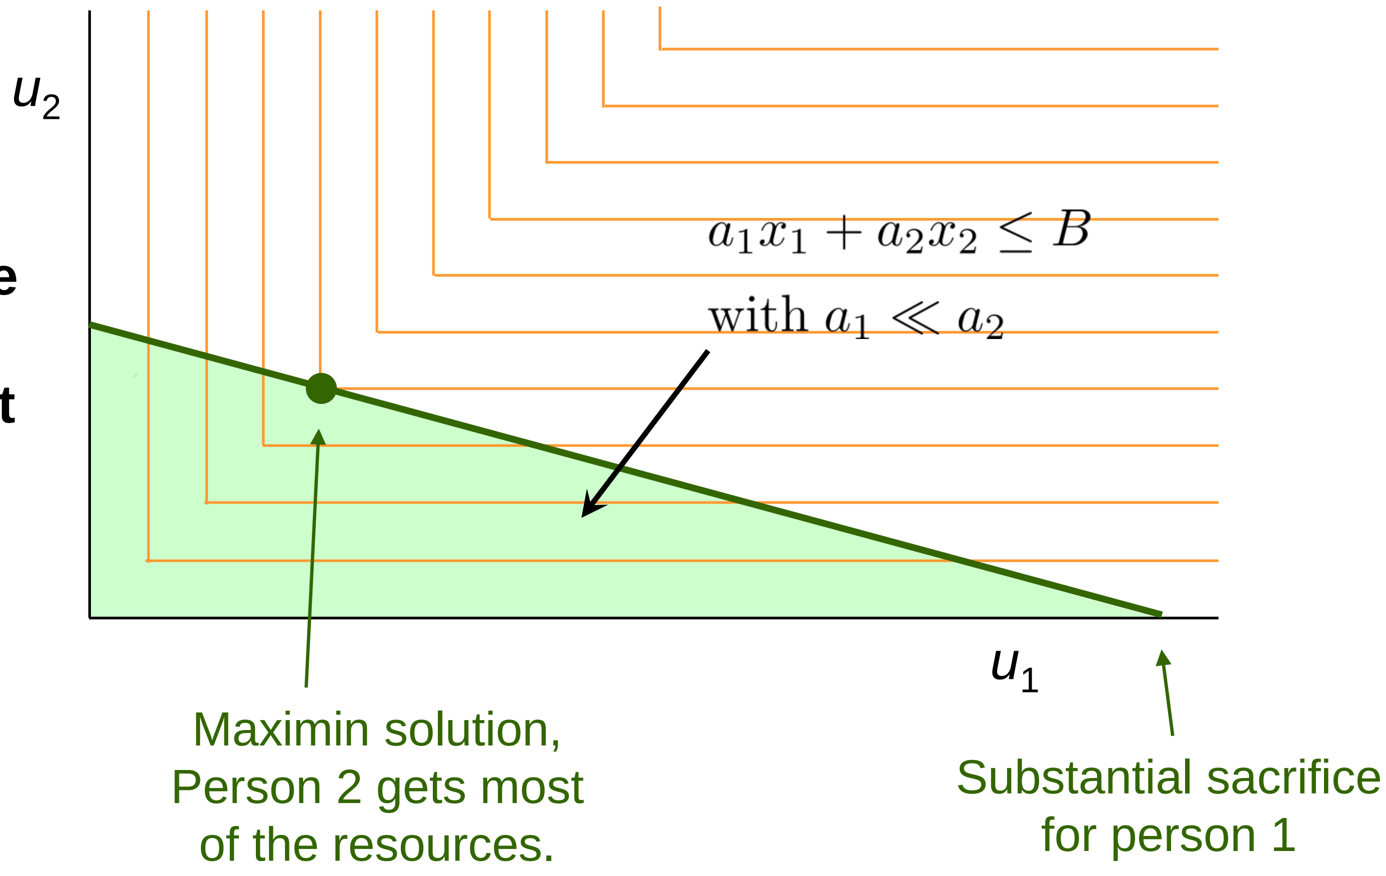

In a medical context, patient 1 is reduced to same level of suffering as seriously ill patient 2.

Maximin

Maximize minimum utility: $W(\boldsymbol{u}) = \min_i \{u_i\}$

Suggested by social contract argument for **Difference Principle** of John Rawls, which applies only to design of social institutions and distribution of "primary goods."

Optimal solution subject to budget constraint:

- Everyone gets **equal** utility.

Optimal solution subject to resource bounds:

- Can **waste** most of the available resources.

Fairness for the Disadvantaged

Maximin

Example with resource bounds

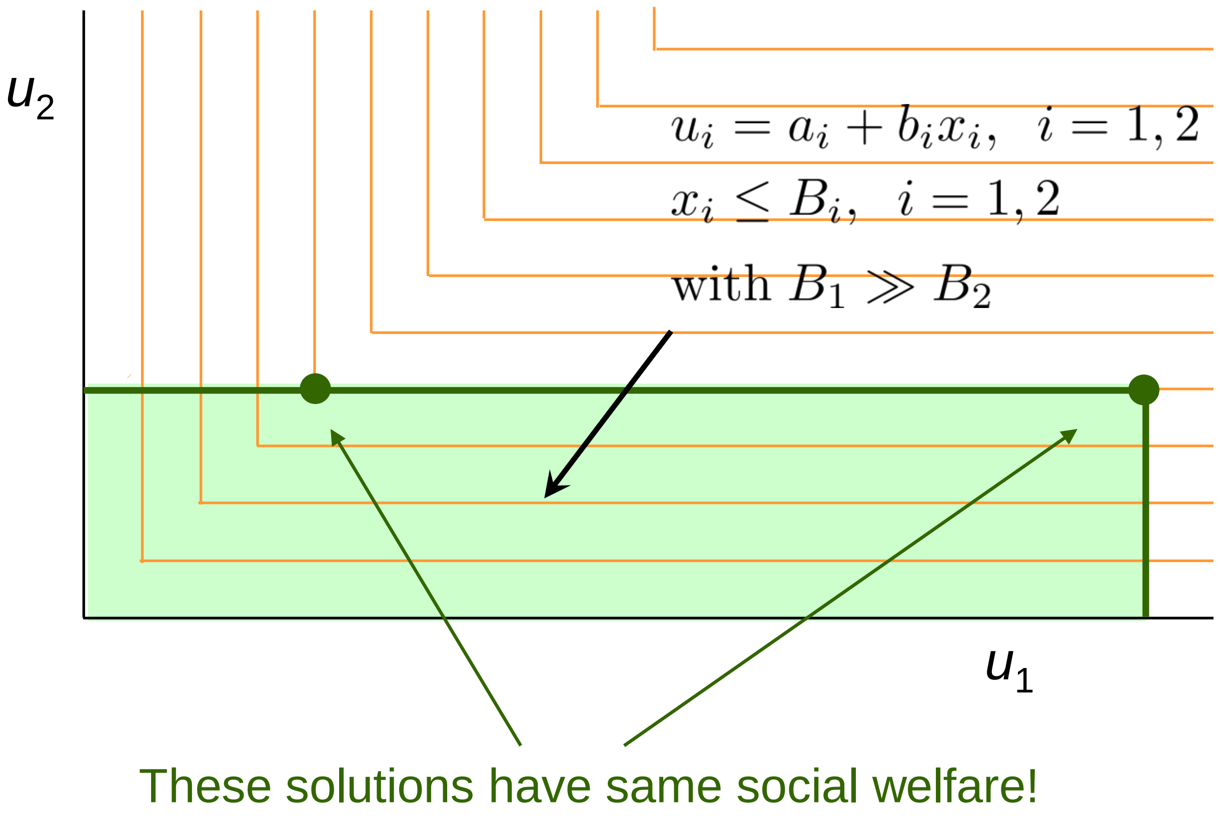

Fairness for the Disadvantaged

Maximin

Example with resource bounds

Remedy**:** use **leximax** solution

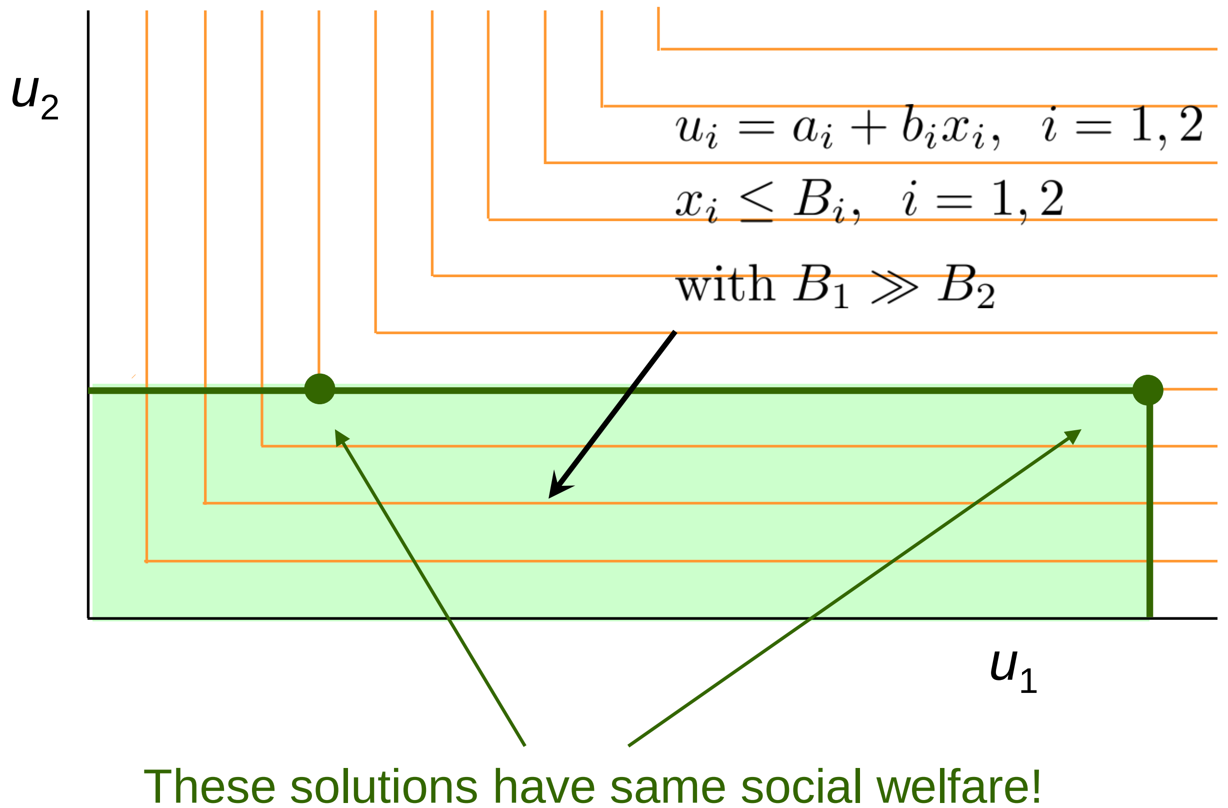

Maximin

Maximize minimum utility: $W(\boldsymbol{u}) = \min_i \{u_i\}$

Regionally decomposable

- Monotonically separable SWF

Positive incentive rate for person *i* = $\dfrac{a_i}{\sum_j a_j}$

- Less efficient parties have greater incentive to improve.

Positive cross-subsidy to all others: $\dfrac{a_i}{\sum_j a_j}$

- Everyone benefits equally from person *i*'s improvement.

Leximax

Maximize smallest utility, then 2nd smallest, etc.

Optimal solution subject to budget constraint:

- Everyone gets **equal** utility.

Optimal solution subject to budget constraint and bounds:

- **No waste** of resources.

Regionally decomposable

- using generalized definition of decomposability

Social Choice Theory

- The economics literature derives social welfare functions from **axioms of rational choice**.
- The social welfare function depends on degree of **interpersonal comparability** of utilities.
- Arrow's impossibility theorem was the first result, but there are many others.

Social Choice Theory

Axioms

Anonymity (symmetry)

Social preferences are the same if indices of u_is are permuted.

Strict pareto

If $\boldsymbol{u} > \boldsymbol{u}'$, then $\boldsymbol{u}$ is preferred to $\boldsymbol{u}'$.

Independence

The preference of $\boldsymbol{u}$ over $\boldsymbol{u}'$ depends only on $\boldsymbol{u}$ and $\boldsymbol{u}'$ and not on what other utility vectors are possible.

Separability

Individuals i for which $u_i = u_i'$ do not affect the relative ranking of $\boldsymbol{u}$ and $\boldsymbol{u}'$.

Social Choice Theory

Interpersonal comparability

- The properties of social welfare functions that satisfy the axioms depend on the degree to which utilities can be **compared** across individuals.

Invariance transformations

- These are transformations of utility vectors that indicate the degree of interpersonal comparability.
- Applying an invariance transformation to utility vectors does not change the **ranking** of distributions.

An invariance transformation has the form $\boldsymbol{\phi} = (\phi_1, \ldots, \phi_n)$, where ϕ_i is a transformation of individual utility i.

Social Choice Theory

Unit comparability.

- Invariance transformation has the form $\phi_i(u_i) = \beta u_i + \gamma_i$
- So, it is possible to compare utility **differences** across individuals:

$$u_i' - u_i > u_j' - u_j \text{ if and only if } \phi_i(u_i') - \phi_i(u_i) > \phi_j(u_j') - \phi_j(u_j)$$

Theorem. Given anonymity, strict pareto, independence axioms, and **unit comparability**, the social welfare criterion must be **utilitarian**.

$$W(\boldsymbol{u}) = \sum_i u_i$$

Social Choice Theory

Level comparability.

- Invariance transformation has the form

$$\boldsymbol{\phi}(\boldsymbol{u}) = \big(\phi_0(u_1), \ldots, \phi_0(u_n)\big)$$
where ϕ_0 is strictly increasing.

- So, it is possible to compare utility **levels** across individuals.

$$u_i > u_j \text{ if and only if } \phi_i(u_i) > \phi_j(u_j)$$

Theorem. Given anonymity, strict pareto, independence, separability axioms, and **level comparability**, the social welfare criterion must be **maximin or minimax**.

Social Choice Theory

Problem with the utilitarian proof.

- The proof assumes that utilities have **no more** than unit comparability.
- This immediately rules out a maximin criterion, since identifying the minimum utility presupposes that utility **levels** can be compared.

Problem with the maximin proof.

- The proof assumes that utilities have **no more** than level comparability.
- This immediately rules out criteria that consider the spread of utilities.
- So, it rules out all the criteria we consider after maximin.

Alpha Fairness

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

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

Mo & Walrand 2000; Verloop, Ayesta & Borst 2010

Solution subject to budget constraint:

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

- **Utilitarian** when $\alpha = 0$, **maximin** when $\alpha \to \infty$
- **Egalitarian** distribution can have same social welfare as **arbitrarily extreme inequality**.
- Can be **derived** from certain axioms.

Lan & Chiang 2011

Alpha Fairness

Example:

Maximum alpha fairness subject to budget constraint

$$u_1 + 2u_2 + \cdots + 8u_8 \leq 100$$

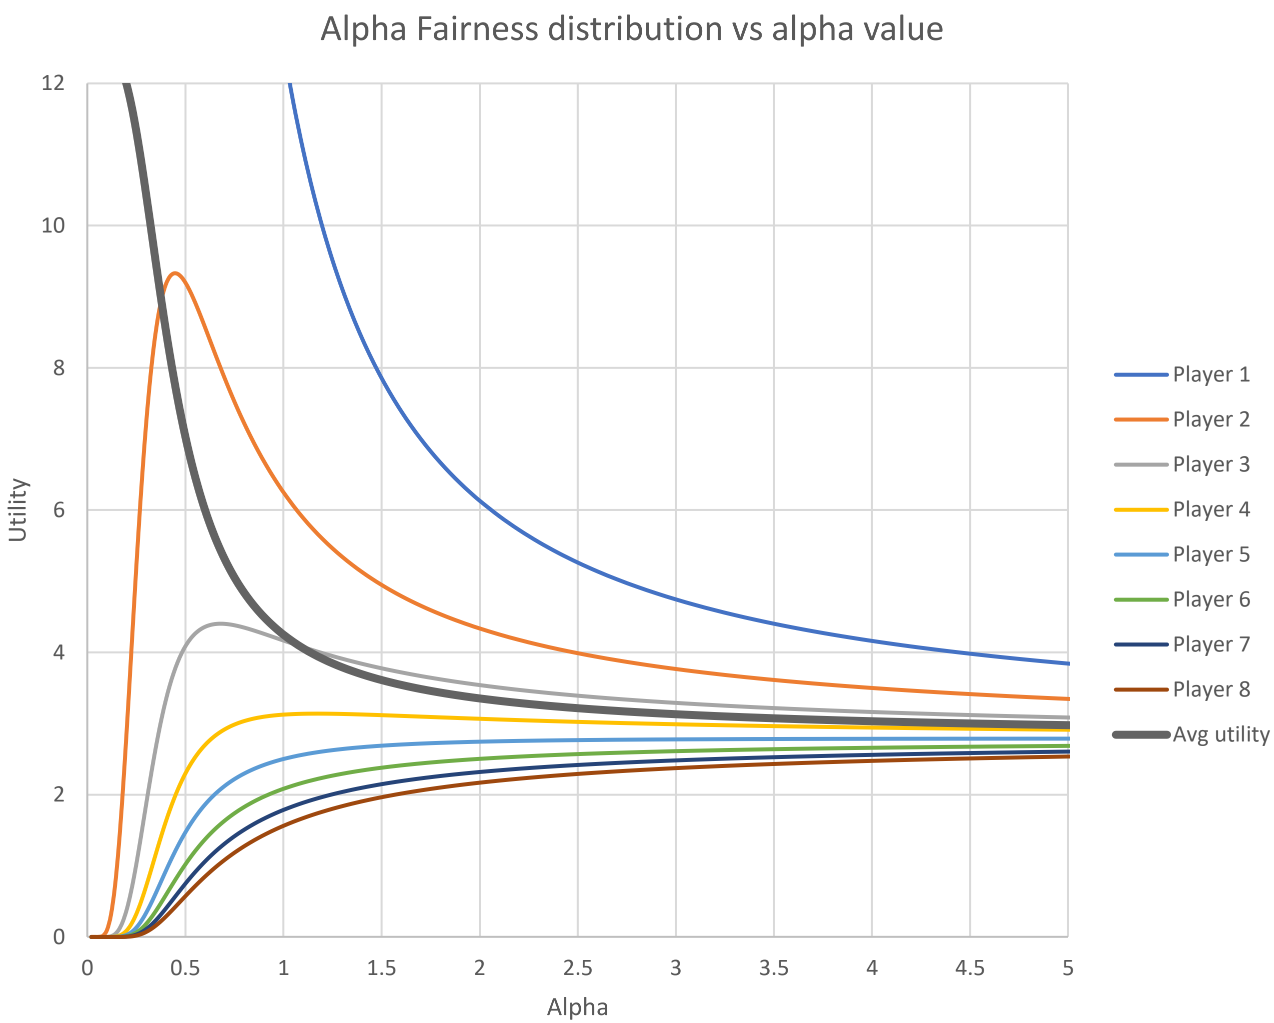

Unclear how to choose α in practice

Alpha Fairness

Regionally decomposable

- Separable SWF → **yes**.

Positive incentive rate for person *i* = $\frac{1}{\alpha} + \left(1 - \frac{1}{\alpha}\right) \frac{a_i^{1-1/\alpha}}{\sum_j a_j^{1-1/\alpha}}$

- Incentive to improve **increases** with current conversion efficiency when $\alpha < 1$, **decreases** when $\alpha > 1$.

Cross-subsidy to others = $\left(1 - \frac{1}{\alpha}\right) \frac{a_i^{1-1/\alpha}}{\sum_j a_j^{1-1/\alpha}}$

- **Negative** when $\alpha < 1$ (**competition**). Efficiency improvements transfer utility **from** other persons
- **Positive** when $\alpha > 1$ (**sharing**), improvements transfer utility **to** others

Proportional Fairness

Nash 1950

Special case of alpha fairness (α = 1)

- Also known as **Nash bargaining solution**, in which case bargaining starts with a default distribution ***d***.

$$W(\boldsymbol{u}) = \sum_i \log(u_i - d_i) \quad \text{or} \quad W(\boldsymbol{u}) = \prod_i (u_i - d_i)$$

Solution subject to budget constraint

- Utility allotted **in proportion to conversion efficiency**.
- Can be **derived** from axiomatic and bargaining arguments.
- Used in engineering applications (telecom, traffic signaling).

Incentive rate = 1

Cross-subsidies = 0

Proportional Fairness

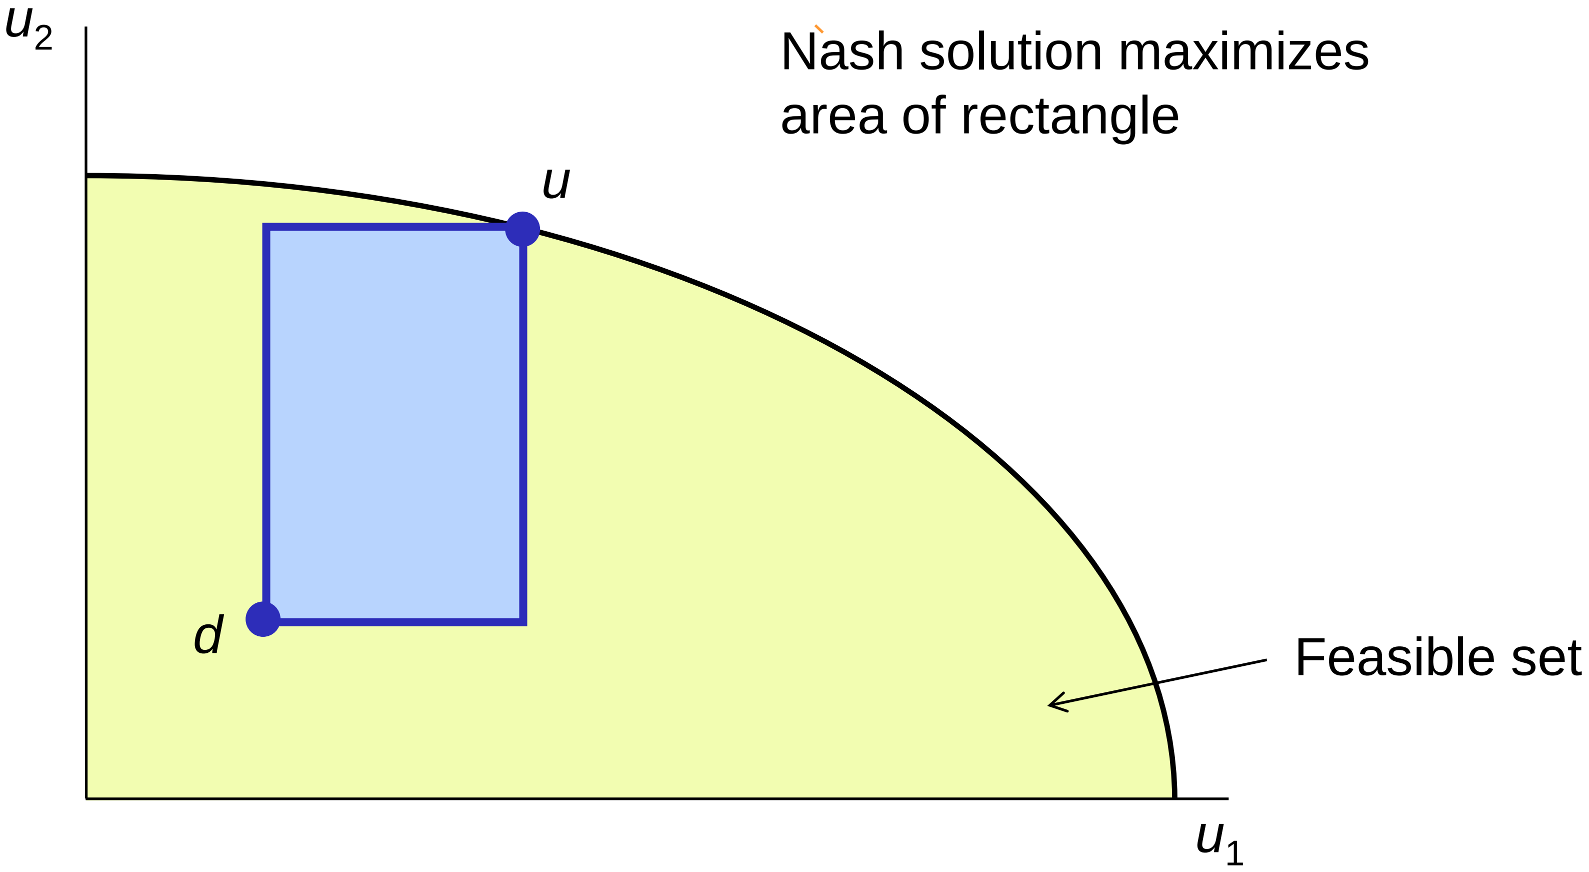

Proportional Fairness

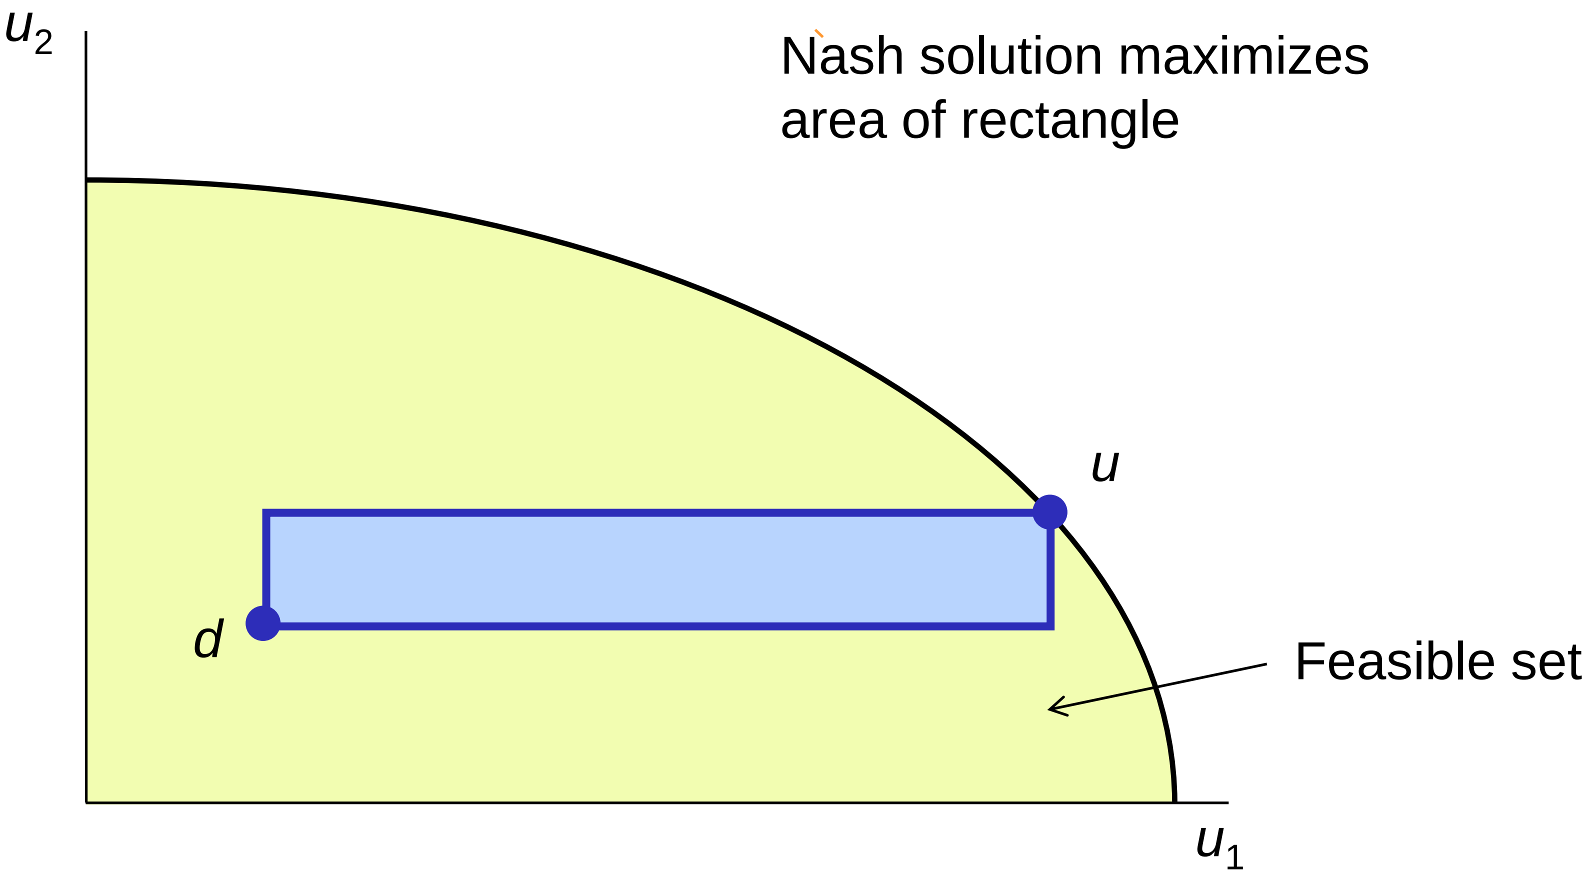

Proportional Fairness

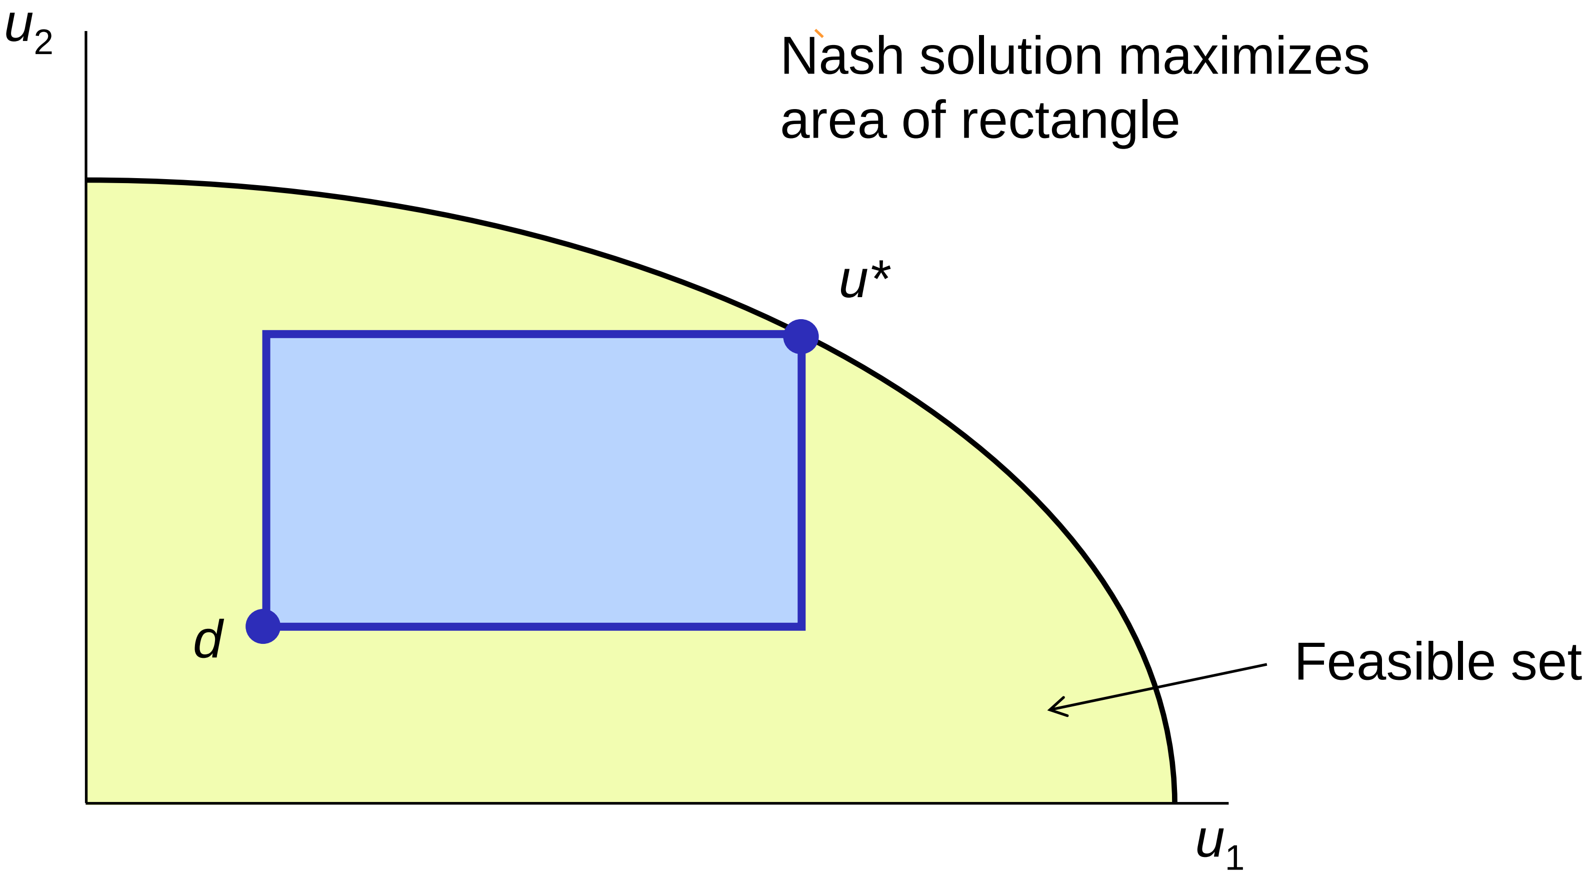

Back to Social Choice Theory

Axiomatic derivation of proportional fairness

From Nash's article, based on:

- **Anonymity, Pareto** and **independence** axioms
- **Scale invariance:** invariance transformation $\phi_i(u_i) = \beta_i u_i$

Nash 1950

Back to Social Choice Theory

Axiomatic derivation of proportional fairness

From Nash’s article, based on:

- **Anonymity, Pareto** and **independence** axioms
- **Scale invariance:** invariance transformation $\phi_i(u_i) = \beta_i u_i$

Nash 1950

Possible problem

Invariance under individual rescaling is better suited to negotiation procedures than assessing just distributions.

Back to Social Choice Theory

Bargaining justifications

“Rational” negotiation converges to the Nash bargaining solution. Assumes an initial utility distribution to which parties return if negotiation fails.

- Finite convergence (assuming a minimum distance between offers), based on a bargaining procedure of Zeuthen.

Harsanyi 1977 | Zeuthen 1930

- Asymptotic convergence based on equilibrium modeling.

Rubinstein 1982 | Binmore, Rubinstein, Wolinsky 1986

Back to Social Choice Theory

Bargaining justifications

"Rational" negotiation converges to the Nash bargaining solution. Assumes an initial utility distribution to which parties return if negotiation fails.

- Finite convergence (assuming a minimum distance between offers), based on a bargaining procedure of Zeuthen.

Harsanyi 1977 | Zeuthen 1930

- Asymptotic convergence based on equilibrium modeling.

Rubinstein 1982 | Binmore, Rubinstein, Wolinsky 1986

Possible problem

Not clear that rational **negotiation** leads to **justice**.

Back to Social Choice Theory

Axiomatic derivation of alpha fairness

- Certain axioms lead to a **family** of SWFs containing **alpha fairness**, along with logarithmic functions (including Theil & Atkinson indices).
- Key to the proof is an **axiom of partition**:

Lan and Chiang 2011

There exists a mean function h such that for any partition $(\boldsymbol{u}_1, \boldsymbol{u}_2)$ of $\boldsymbol{u}$ and any two distributions $\boldsymbol{u}$ and $\boldsymbol{u}'$,

$$\frac{W(t\boldsymbol{u})}{W(t\boldsymbol{u}')} = h\left(\frac{W(\boldsymbol{u}_1)}{W(\boldsymbol{u}_1')}, \frac{W(\boldsymbol{u}_2)}{W(\boldsymbol{u}_2')}\right)$$

where $t > 0$ is an arbitrary scalar. This implies that h must be a geometric or power mean.

Back to Social Choice Theory

Axiomatic derivation of alpha fairness

- Certain axioms lead to a **family** of SWFs containing **alpha fairness**, along with logarithmic functions (including Theil & Atkinson indices).
- Key to the proof is an **axiom of partition**:

Lan and Chiang 2011

There exists a mean function h such that for any partition $(\boldsymbol{u}_1, \boldsymbol{u}_2)$ of $\boldsymbol{u}$ and any two distributions $\boldsymbol{u}$ and $\boldsymbol{u}'$,

$$\frac{W(t\boldsymbol{u})}{W(t\boldsymbol{u}')} = h\left(\frac{W(\boldsymbol{u}_1)}{W(\boldsymbol{u}_1')}, \frac{W(\boldsymbol{u}_2)}{W(\boldsymbol{u}_2')}\right)$$

where $t > 0$ is an arbitrary scalar. This implies that h must be a geometric or power mean.

Possible problem

It is hard to interpret the axiom of partition.

Example of Alpha Fairness

Investment in electric generating capacity and transmission - Liberia

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

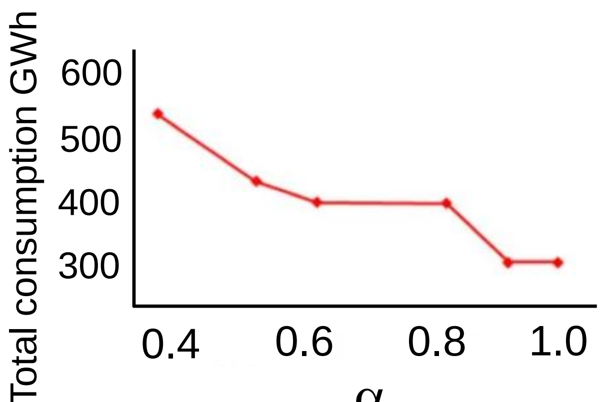

Sackey, Nock, Cao, Armanios, Davis 2023

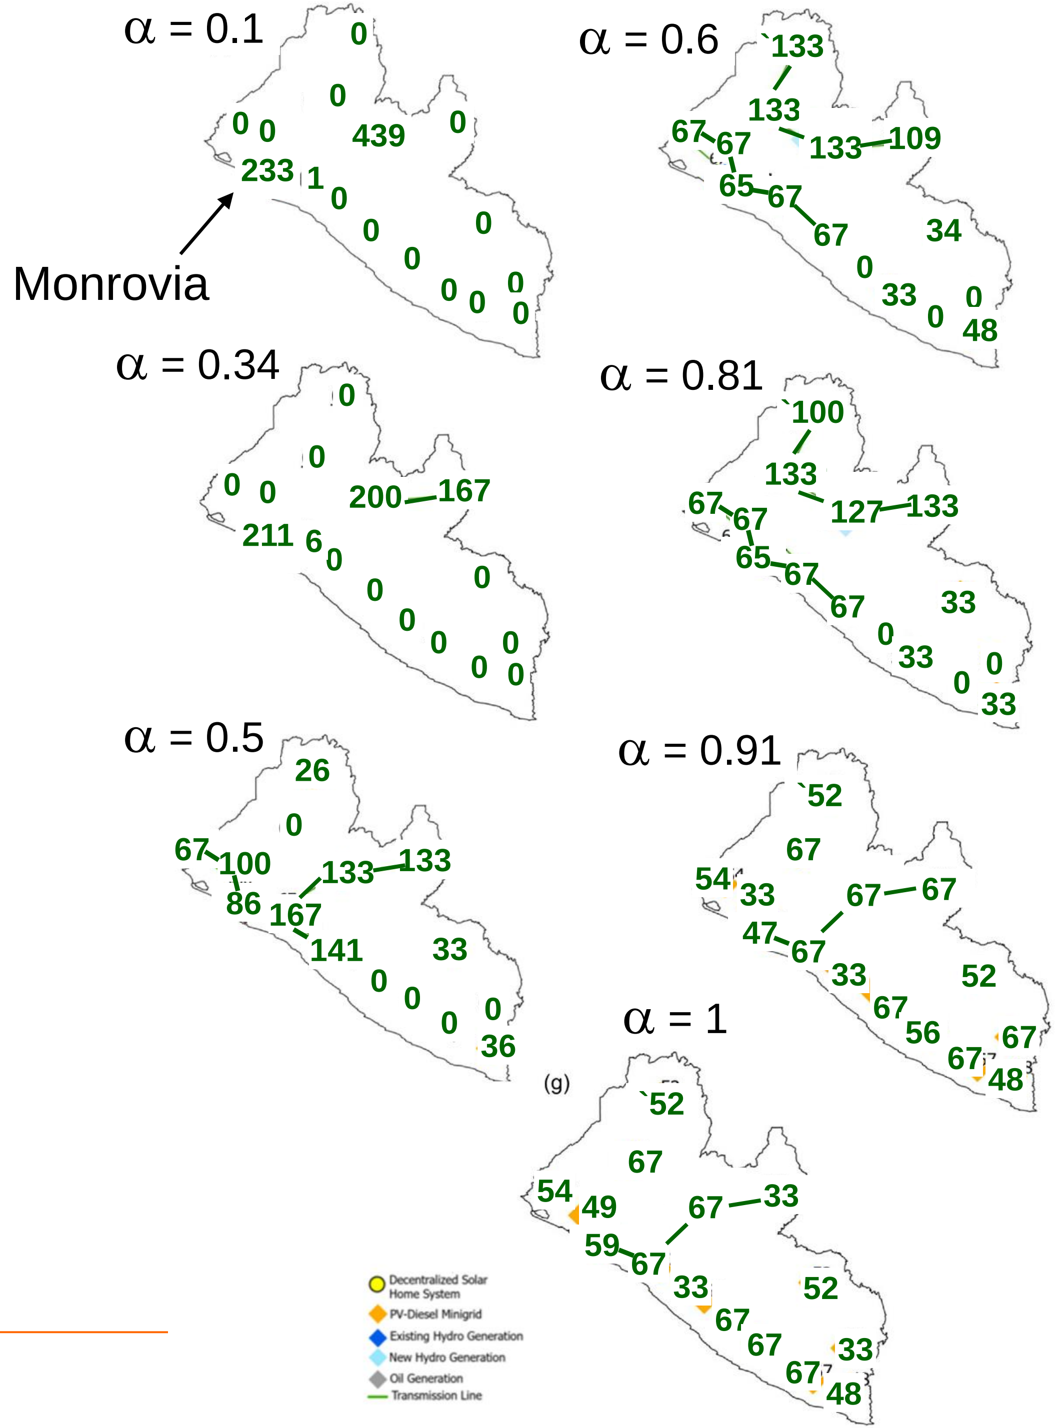

Example of Alpha Fairness

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 showing 9 hypothetical maps to U.S. engineering graduate students

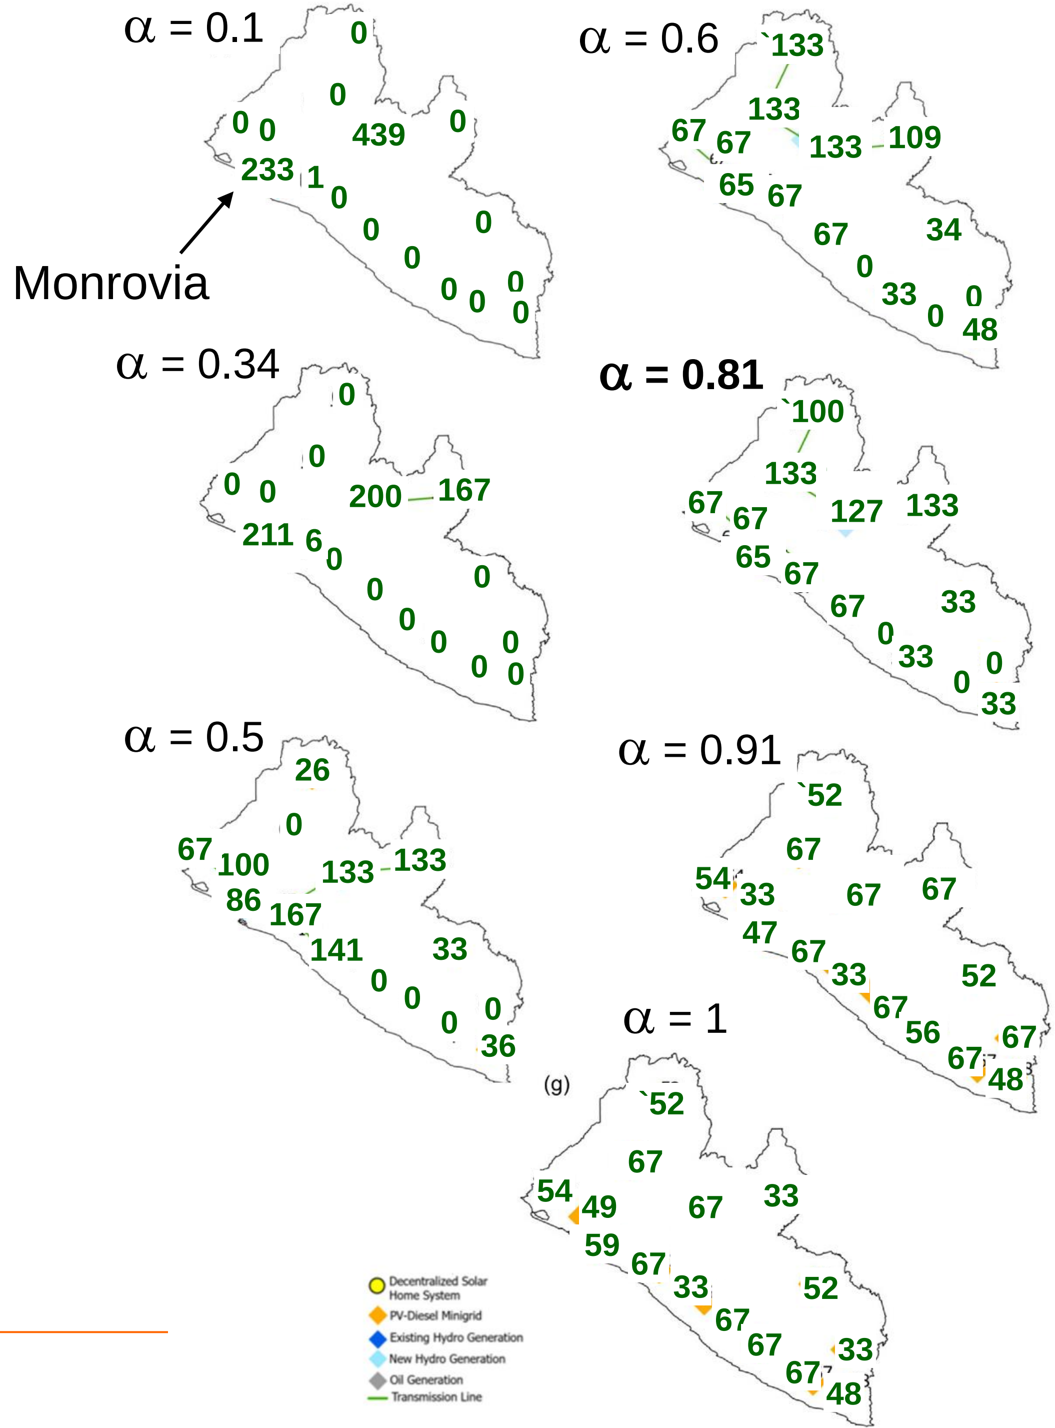

Sackey, Nock, Cao, Armanios, Davis 2023

Example of Alpha Fairness

Investment in electric generating capacity and transmission

- Application in Liberia
- Pure efficiency objective neglects the hinterland.
- Emphasis on fairness neglects urban dwellers.
- Elicited value of α = 0.81
  - Based on showing 9 hypothetical maps to U.S. engineering graduate students

Fairness preference frequencies

Based on pairwise comparisons

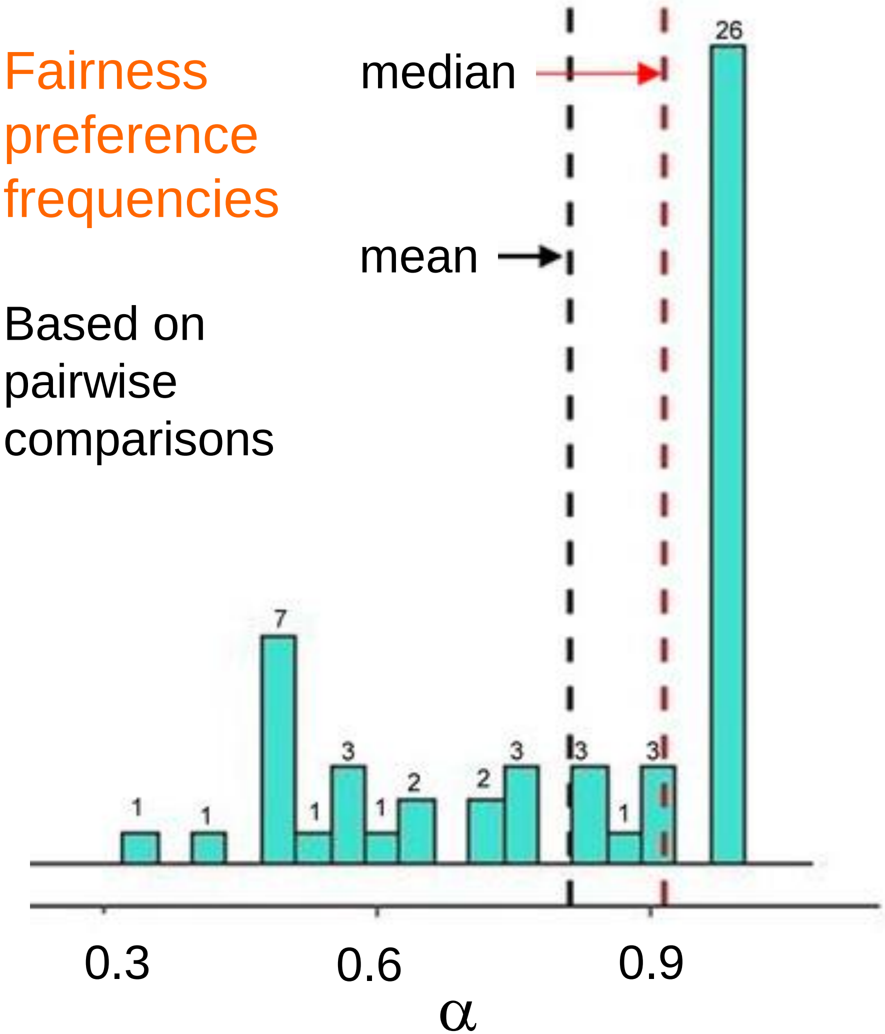

Sackey, Nock, Cao, Armanios, Davis 2023

Kalai-Smorodinsky Bargaining

- Begins with a critique of the Nash bargaining solution.

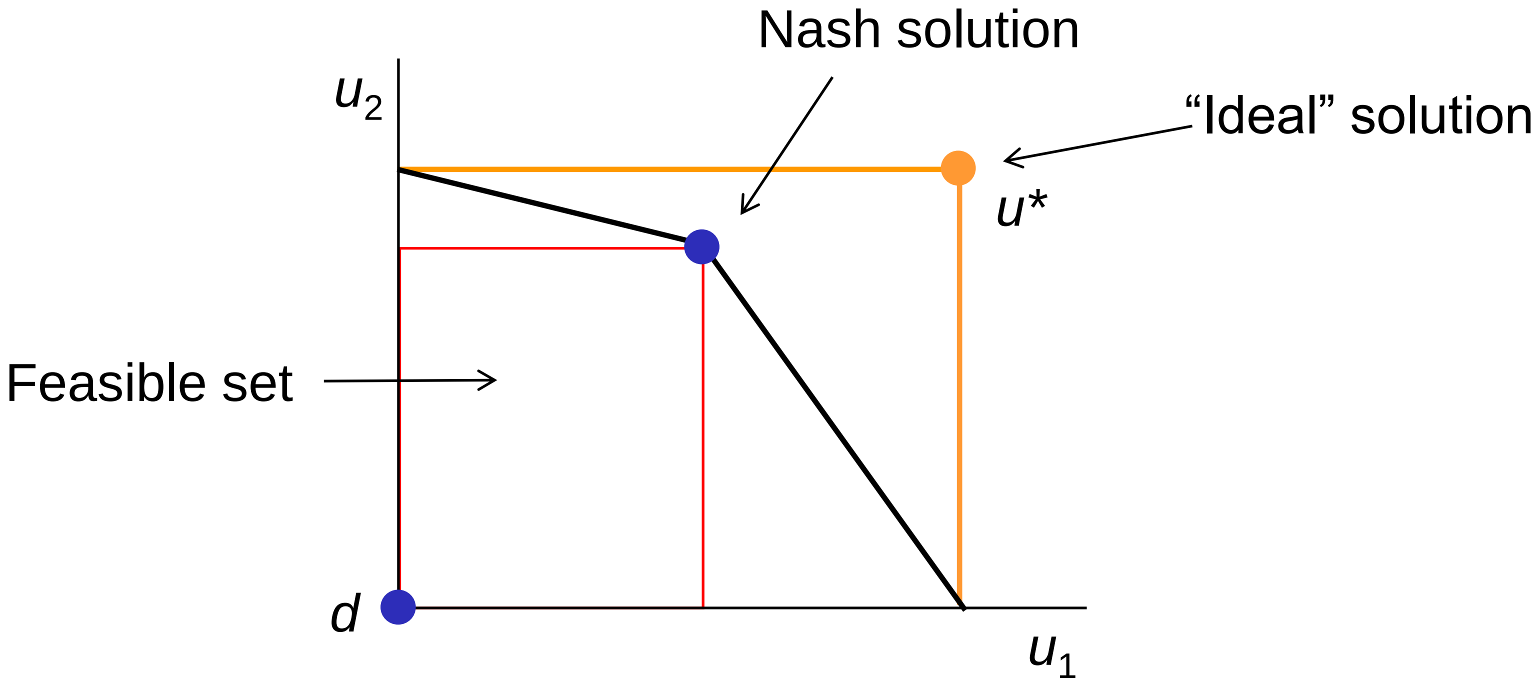

Kalai-Smorodinsky Bargaining

- Begins with a critique of the Nash bargaining solution.
- The new Nash solution is **worse** for player 2 even though the feasible set is **larger**.

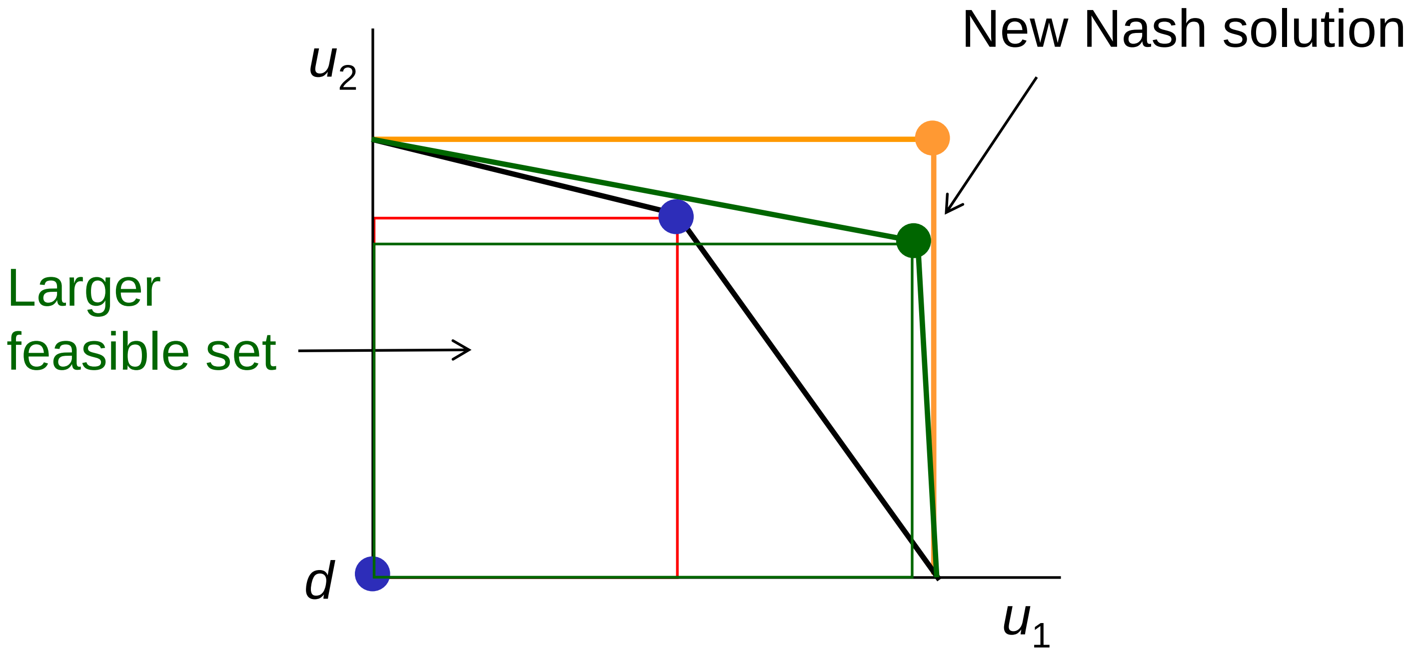

Kalai-Smorodinsky Bargaining

- **Proposal**: Bargaining solution is pareto optimal point on line from *d* to ideal solution.

Kalai & Smorodinksy 1975

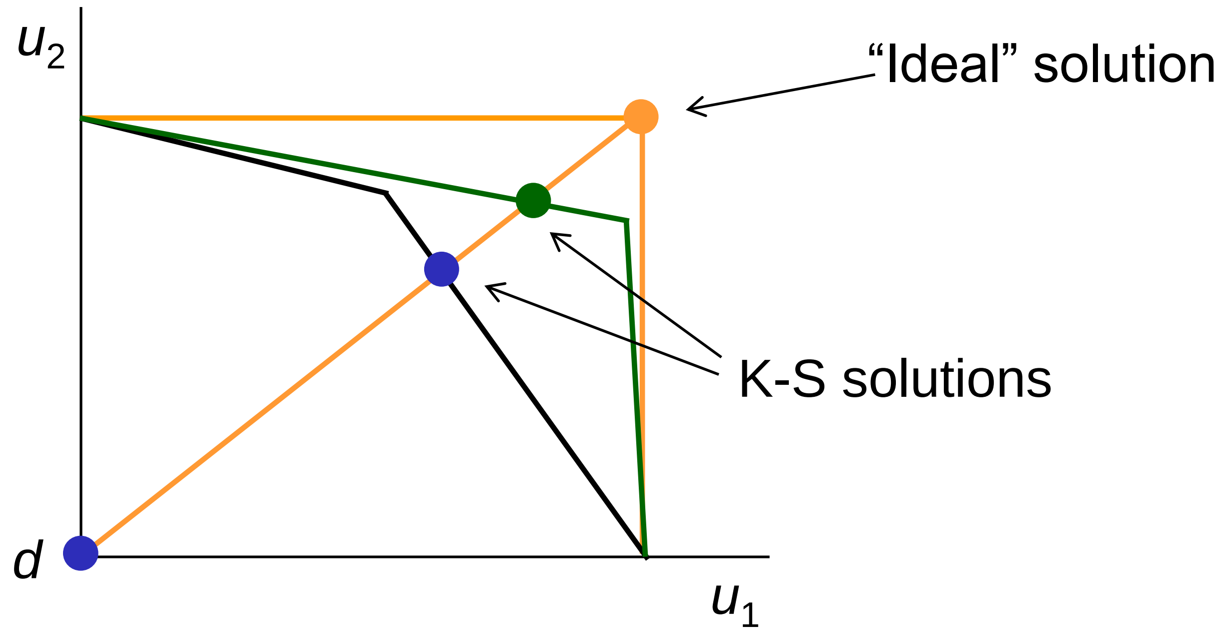

Kalai-Smorodinsky Bargaining

$$\max_{\beta,\boldsymbol{x},\boldsymbol{u}} \left\{ \beta \;\middle|\; \boldsymbol{u} = (1-\beta)\boldsymbol{d} + \beta\boldsymbol{u}^{\max},\ (\boldsymbol{u},\boldsymbol{x}) \in S,\ \beta \le 1 \right\}$$

Solution subject to budget constraint

- Same as proportional fairness.
- Seems reasonable for **price or wage negotiation.**
- Defended by some social contract theorists (e.g., "contractarians")

Gauthier 1983, Thompson 1994

Regionally decomposable…

- **…if collapsible**
  - (i.e., if it is never optimal for central authority to **take** resources **from** regions, which can be checked by simple algebraic test)

Threshold Methods

Combining utility and maximin

- **Utility threshold:** Use a maximin criterion until the utility cost becomes too great, then switch some to a utilitarian criterion.
  - Fairness is a primary concern, but without sacrificing too much utility.
  - As in a medical context, emergency facility location, task assignment.

Williams & Cookson 2000

Threshold Methods

Combining utility and maximin

- **Utility threshold:** Use a maximin criterion until the utility cost becomes too great, then switch some to a utilitarian criterion.
  - 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 to a maximin criterion.
  - Use when efficiency is the primary concern, but without excessive sacrifice by any individual.
  - As in telecommunications, disaster recovery, traffic control..

Williams & Cookson 2000

Utility Threshold

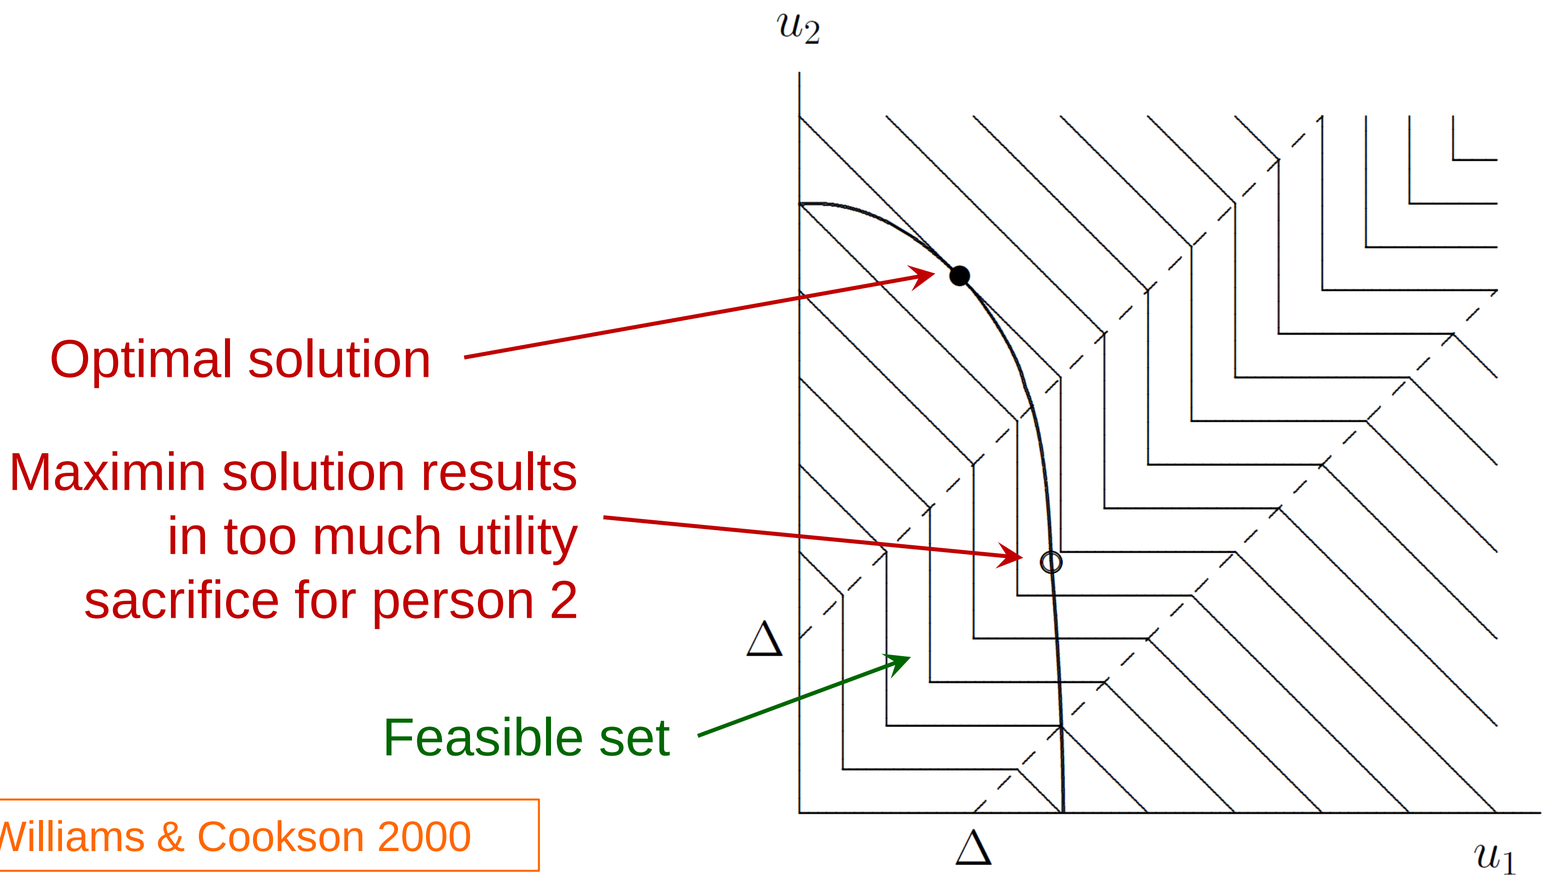

Williams & Cookson 2000

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

Utility Threshold

Generalization to *n* persons

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

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

JH & 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**.

Solution subject to budget constraint

- Purely **utilitarian** for smaller values of Δ, **maximin** for larger values.

Utility Threshold

Theorem. When maximizing the SWF subject to a **budget constraint**, the optimal solution is purely **maximin** or **purely utilitarian**.

Purely maximin if

$$\Delta \geq B\left(\frac{1}{a_{\langle 1\rangle}} - \frac{n}{\sum_i a_i}\right)$$

Here, parties have **similar** treatment costs, or Δ is **large**.

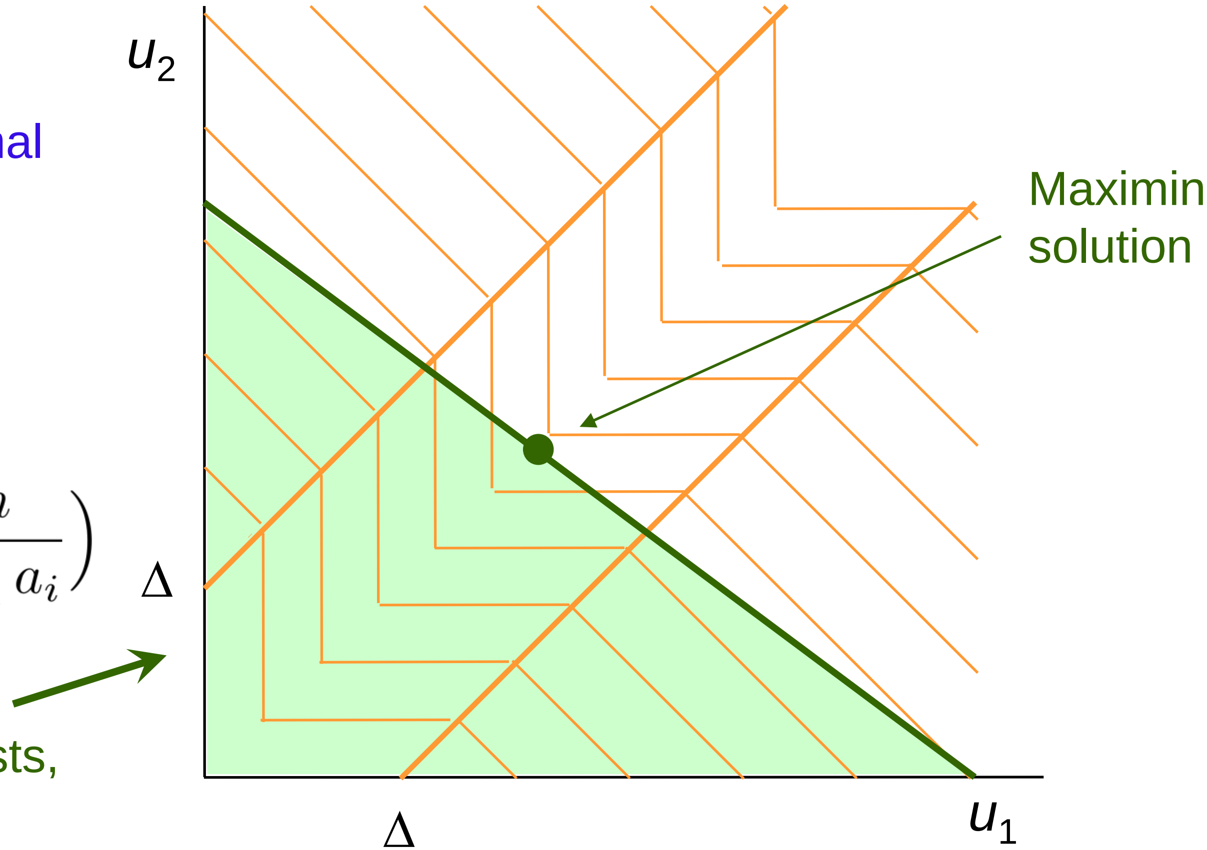

Utility Threshold

Theorem. When maximizing the SWF subject to a **budget constraint**, the optimal solution is purely **maximin** or **purely utilitarian**.

Purely utilitarian if

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

Here, parties have **very different** treatment costs, or Δ is **small**.

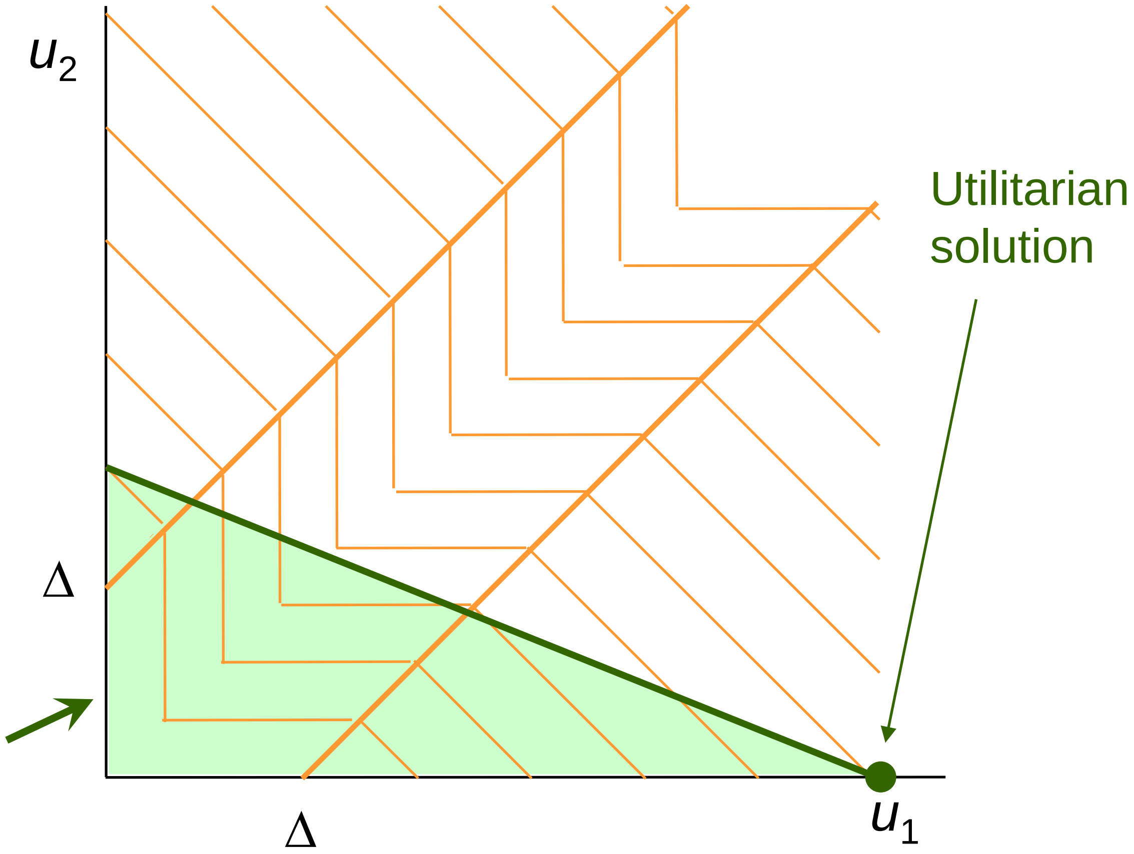

Utility Threshold

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.

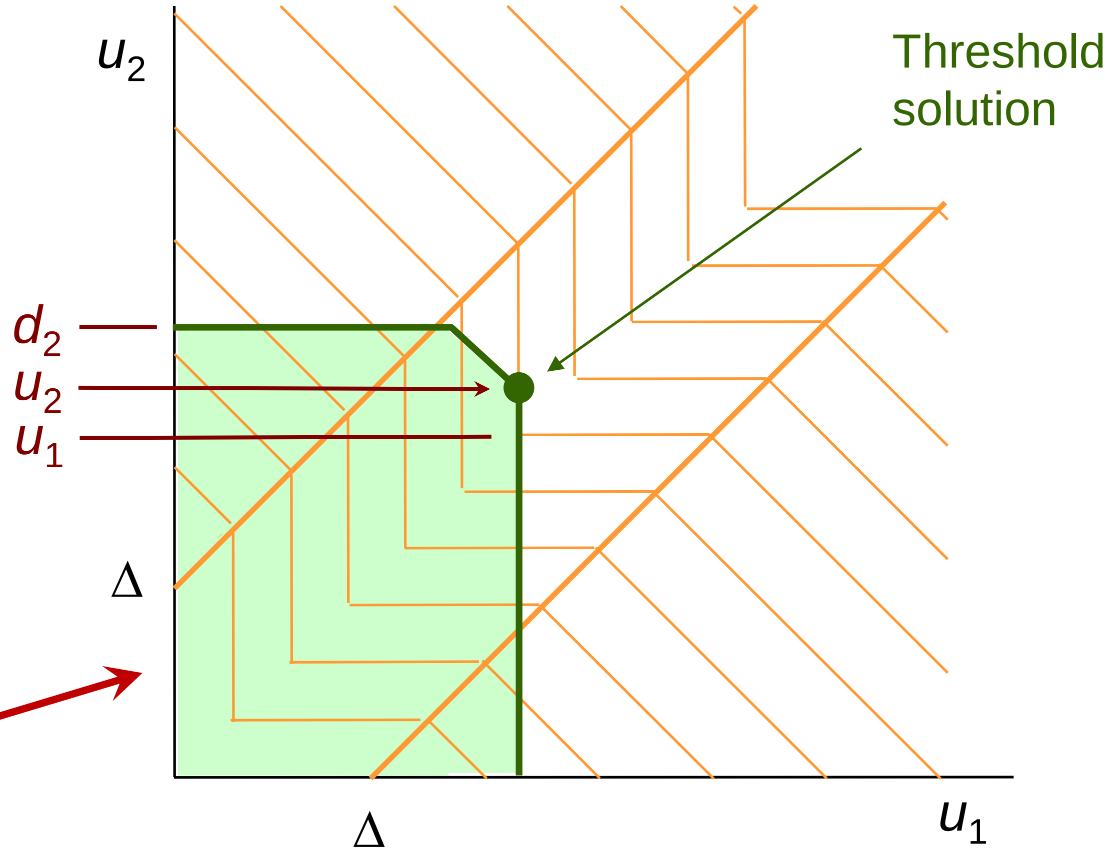

Utility Threshold

Not regionally decomposable

- This could be an advantage or disadvantage.

Incentive and cross-subsidy rates:

- Same as **utilitarian** (for small Δ) or **maximin** (for large Δ)

Equity Threshold

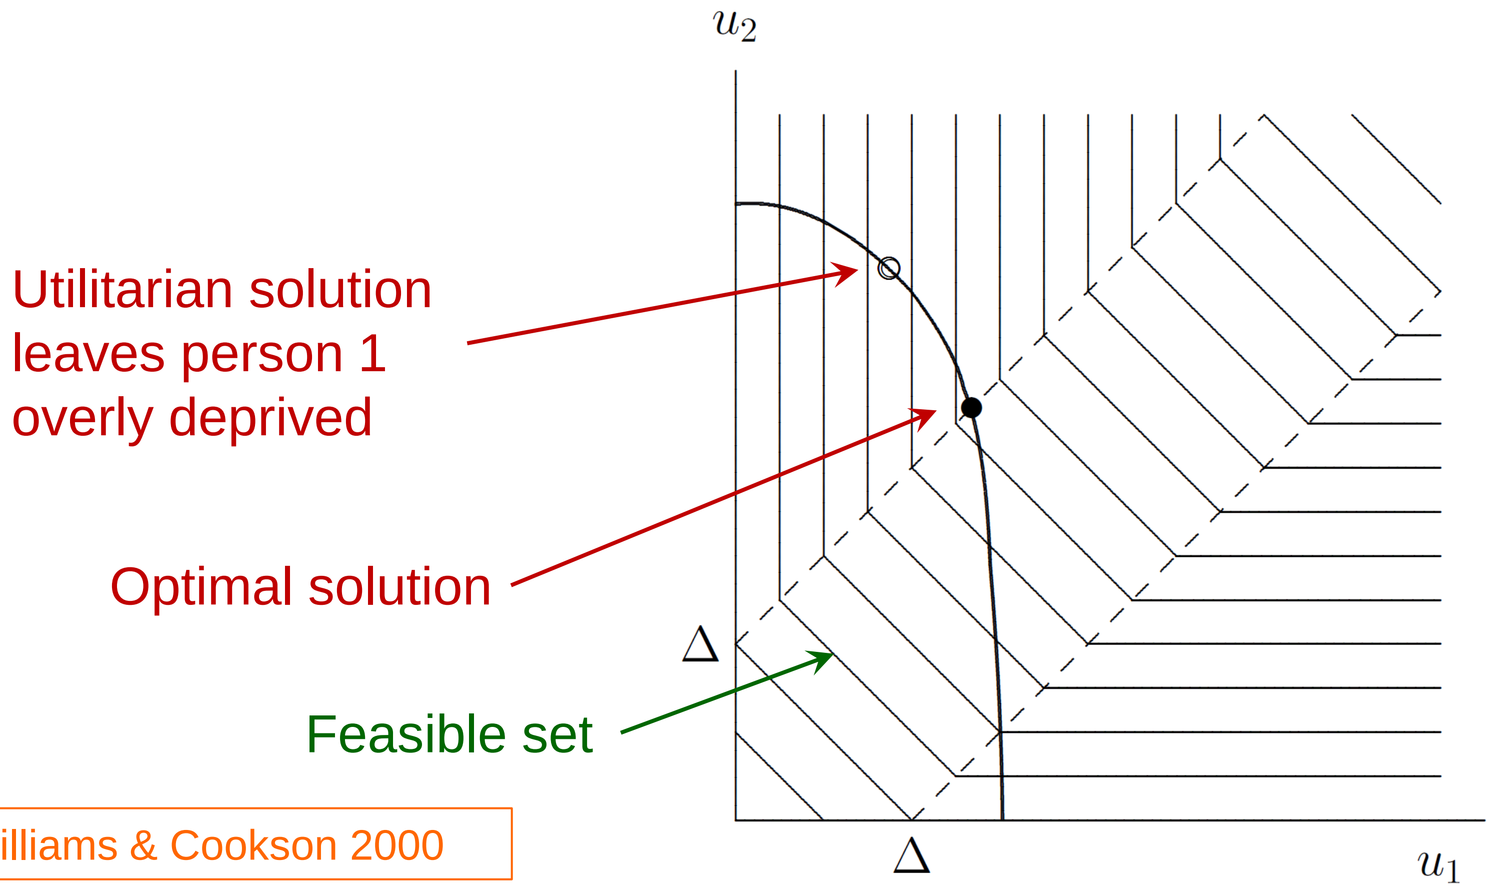

Williams & Cookson 2000

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

Equity Threshold

Generalization to *n* persons

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

- Δ is chosen so that well-off individuals **do not deserve more utility** unless utilities within Δ of smallest are also increased.
- Values **reversed**: $\Delta = \infty$ corresponds to utilitarian, $\Delta = 0$ to maximin.

Solution subject to budget constraint

- For **large** (more utilitarian) values of Δ, **more efficient** individuals get **utility** $\boldsymbol{\Delta}$, less efficient get **zero**.
- For **small** (more egalitarian) values of Δ, **everyone** gets something, but **more efficient** individuals get $\boldsymbol{\Delta}$ **more utility** than less efficient.

Equity Threshold

Example:

Maximum equity threshold SWF subject to budget constraint

$$u_1 + 2u_2 + \cdots + 8u_8 \leq 100$$

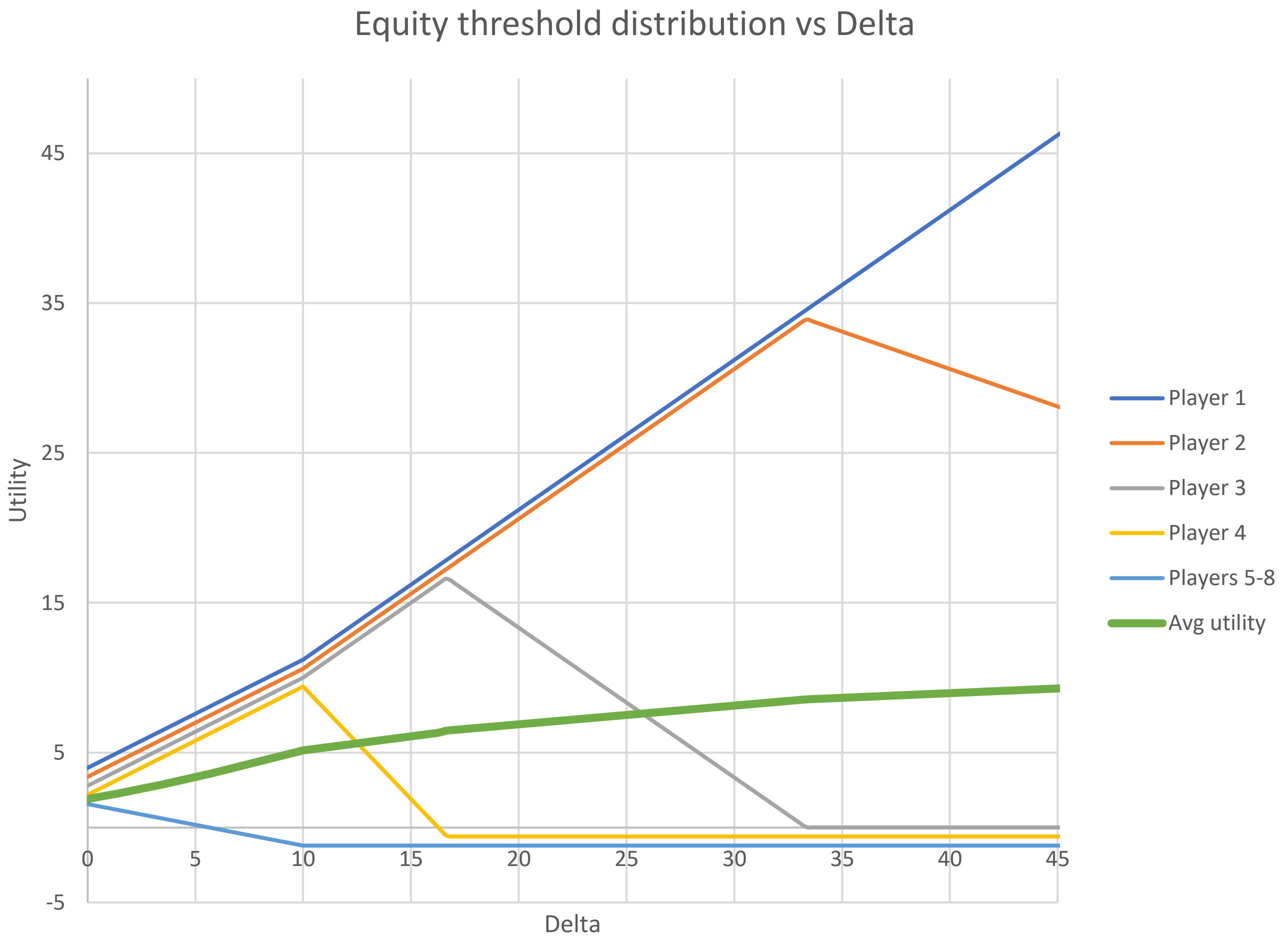

Equity Threshold

Not regionally decomposable

Incentive rate:

- For **large** (more utilitarian) Δ, **rate = 1** for **one person** with a certain intermediate utility level, zero for others
- For **small** (more egalitarian) Δ, rate is **positive**: $\frac{a_i}{\sum_j a_j}$ for any individual *i*.

Cross-subsidies:

- For **large** (more utilitarian) Δ, subsidies are **zero,** except **positive** for the **one person** with an intermediate utility level, who benefits from the improvements of some others (namely, those with greater efficiencies).
- For **small** (more egalitarian) Δ, positive subsidies for all: $\frac{a_i}{\sum_j a_j}$

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

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

Chen & JH 2021

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

Solution subject to budget constraint

- **The *m* most efficient** individuals receive **equal** utility $\dfrac{a_i}{\sum_{j=1}^{m} a_j}$, others **zero**.
- Larger Δ spreads utility over more individuals (larger *m*).

Utility Threshold with Leximax

Example:

Maximum utility threshold SWF with leximax subject to budget constraint

$$u_1 + 2u_2 + \cdots + 8u_8 \leq 100$$

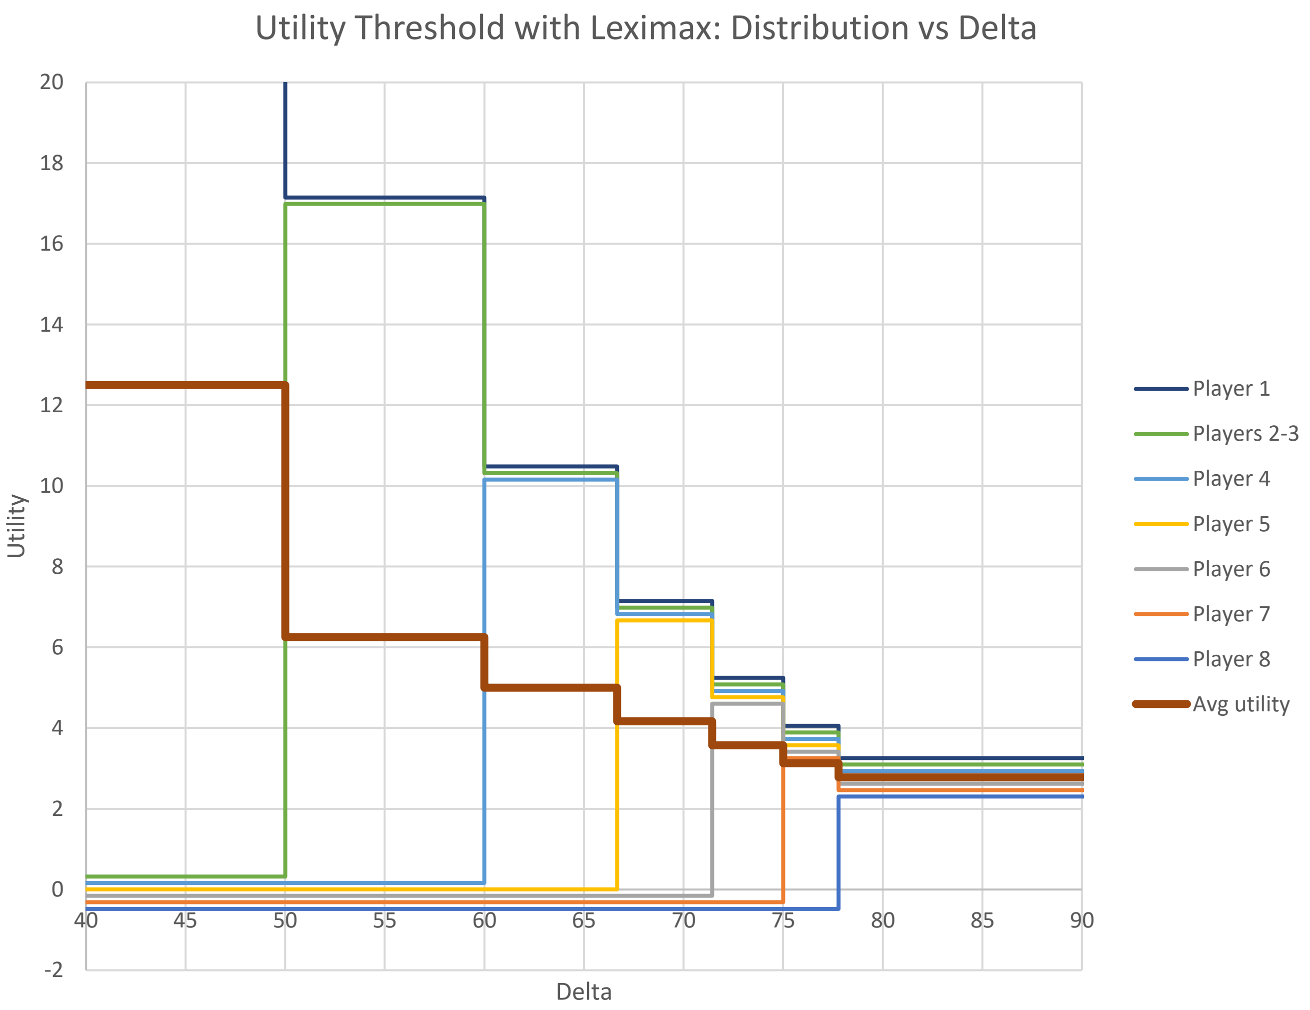

Utility Threshold with Leximax

Not regionally decomposable

Incentive rate:

- Individuals who receive **positive utility** have **positive** rate $\frac{a_i}{\sum_{j=1}^{m} a_j}$, others **zero**

Cross-subsidies:

- **Positive** subsidies $\frac{a_i}{\sum_{j=1}^{m} a_j}$ to those who receive positive utility
- **Zero** for others.

Example of Utility Threshold with Leximax

- Select earthquake shelter locations in Istanbul.
- Utility = negative distance of each neighborhood to nearest shelter, subject to limited budget.
- 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

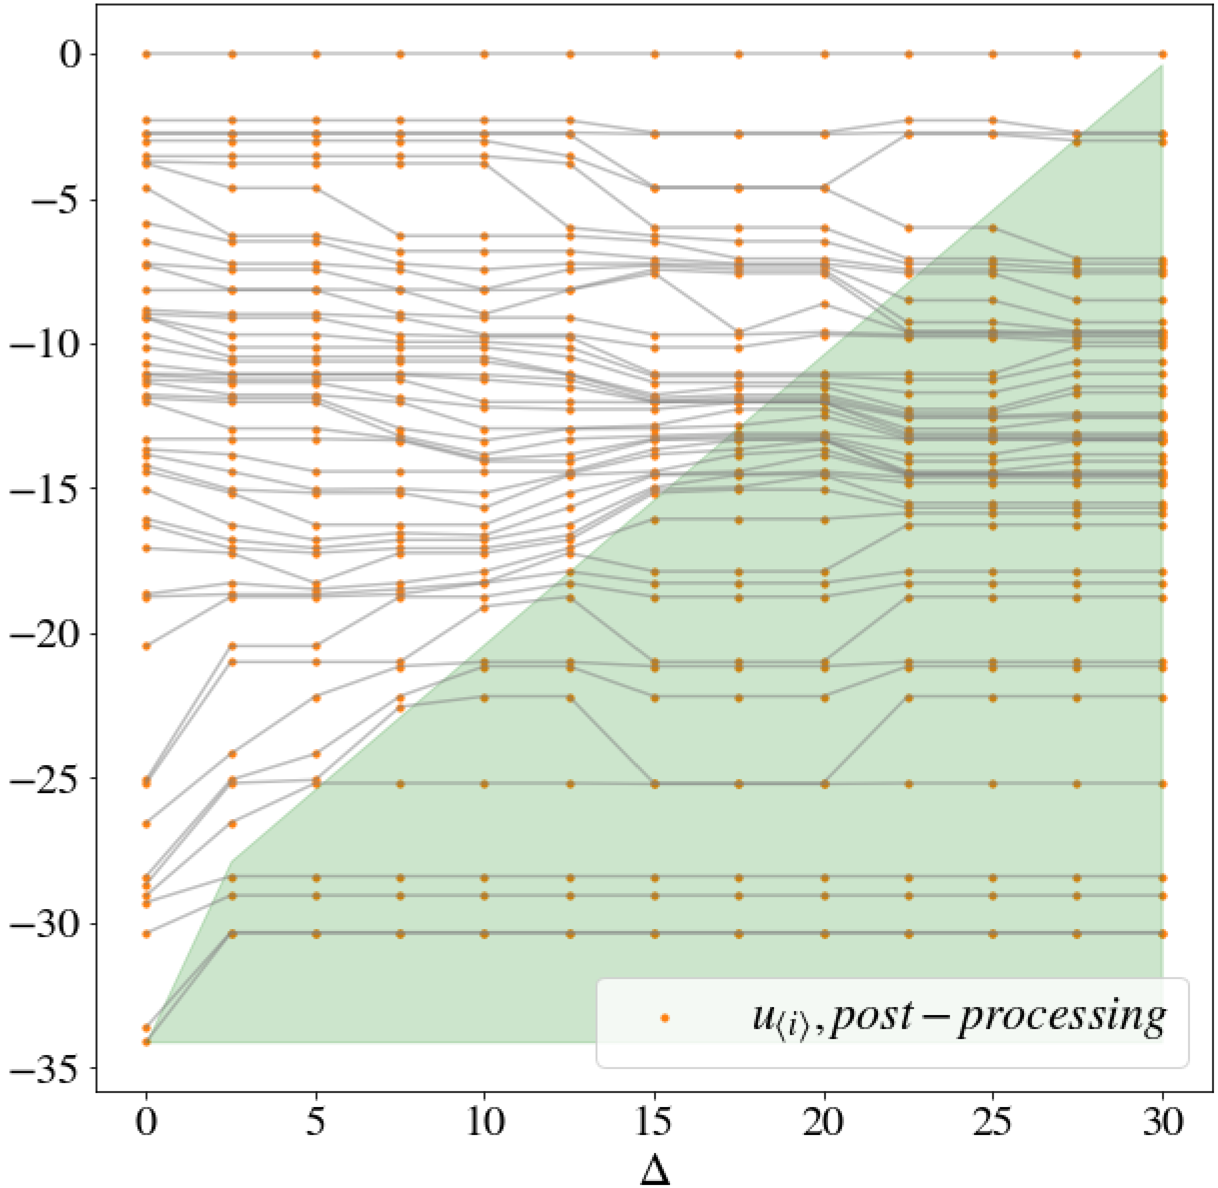

Properties of Fair Solutions

Social welfare criterion	**Solution structure** with simple budget constaint	Special comment
Utilitarian	Most efficient party gets **everything**	Traditional objective
Maximin/leximax	Everyone gets **equal** utility	Leximax **avoids wasting** utility
Alpha fairness	Fairness **increases with α**	**Utilitarian** when $\alpha = 0$, **maximin** when $\alpha \to \infty$
Kalai-Smorodinsky	**Same solution** as alpha fairness with $\alpha = 1$ (**proportional fairness**)	Utility allotment is **proportional to efficiency**
Utility threshold with maximin	**Purely utilitarian or maximin**, depending on Δ	Interesting structure when **bounds** are added
Equity threshold with maximin	More efficient parties receive **Δ more** than less efficient parties	Least efficient parties receive **zero**
Utility threshold with leximax	**More efficient** parties receive **equal utility**, others zero	For larger Δ, more parties receive utility but **smaller allotment**

Properties of Fair Solutions

Social welfare criterion	**Regionally decomposable?**	**Incentives and sharing** with simple budget constaint
Utilitarian	**Yes**	**Only most efficient party** incentivized to improve efficiency, **no sharing**
Maximin/leximax	**Yes**	**Less efficient** parties have **greater incentive** to improve, benefits **shared equally**
Alpha fairness	**Yes**	**Less efficient** parties have **greater incentive. Competitive** when $\alpha < 1$, **cooperative** when $\alpha > \infty$
Kalai-Smorodinsky	**Yes**, if collapsible	Same as **proportional fairness** ($\alpha = 1$)
Utility threshold with maximin	**No**	Same as **utilitarian** or **maximin**, depending on Δ
Equity threshold with maximin	**No**	For larger Δ, only **one party** incentivized to improve and receives all benefits. For smaller Δ, **all** are incentivized and benefit.
Utility threshold with leximax	**No**	Parties who receive **positive utility** are incentivized to improve and **share** benefits of efficiency improvement.


Questions or
comments?