Structural Properties of Fair Solutions

John Hooker
Carnegie Mellon University

Joint work with

Özgün Elçi
Amazon

Peter Zhang
Carnegie Mellon University

INFORMS 2023

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

- 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
$$u_1 + u_2$$

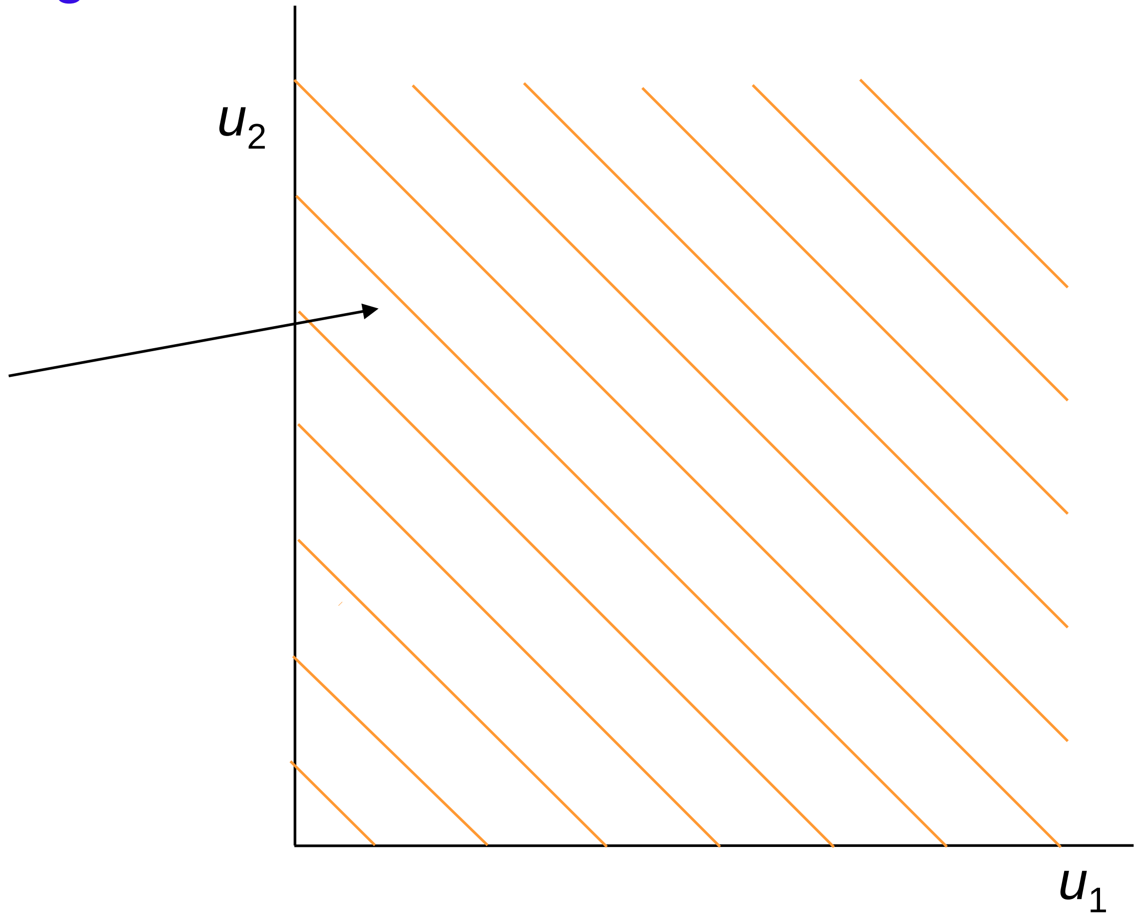

Maximize Utility?

Utility maximizing
distribution
for 2 persons,
subject to
budget constraint

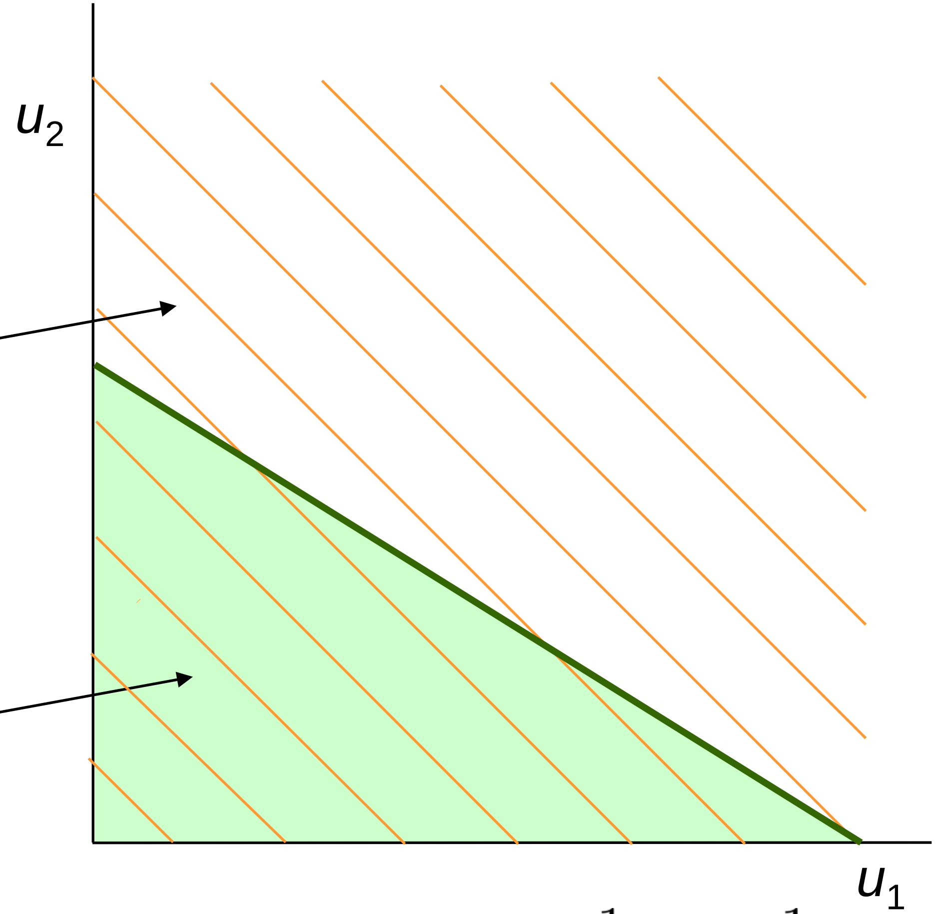

Utility contours

$$u_1 + u_2$$

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

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

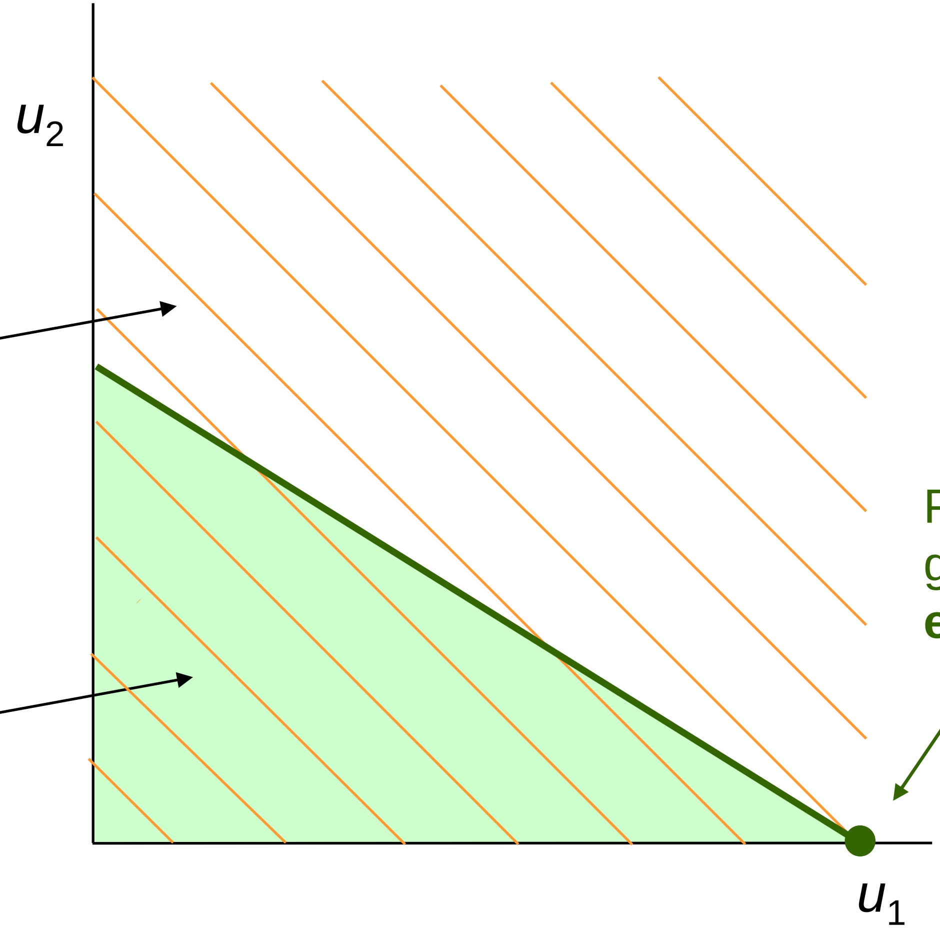

Person 1 gets **everything**!

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

Maximize Utility?

Add **bounds** on individual utilities…

Utility contours

$$u_1 + u_2$$

Feasible region

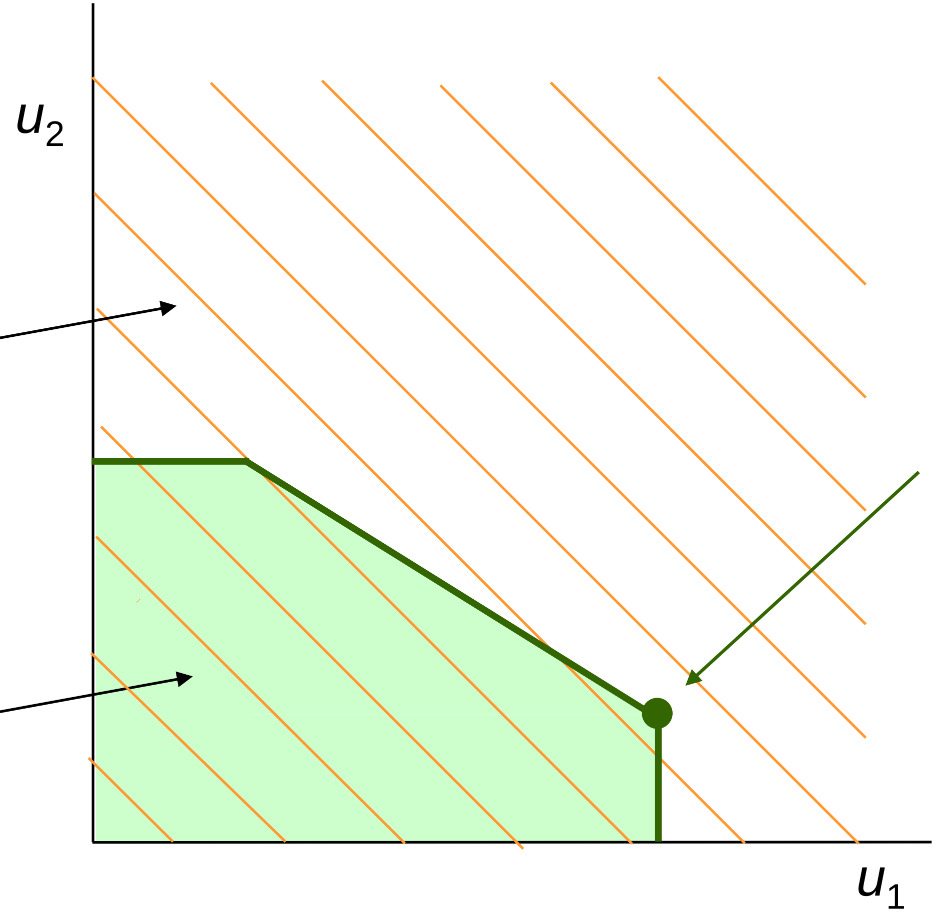

The Problem

- True, these constraints are simplistic...
  - ...and such extreme solutions **rarely occur** in practice.

The Problem

- True, these constraints are simplistic...
    - ...and such extreme solutions **rarely occur** in practice.
    - This is only because complex constraints **happen to rule out** extremely unfair solutions.
    - The constraints only **conceal the basic inadequacy** of the objective function!
- We need an objective function that **balances utility and fairness**.

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**.

Modeling Fairness

- This is an ***ex post*** approach
  - …as opposed to the traditional ***ex ante*** approach of **social choice theory**
  - …which derives fairness criteria from **axioms of rational choice** or **bargaining arguments**.
  - These make strong **assumptions** that are unrealistic or difficult to assess in practice.

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

Reciprocals of conversion efficiencies

Individual utilities

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

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.

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**.

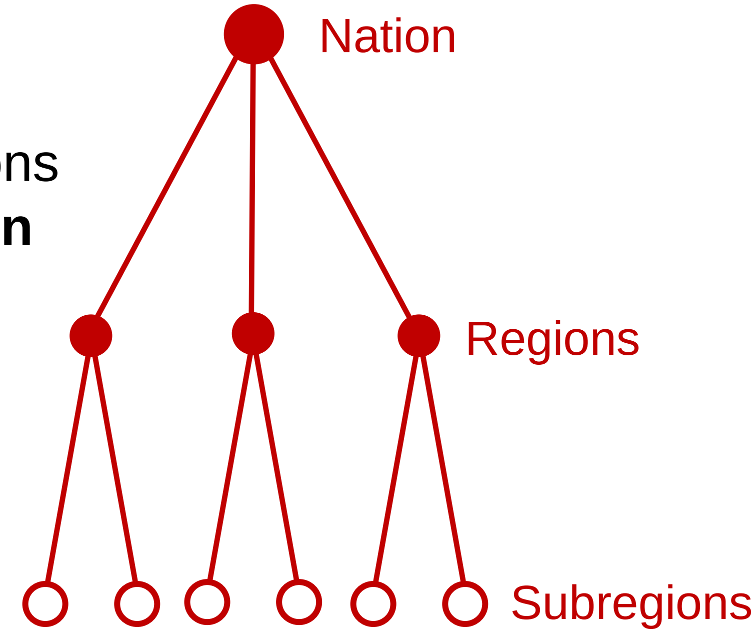

Hierarchical Distribution

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

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

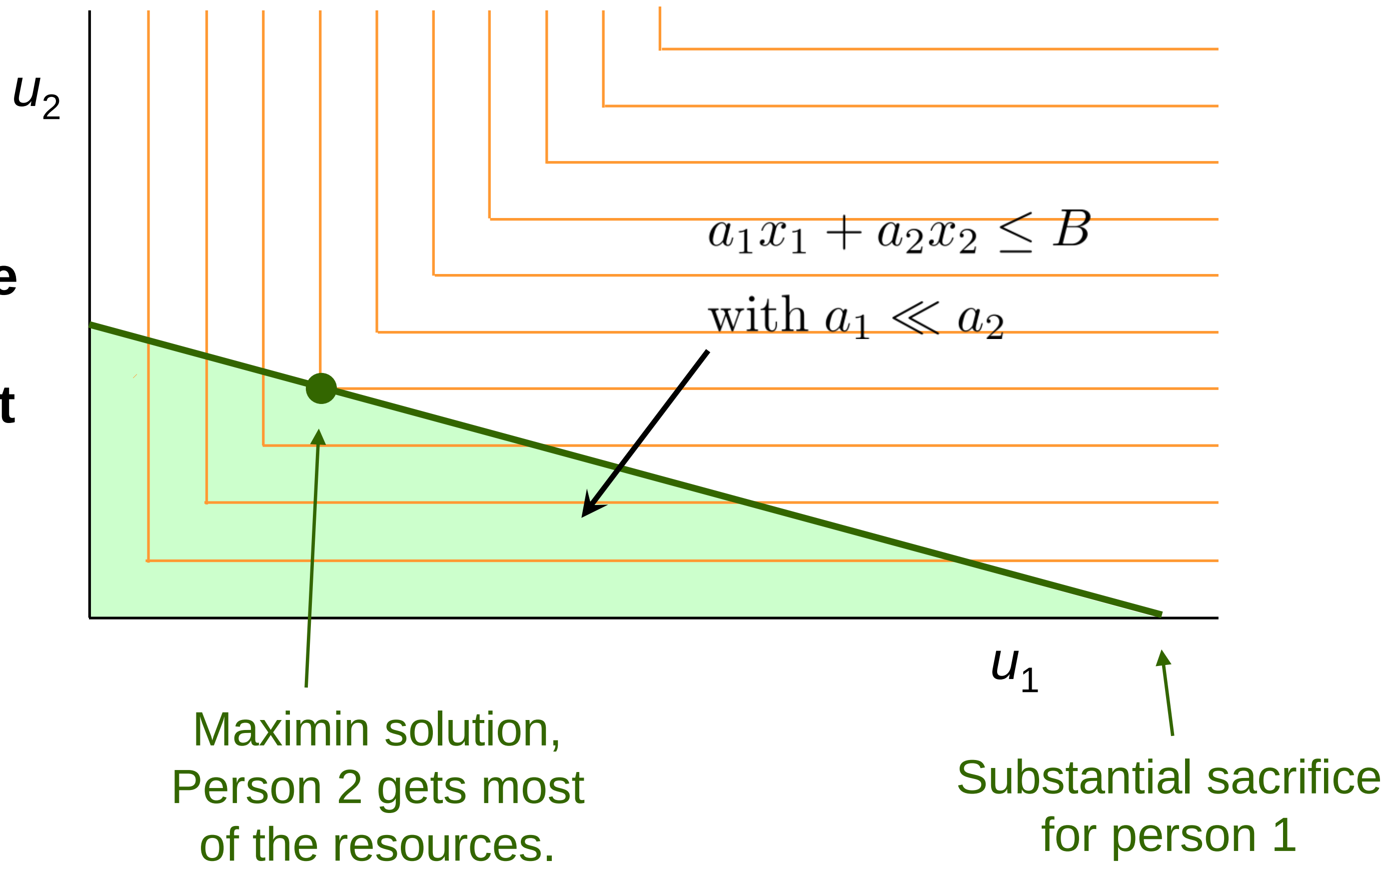

In a medical context, patient 1 is reduced to same level of suffering as seriously ill patient 2.

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

Medical example with resource bounds

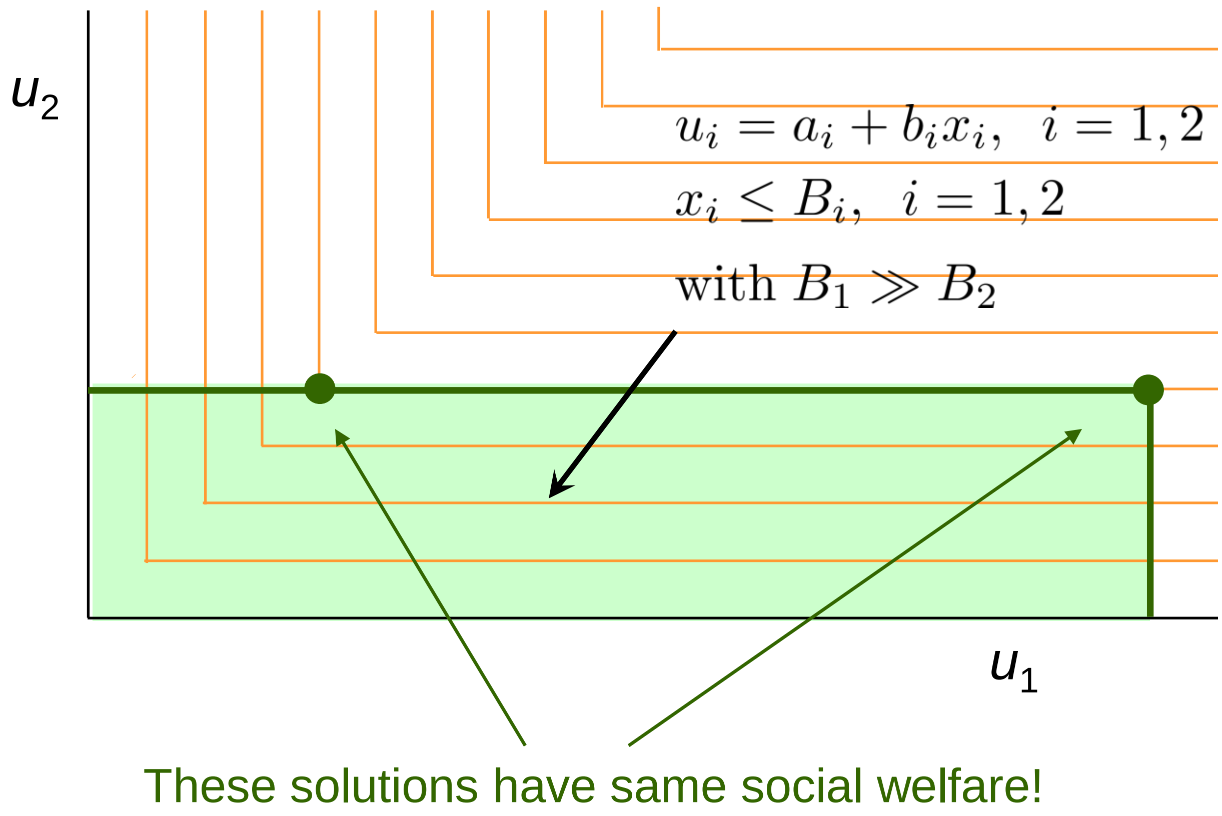

Fairness for the Disadvantaged

Maximin

Medical example with resource bounds

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

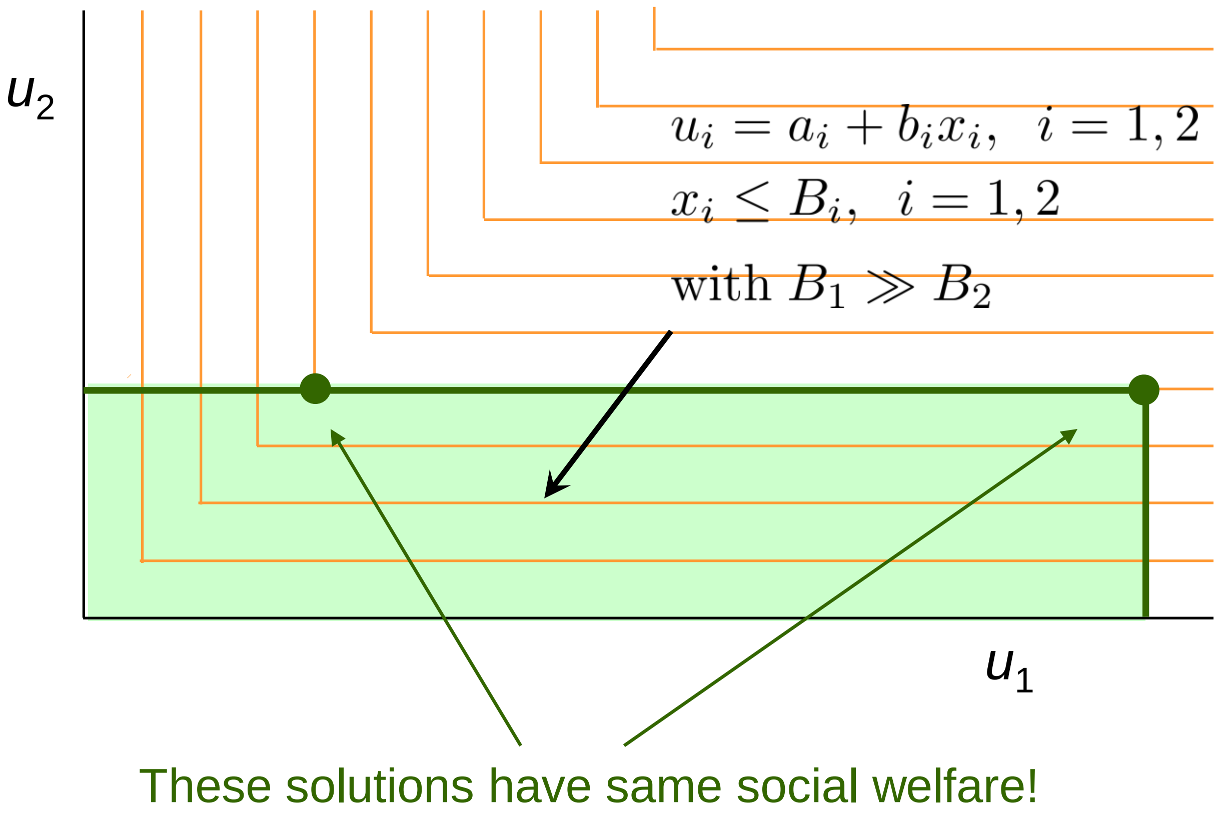

Maximin

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

Regionally decomposable?

- Monotonically separable SWF → **yes**.

Incentive rate for person *i* ? $\dfrac{a_i}{\sum_j a_j}$

- Less efficient parties have greater incentive to improve.

Cross-subsidy rate? $\dfrac{a_i}{\sum_j a_j}$

- 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?

- **Yes** (using generalized definition of decomposability)

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

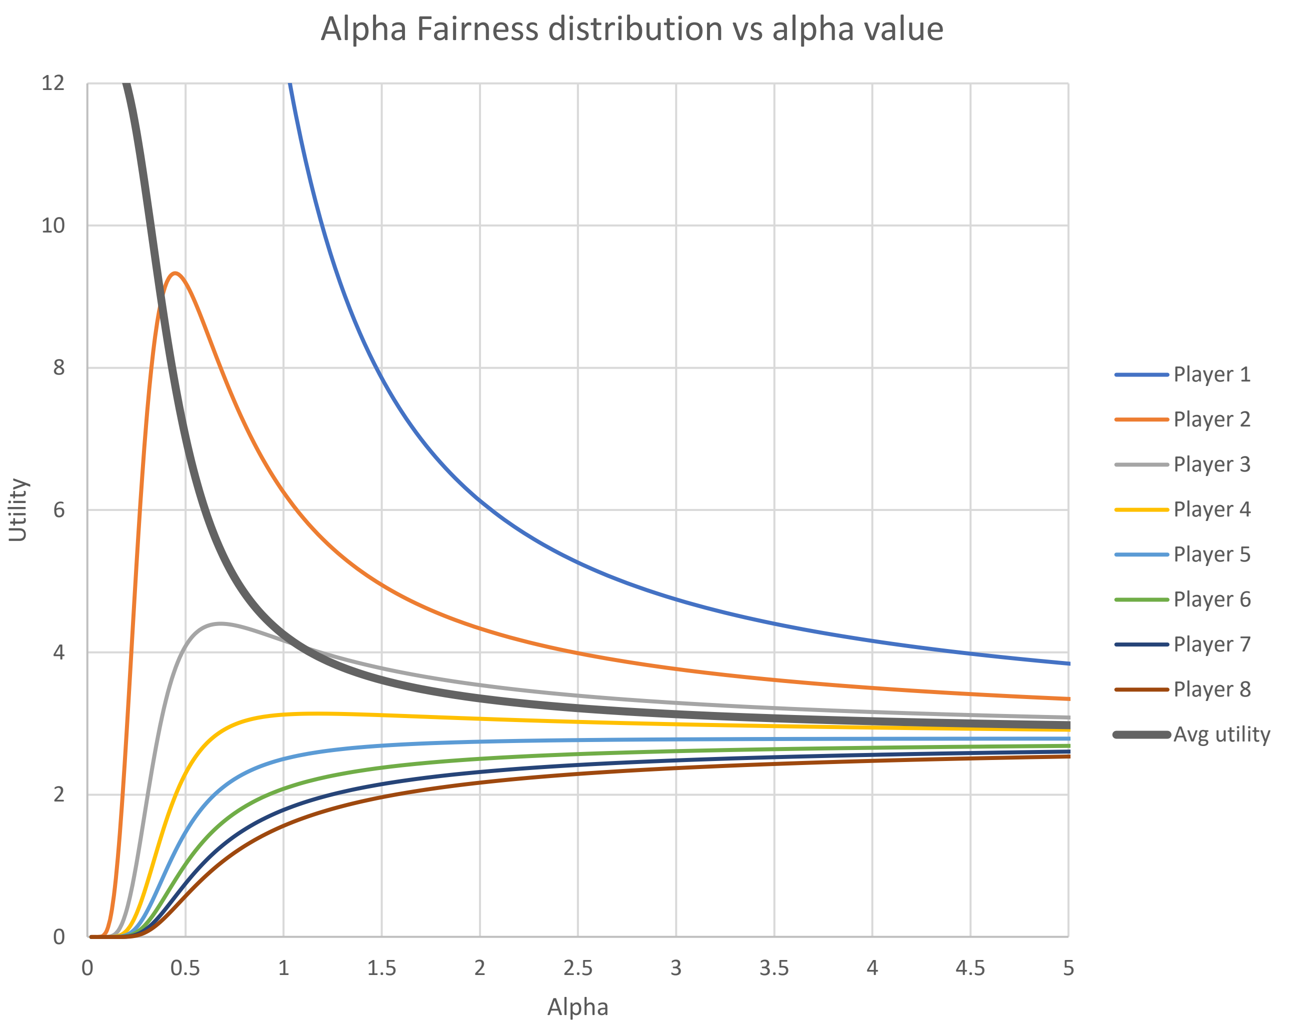

Alpha Fairness

Regionally decomposable?

- Separable SWF → **yes**.

Incentive rate for person *i* :

$$\frac{1}{\alpha} + \left(1 - \frac{1}{\alpha}\right) \frac{a_i^{1-1/\alpha}}{\sum_j a_j^{1-1/\alpha}}$$

- **More efficient** persons have **greater** incentive to improve efficiency when $\alpha < 1$, **less** incentive when $\alpha > 1$.

Cross-subsidy rates:

$$\left(1 - \frac{1}{\alpha}\right) \frac{a_i^{1-1/\alpha}}{\sum_j a_j^{1-1/\alpha}}$$

- When $\alpha < 1$ (**competition**), efficiency improvements transfer utility **from** other persons
- 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-subsidy rates = 0

Proportional Fairness

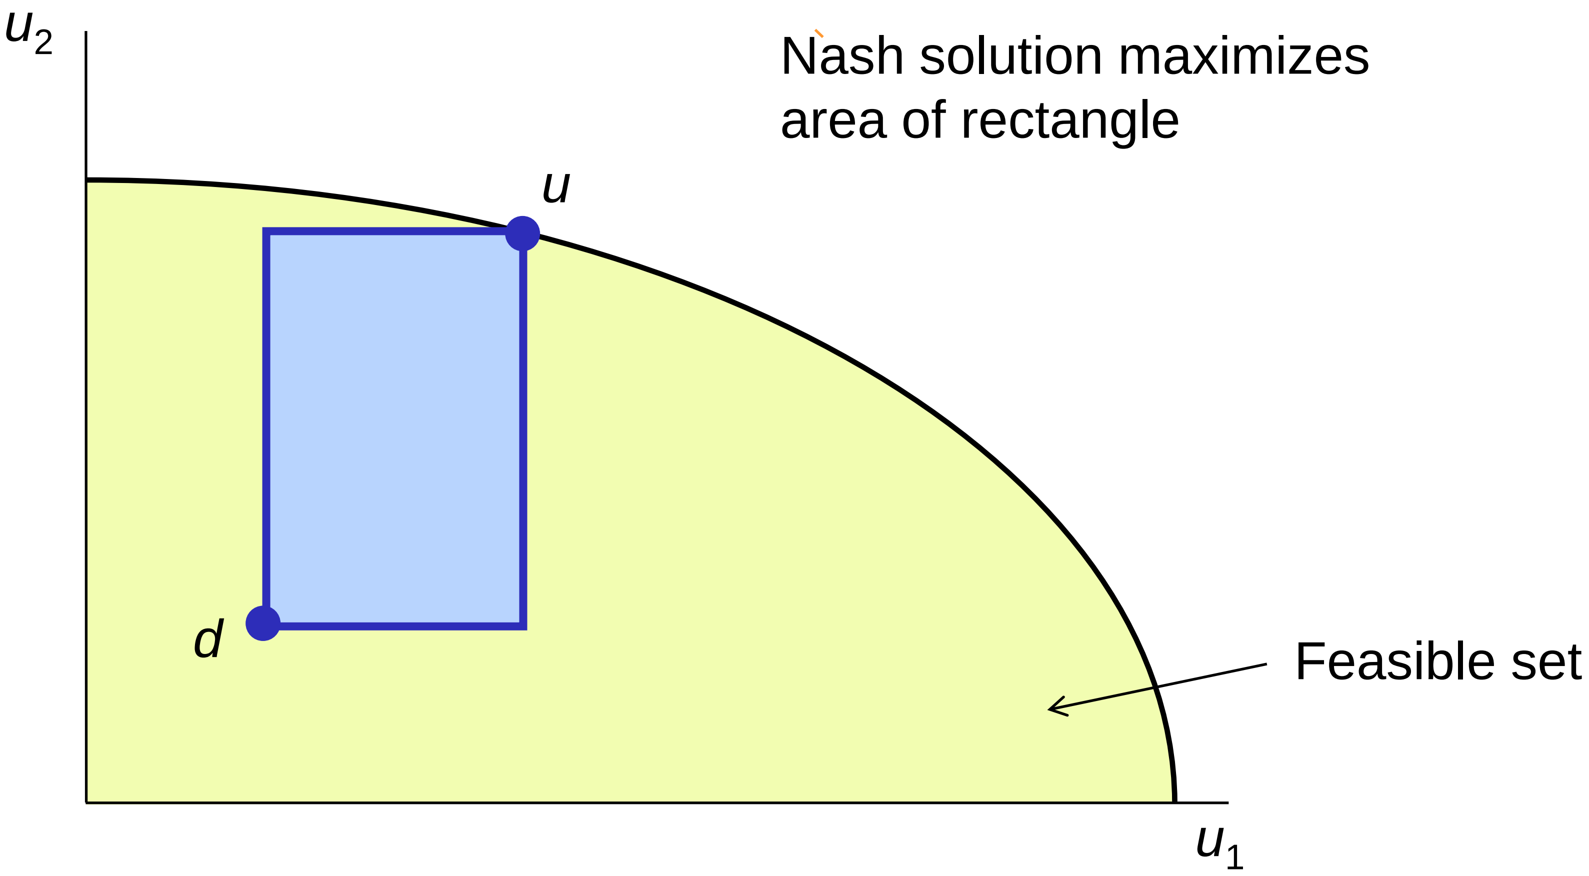

Proportional Fairness

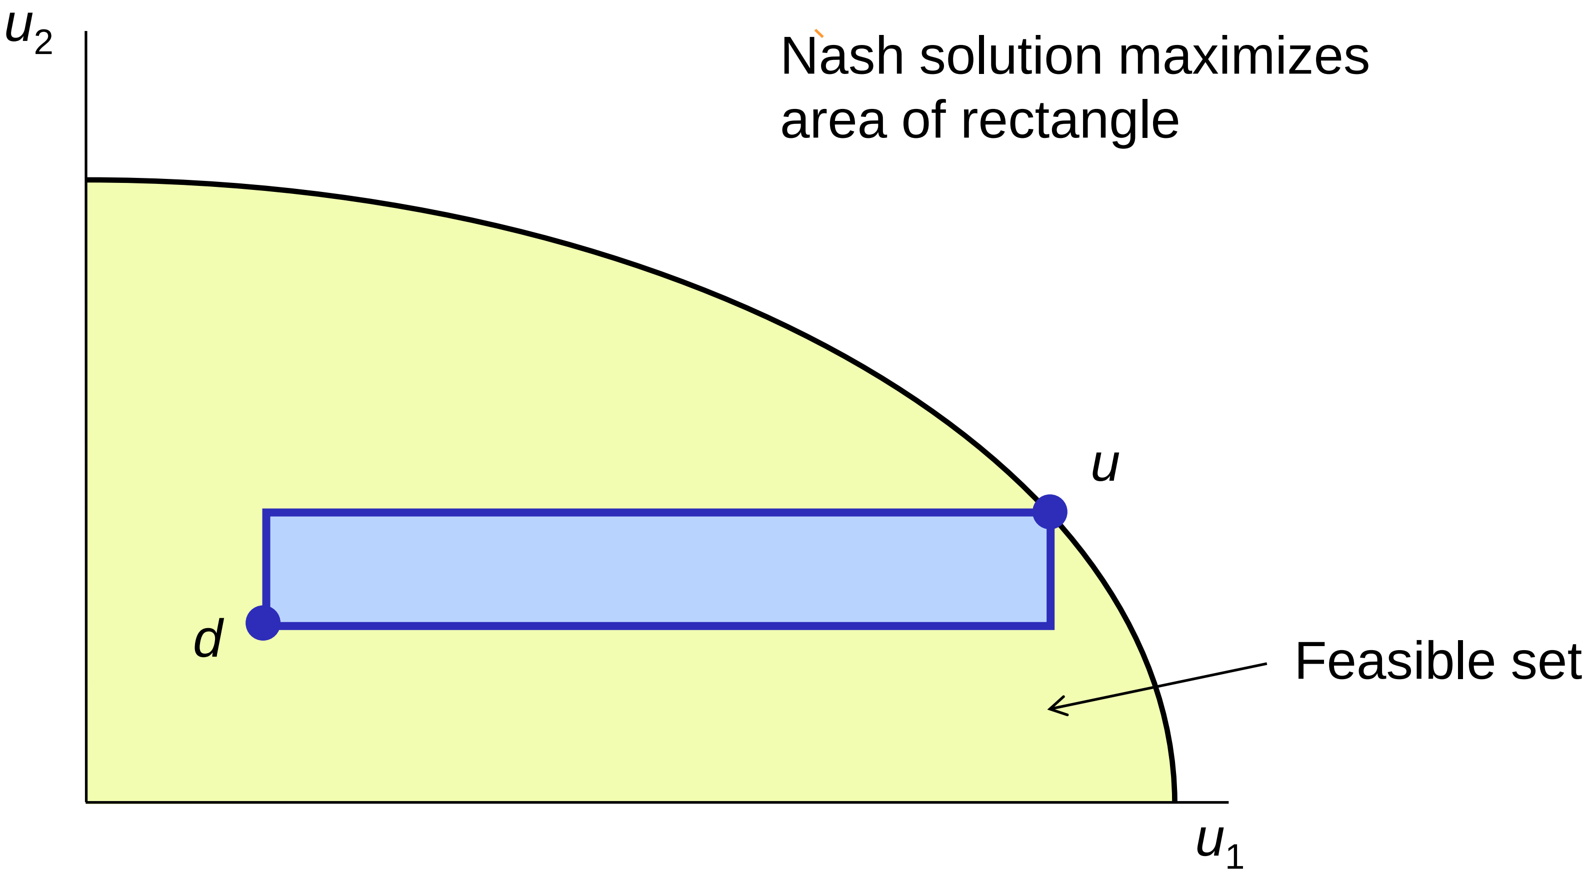

Proportional Fairness

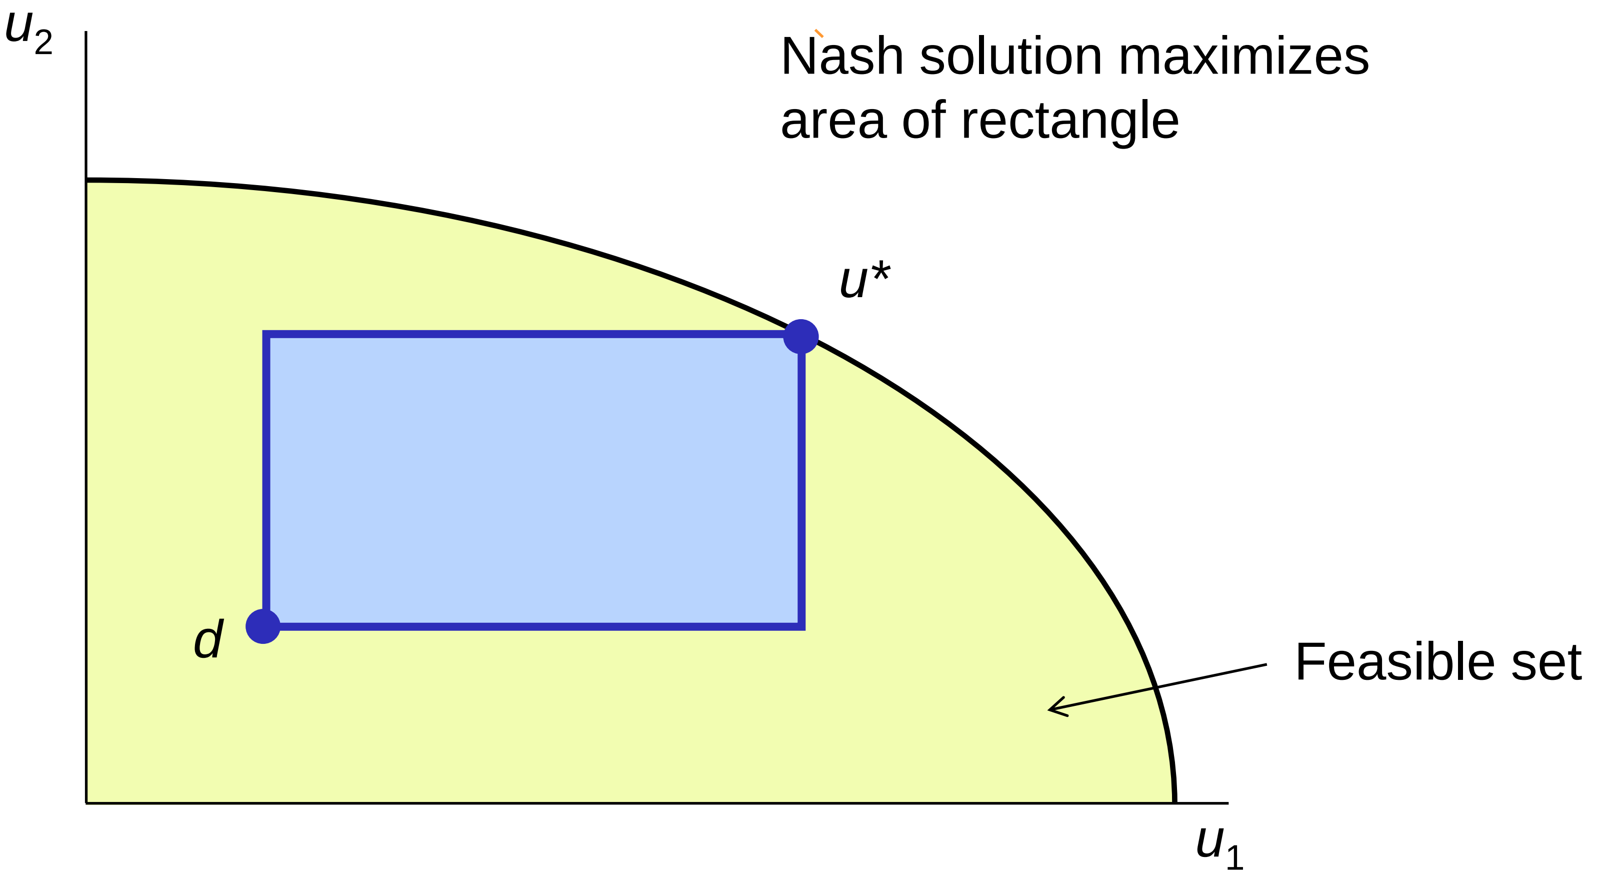

Kalai-Smorodinsky Bargaining

- Begins with a critique of the Nash bargaining solution.

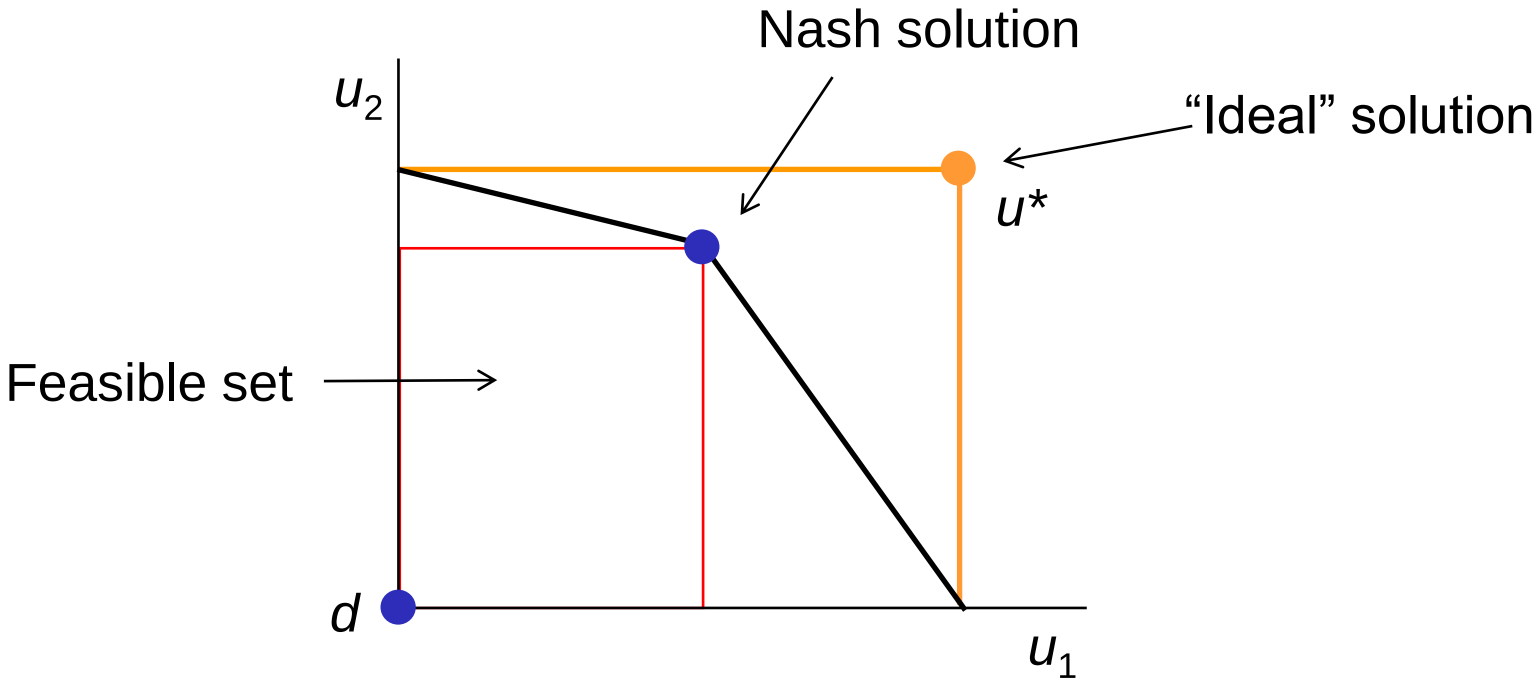

Kalai-Smorodinsky Bargaining

- 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**.

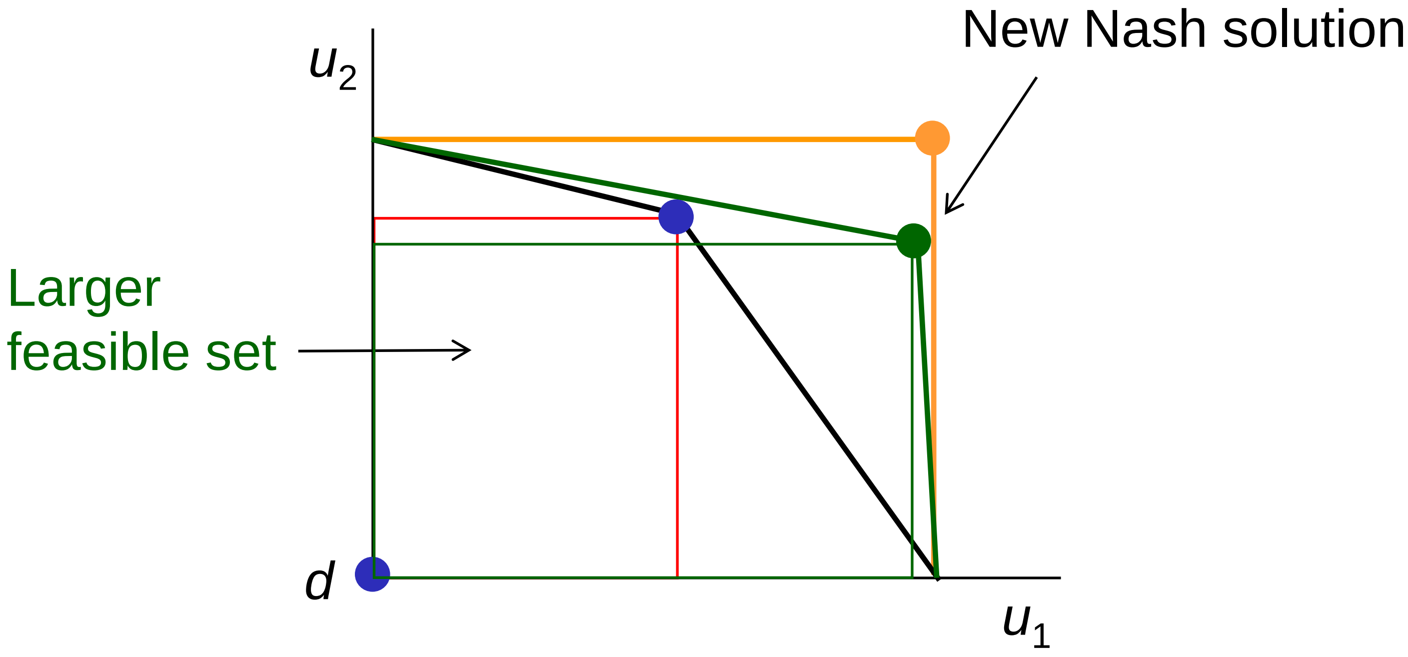

Kalai-Smorodinsky Bargaining

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

Kalai & Smorodinksy 1975

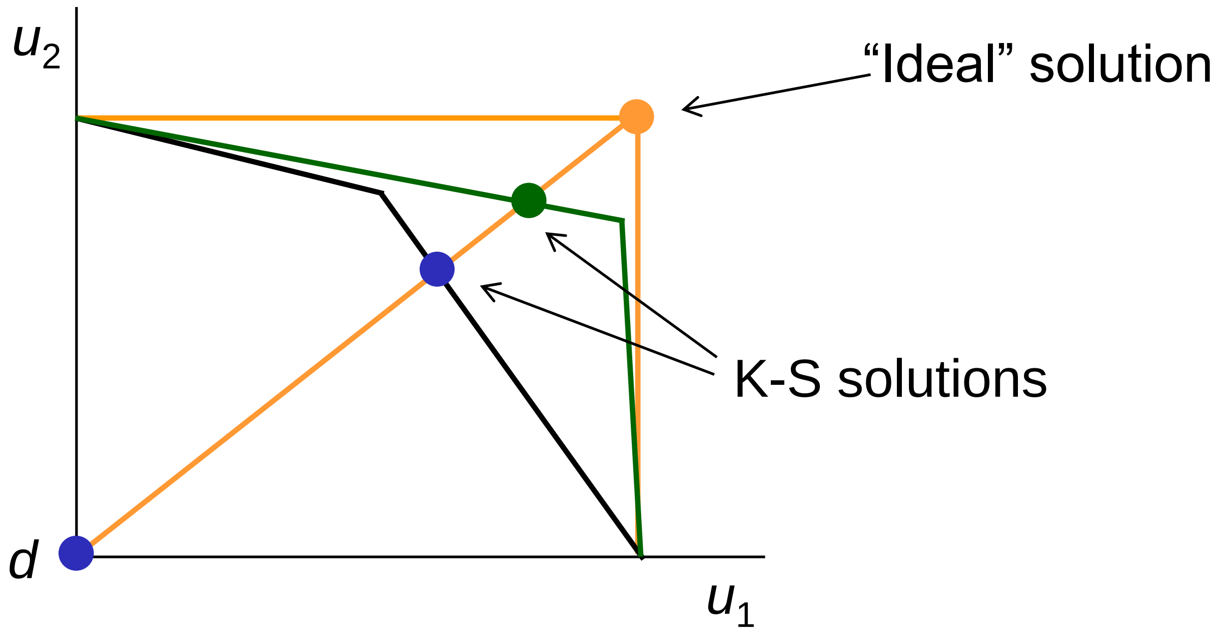

Kalai-Smorodinsky Bargaining

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

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?

- **Yes, if collapsible**
  - (i.e., 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, 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, 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

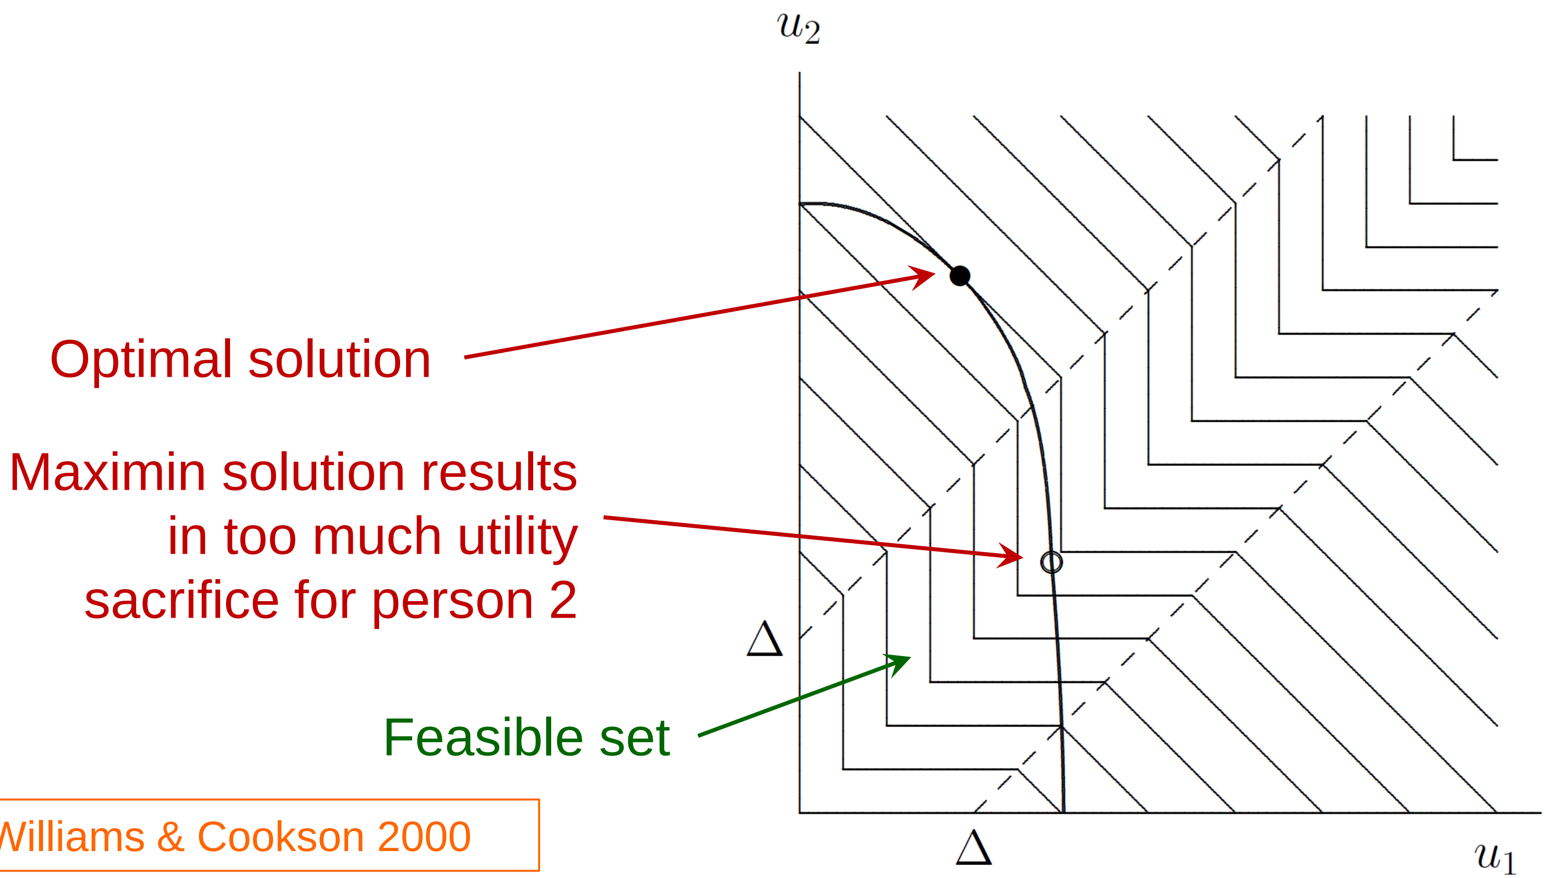

Williams & Cookson 2000

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

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

Solution subject to budget constraint

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

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**.

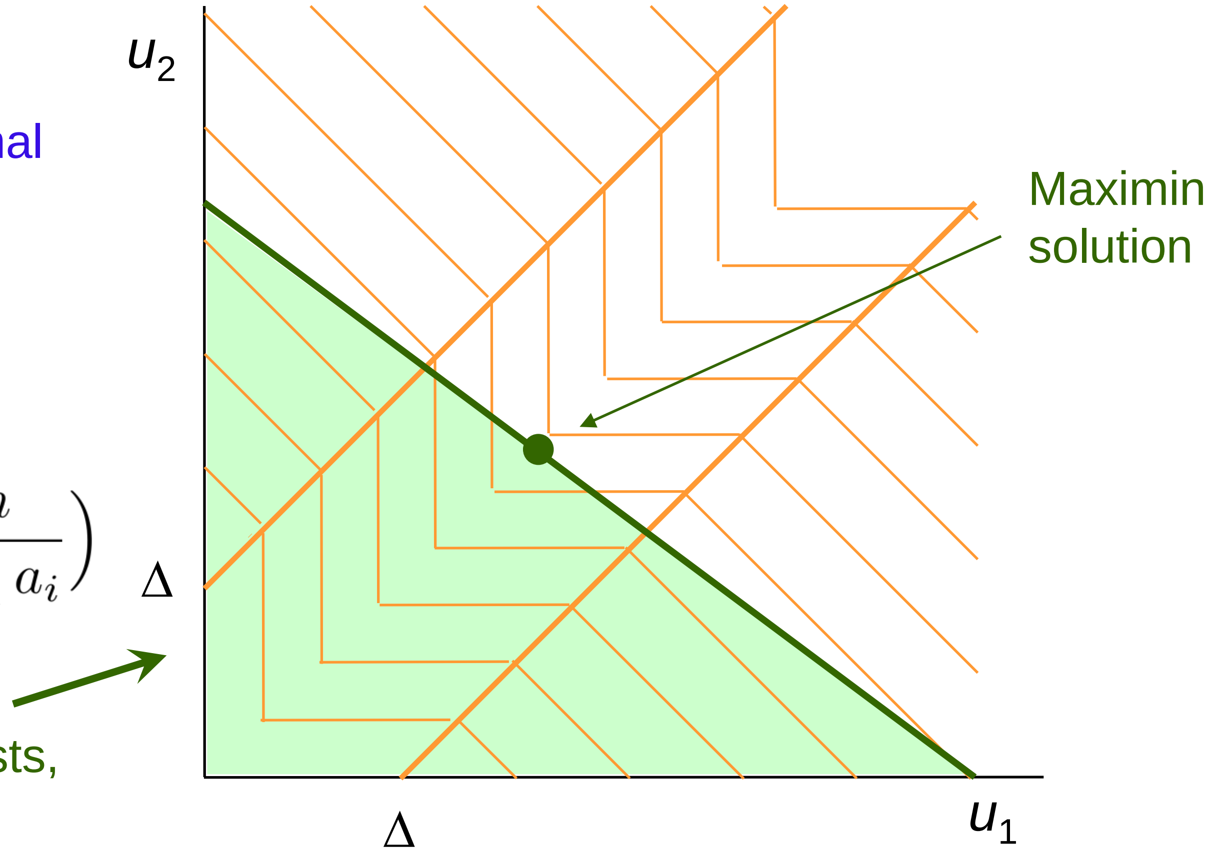

Utility Threshold

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**.

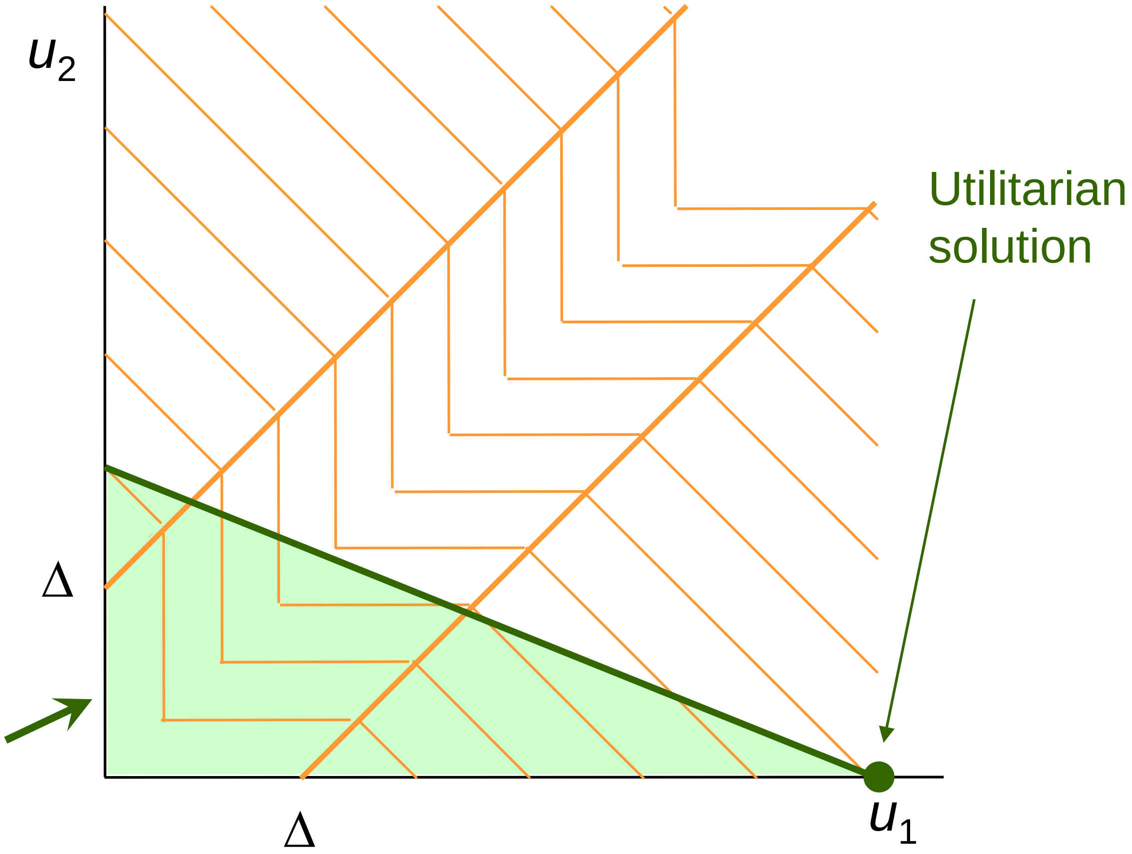

Utility Threshold

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.

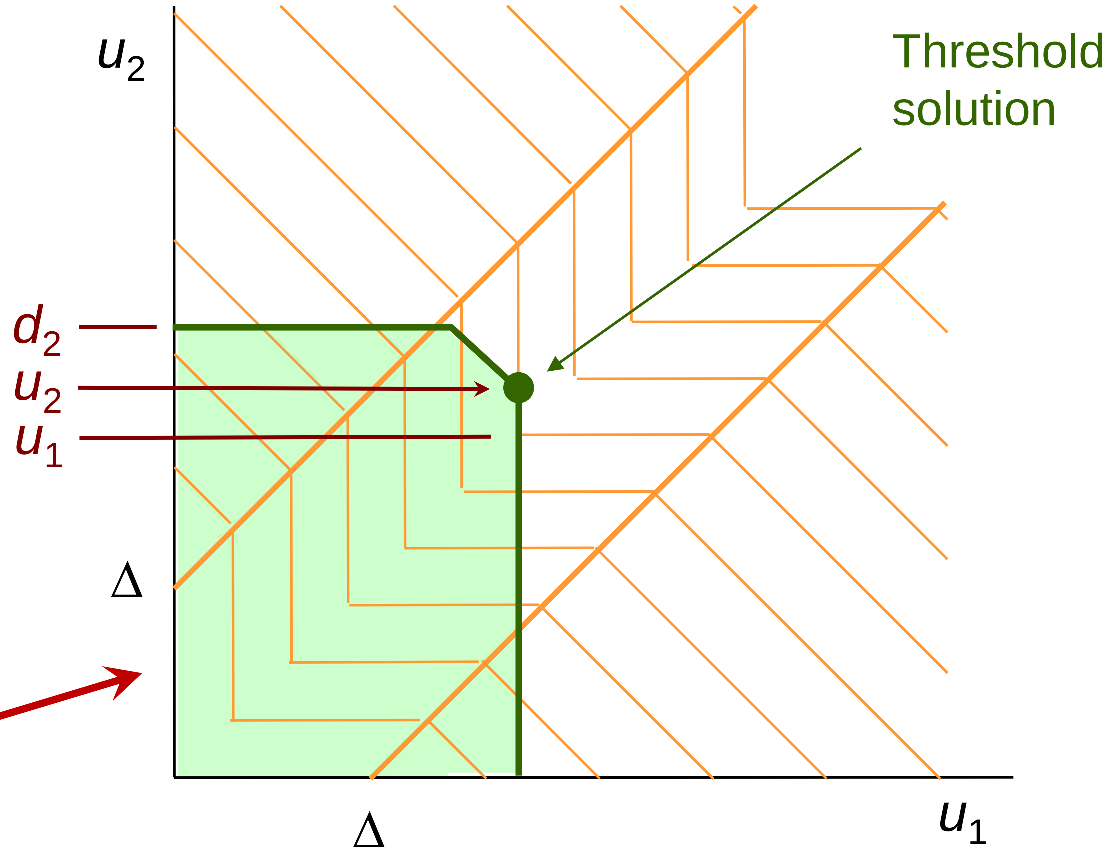

Here, **one** utility u_2 is **strictly between** upper bound d_2 and the smallest utility u_1.

Utility Threshold

Regionally decomposable?

- **No**
- 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

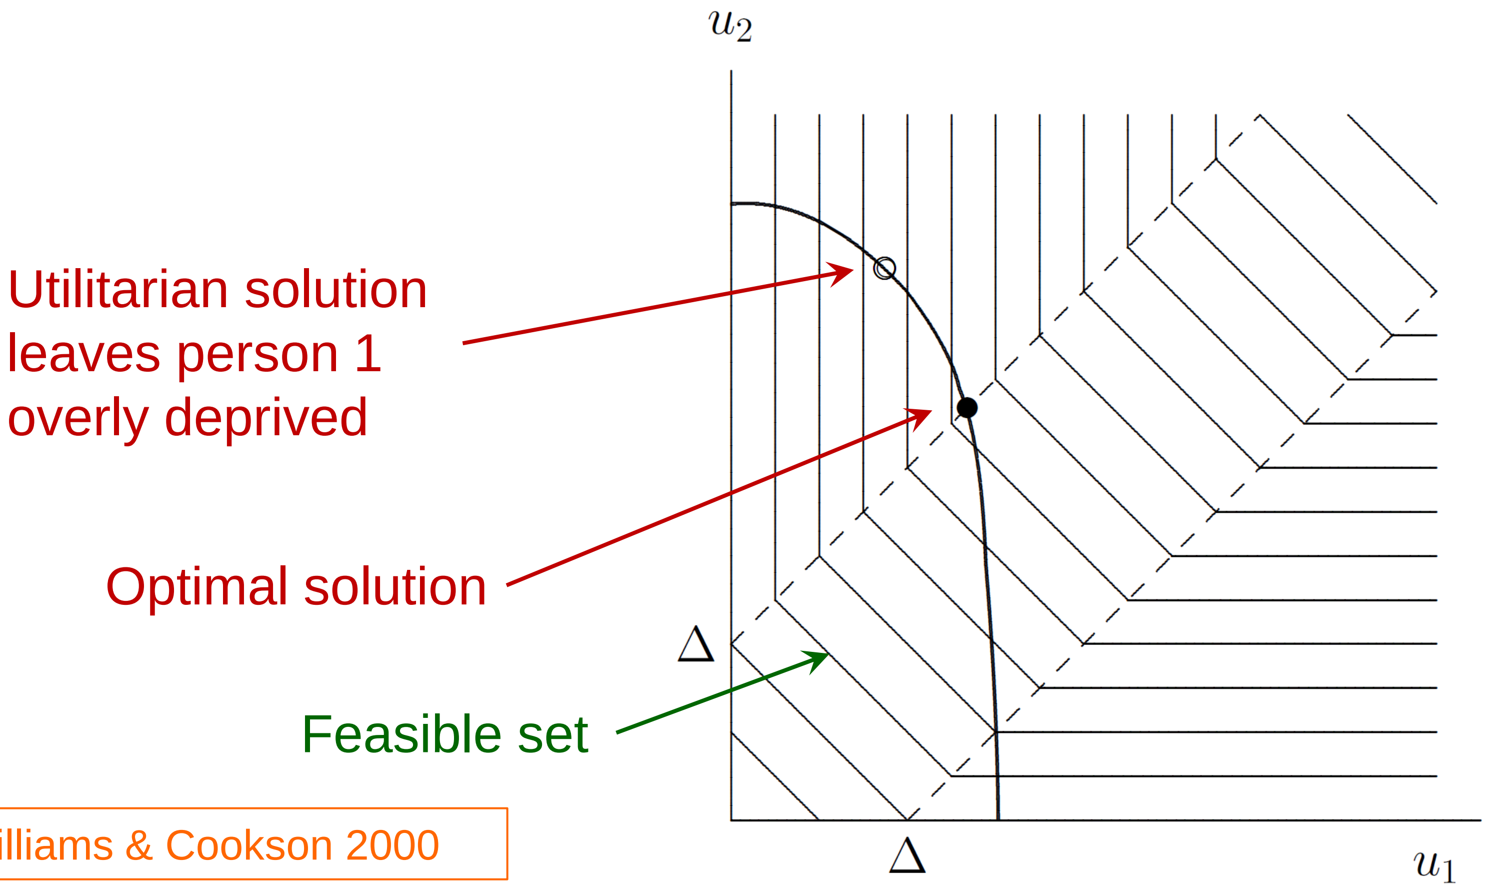

Williams & Cookson 2000

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

Equity Threshold

Generalization to *n* persons

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

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.
- 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.

Equity Threshold

Example:

Maximum equity threshold SWF subject to budget constraint

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

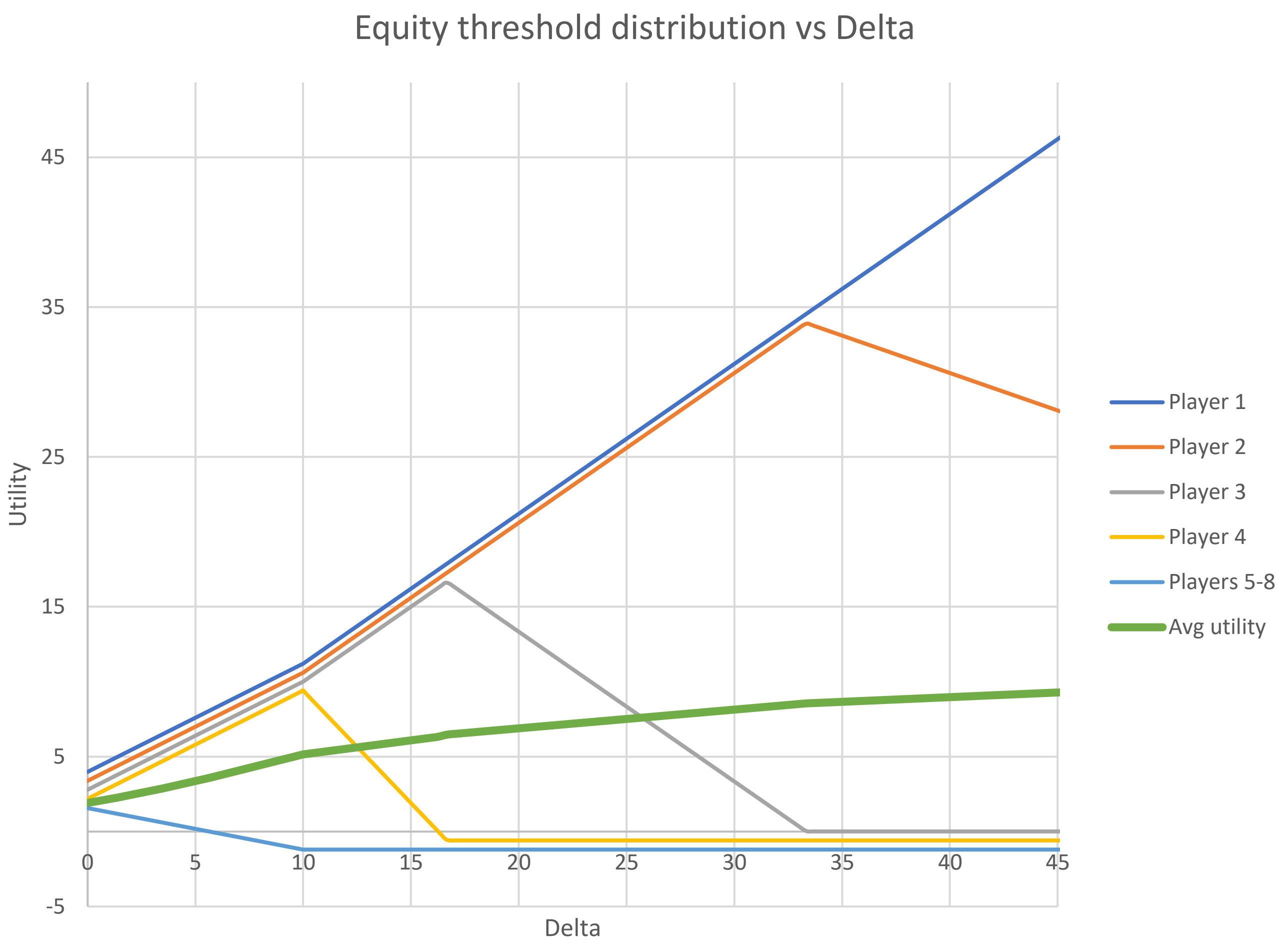

Equity Threshold

Regionally decomposable?

- **No**

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 $\frac{a_i}{\sum_j a_j}$ for any individual *i*.

Cross-subsidy rates:

- For **large** (more utilitarian) Δ, only the **one person** with a certain intermediate utility level benefits from the improvements of others (namely, those with greater efficiencies).
- For **small** (more egalitarian) Δ, all rates are $\frac{a_i}{\sum_j a_j}$

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

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

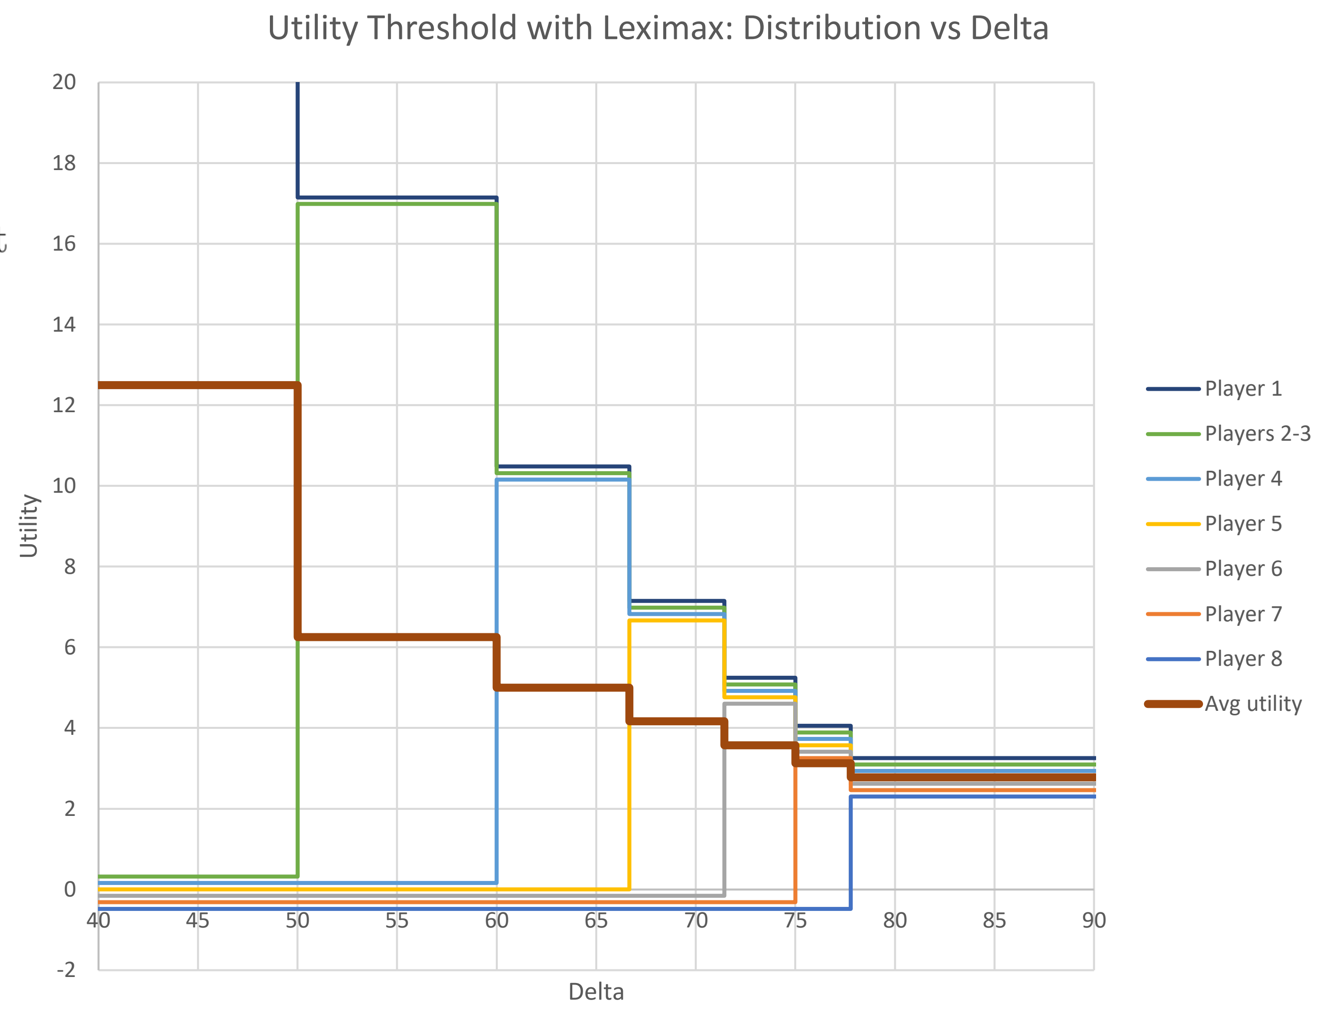

Utility Threshold with Leximax

Regionally decomposable?

- **No**

Incentive rate:

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

Cross-subsidy rates:

- Rates among individuals who receive positive utility are $\frac{a_i}{\sum_{j=1}^{m} a_j}$, others are **zero**.

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.

Cross-fertilization

- Research in **optimization** and **other fields** can be mutually beneficial.
- For example,

Linear programming and duality theory ↔ Economics and game theory

Cross-fertilization

- Research in **optimization** and **other fields** can be mutually beneficial.
- For example,

Linear programming and duality theory ⟷ Economics and game theory

- Potentially,

Post hoc analysis of social welfare optimization ⟷ Ethics and distributive justice


Questions or
comments?