Balancing Fairness and Efficiency in the Allocation of Indivisible Goods

Karen Frilya Celine (karen.celine@u.nus.edu)

National University of Singapore

Fair Division

- We allocate indivisible resources among a group of agents.
- How to ensure fairness and efficiency?
- Example: inheritance.

Fairness vs Efficiency

- The conflict between fairness and efficiency has been widely studied in economics (Le Grand, 1991) and fair division (Brams & King, 2005).

Motivating Example

	Good 1	Good 2	Good 3
Agent 1	40	30	30
Agent 2	30	40	30
Agent 3	0	50	50

	Good 1	Good 2	Good 3
Agent 1	40	30	30
Agent 2	30	40	30
Agent 3	0	50	50

- Allocation 1:
  - Efficient:
    - Maximum utilitarian welfare (MUW): 40 + 0 + 100 = 140.
- Allocation 2:
  - Fair:
    - Balanced: each agents gets the same number of goods.
    - Envy-free (EF): each agent prefers her own good the most.
    - Envy-free up to one good (EF1): follows from EF.
  - Efficient:
    - Maximum egalitarian welfare (MEW): min(40, 40, 50) = 40.
    - Maximum Nash welfare (MNW): 40 · 40 · 50 = 80,000.

Key Points

- "Fair" and "efficient" can be defined in different ways.
- Depending on the definition, fairness and efficiency can be conflicting or aligned.
- How do we measure the degree of compatibility between specific fairness and efficiency criteria?

The Egalitarian Price of Fairness

- The price of fairness (POF) quantifies the loss of social welfare in the worst case.
  - Social welfare can be defined as:
    - utilitarian welfare (sum of utilities over all agents),
    - egalitarian welfare (minimum utility over all agents),
    - etc.

Results

Property		Price of fairness: Egalitarian[1]	Price of fairness: Utilitarian[2]
EF1		$\Theta(n)$	$\Theta(\sqrt{n})$
Balanced		n	$\Theta(\sqrt{n})$
Round-robin		$\Theta(n)$	n
MNW	$(n=2)$	≈ 2	≈ 1.2
MNW	$(n \geq 3)$	∞	$\Theta(n)$
MUW		∞	1
MEW		1	$\Theta(n)$

[1]Celine et al., 2023

[2]Barman et al., 2020; Bei et al., 2019

- Egalitarian welfare is arguably fairer than utilitarian welfare.
  - Egalitarian welfare cares about the poorest agent.
- Egalitarian POF is surprisingly higher than utilitarian POF.

The Fairness of Additive Welfarist Rules

- Additive welfarist rule with function f chooses an allocation that maximizes

$$\sum_{i=1}^{n} f(u_i(A_i)).$$

- Examples:
  - Maximum utilitarian welfare (MUW): $f(x) = x$.
  - Maximum Nash welfare (MNW): $f(x) = \log x$.
- We study which additive welfarist rules guarantee EF1 in certain subclasses of instances.

Results

- Real-valued:
  - No restriction: MNW is the only additive welfarist rule guaranteeing EF1 (Caragiannis et al., 2019; Suksompong, 2023).
  - We extended this result to more restricted subclasses (Celine et al., 2025):
    - identical-good instances, two-value instances, and normalized instances with at least three agents.
- Integer-valued:
  - There are infinitely many other additive welfarist rules guaranteeing EF1 (Celine et al., 2025).
    - Example: $f(x) = \log(x + c)$ for any $c \in [0, 1]$.

Future Work

- Other settings: e.g. weighted fair division.
- Combining the two: compare different additive welfarist rules guaranteeing EF1 using a concept similar to the price of fairness.