Fairness in allocation problems

Ioannis Caragiannis

University of Patras

Seminar at GSSI, L'Aquila, 2-4 July, 2018

An ancient problem

- **Cake cutting**
  - Input: **agents** with different **preferences** for parts of the cake
  - Goal: **divide** the cake in a **fair** manner
- Mathematical formulations initiated by Steinhaus, Banach, & Knaster (1948)
- Basic algorithm/protocol: **cut-and-choose**

Cake cutting

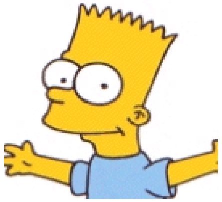

Cake cutting

Value of the agent for the piece of the cake at the left of the cut

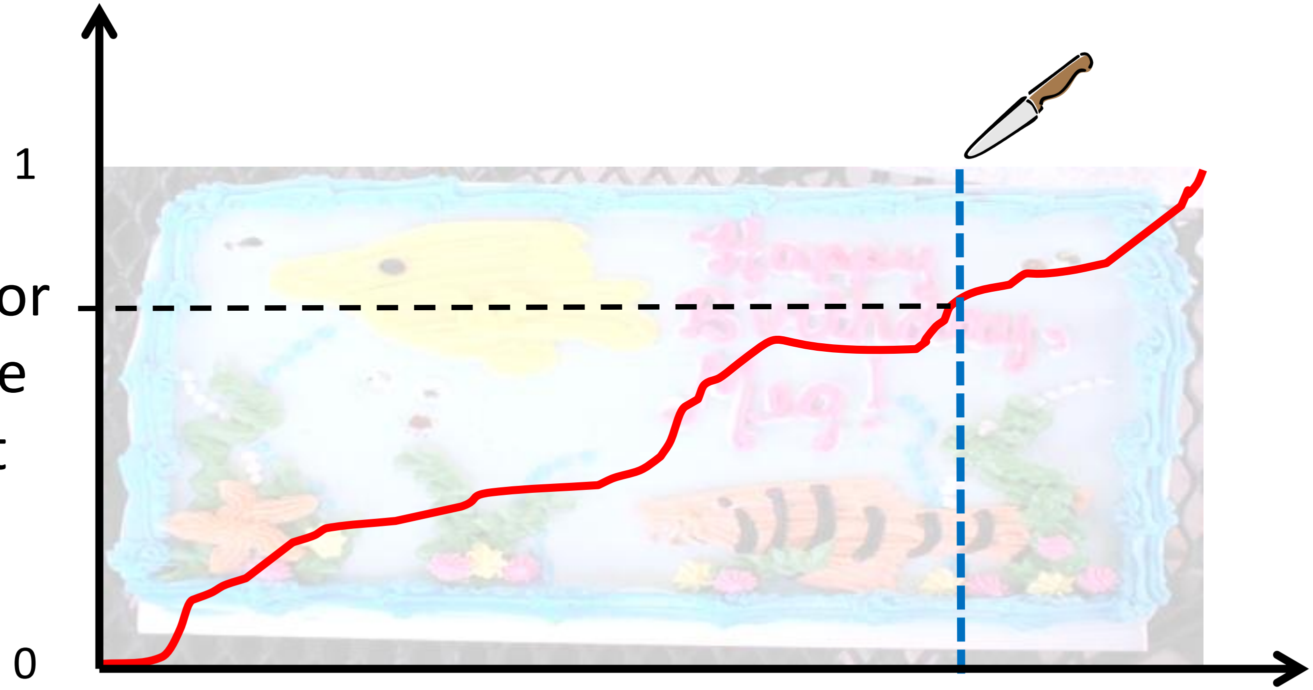

Cake cutting

- Cut-and-choose: Lisa **cuts**, Bart **chooses** first

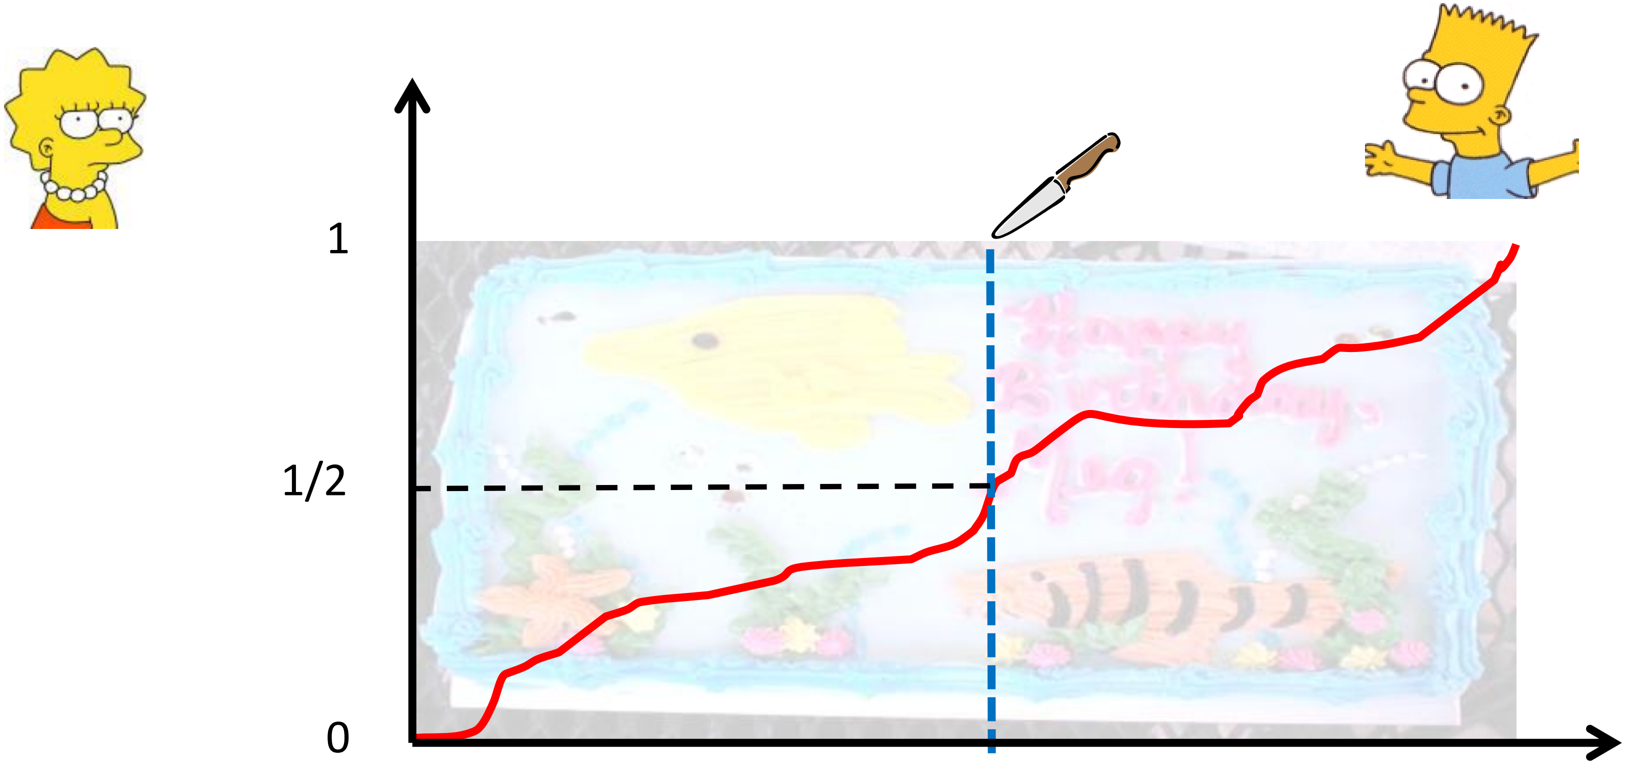

Allocations of goods

- **Indivisible** goods

- **Agents** with additive **valuations** for goods

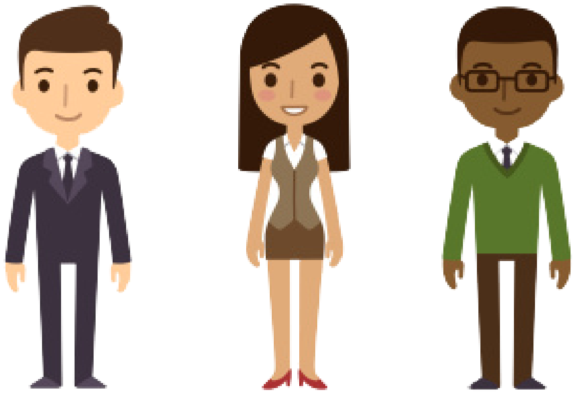

- Goal: **divide** the goods **fairly**

An allocation problem

					
	\$1000	\$200	\$600	\$100	\$100
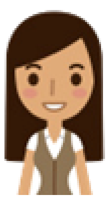	\$700	\$500	\$100	\$400	\$300
	\$500	\$700	\$400	\$200	\$200

An allocation problem




	$1000	$200	$600	**$100**	$100
	**$700**	**$500**	**$100**	$400	$300
	$500	$700	$400	$200	**$200**

Allocation problems: some history

- Ancient Egypt:
  - Land division around Nile (i.e., of the most fertile land)
- Ancient Greece:
  - Sponsorships in theatrical performances
- First references to cut-and-choose protocol
  - Theogony (Hesiod, 8^{th} century B.C.): run between Prometheus and Zeus
  - Bible: run between Abraham and Lot

Related implementations/tools

- http://www.spliddit.org
  - Algorithms for various classes of problems (allocations of goods, rent division, etc.)
  - Ariel Procaccia
- http://www.nyu.edu/projects/adjustedwinner/
  - Implementation of the "Adjusted Winner" algorithm for two agents
  - Steven Brams & Alan Taylor
- http://www.math.hmc.edu/~su/fairdivision/calc/
  - Algorithms for allocating goods
  - Francis Su

Basic notions

Formally …

- n **agents**
- A set of **goods** G
- Agent i has **valuation** $v_i(g)$ for good g
- Utilities are additive, i.e.,

$$v_i(S) = \sum_{g \in S} v_i(g)$$

- **Allocation**: a partition $A=(A_1, \ldots, A_n)$ of the goods in G

What does “fairly” mean?

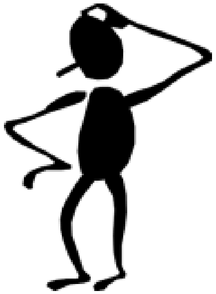

- **Fairness notions**
  - Envy freeness
  - Proportionality

What does "**fairly**" mean?

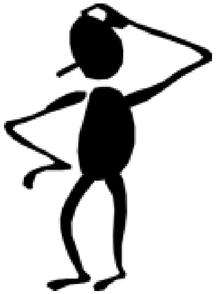

- **Fairness notions**
  - **Envy freeness**: every agent prefers her own bundle to the bundle of any other agent

$$\forall j, i, v_i(A_i) \geq v_i(A_j)$$

EF: an example

			ROLLS ROYCE		
	\$1000	\$200	\$600	\$100	\$100
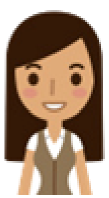	\$700	\$500	\$100	\$400	\$300
	\$500	\$700	\$400	\$200	\$200

EF: an example

	$\boxed{\textbf{\$1000}}$	\$200	\$600	\$100	\$100
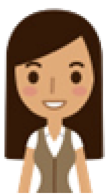	\$700	\$500	\$100	\$400	\$300
	\$500	\$700	\$400	\$200	\$200

EF: an example

					
	**$1000**	$200	$600	$100	$100
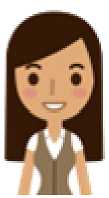	$700	$500	$100	**$400**	**$300**
	$500	**$700**	**$400**	$200	$200

What does “fairly” mean?

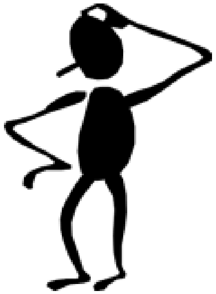

- Fairness notions
  - **Envy freeness**: every agent prefers her own bundle to the bundle of any other agent
  - **Proportionality**: every agent feels that she gets at least 1/n-th of the goods

$$\forall i, v_i(A_i) \geq \frac{1}{n} v_i(G)$$

Proportionality: an example

					
	$1200	$200	$300	$200	$100
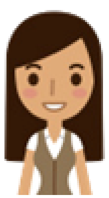	$800	$500	$200	$300	$200
	$800	$400	$400	$300	$100

Proportionality: an example

			ROLLS ROYCE		
	**$1200**	$200	$300	$200	$100
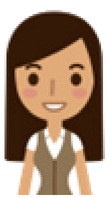	$800	**$500**	$200	$300	**$200**
	$800	$400	**$400**	**$300**	$100

What does “fairly” mean?

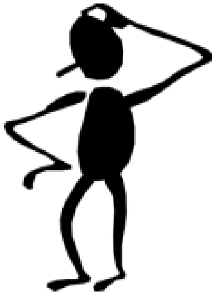

- **Fairness notions**
  - **Envy freeness**: every agent prefers her own bundle to the bundle of any other agent
  - **Proportionality**: every agent feels that she gets at least 1/n-th of the goods

Properties

- **Theorem**: EF implies Proportionality

Properties

- **Theorem**: EF implies Proportionality
- **Proof**: Since agent i does not envy any other agent, $\forall j \neq i, v_i(A_i) \geq v_i(A_j)$

Properties

- **Theorem**: EF implies Proportionality
- **Proof**: Since agent i does not envy any other agent, $\forall j \neq i, v_i(A_i) \geq v_i(A_j)$
  Trivially, $v_i(A_i) \geq v_i(A_i)$

Properties

- **Theorem**: EF implies Proportionality
- **Proof**: Since agent i does not envy any other agent, $\forall j \neq i, v_i(A_i) \geq v_i(A_j)$

  Trivially, $v_i(A_i) \geq v_i(A_i)$

  Summing all these n inequalities, we get

$$n \cdot v_i(A_i) \geq \sum_{j=1}^{n} v_i(A_j) = v_i(G)$$

Properties

- **Theorem**: EF implies Proportionality
- **Proof**: Since agent i does not envy any other agent, $\forall j \neq i, v_i(A_i) \geq v_i(A_j)$

  Trivially, $v_i(A_i) \geq v_i(A_i)$

  Summing all these n inequalities, we get

$$n \cdot v_i(A_i) \geq \sum_{j=1}^{n} v_i(A_j) = v_i(G)$$

and, equivalently, $\boxed{v_i(A_i) \geq \frac{1}{n} v_i(G)}$

Properties

- **Theorem**: For 2 agents, Proportionality is equivalent to EF

Properties

- **Theorem**: For 2 agents, Proportionality is equivalent to EF
- **Proof**: Since $v_1(A_1) \geq v_1(G)/2$, it must also be $v_1(A_2) \leq v_1(G)/2$, i.e., $v_1(A_1) \geq v_1(A_2)$.

Proportionality may not imply EF for more than two agents

				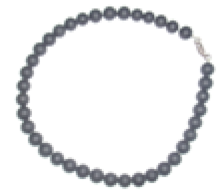	
	\$800	\$300	\$300	\$300	\$300
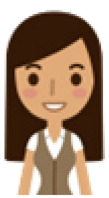	\$800	\$500	\$200	\$300	\$200
	\$800	\$400	\$400	\$300	\$100

Proportionality may not imply EF for more than two agents

				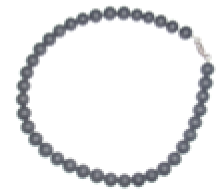	
	**$800**	$300	$300	$300	$300
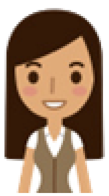	$800	**$500**	$200	$300	**$200**
	$800	$400	**$400**	**$300**	$100

Fairness vs. Efficiency

A motivating example

goods

agents

				
	\$3	\$0	\$5	**\$12**
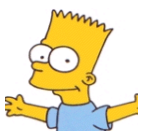	**\$0**	**\$8**	**\$8**	\$4

 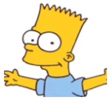

allocation ({  }, {    } is EF

A motivating example

goods

agents

				
	**$3**	$0	$5	**$12**
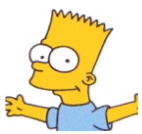	$0	**$8**	**$8**	$4

 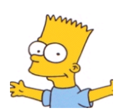

allocation ({}, {  } is EF

allocation ({ }, { } is EF and, in a sense, better!

Efficiency

- **Economic** efficiency
  - Pareto-optimality
  - Social welfare maximization
- **Computational** efficiency
  - Polynomial-time computation
  - Low query complexity

Efficiency

- **Economic** efficiency
  - Pareto-optimality
  - Social welfare maximization
- **Computational** efficiency
  - Polynomial-time computation
  - Low query complexity

a property of allocations

a property of allocation algorithms/protocols

Warming up: Pareto-optimality vs fairness

- Definition: an allocation $A = (A_1, A_2, \ldots, A_n)$ is called **Pareto-optimal** if there is no allocation $B = (B_1, B_2, \ldots, B_n)$ such that $v_i(B_i) \geq v_i(A_i)$ for every agent i and $v_{i'}(B_{i'}) > v_{i'}(A_{i'})$ for some agent i'

- Informally: there is no allocation in which **all agents are at least as happy** and **some agent is strictly happier**

Envy-freeness vs. Pareto-optimality

goods

agents

				
	$3	$0	$5	**$12**
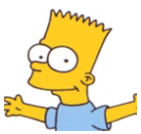	**$0**	**$8**	**$8**	$4

- Observation: In a Pareto-optimal allocation, agent  does not get  and agent 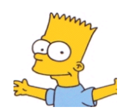 does not get 

Envy-freeness vs. Pareto-optimality

goods

agents				
	\$3	\$0	\$5	**\$12**
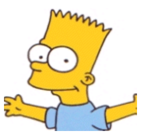	**\$0**	**\$8**	**\$8**	\$4

- Observation: In a Pareto-optimal allocation, agent  does not get  and agent 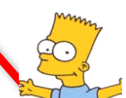 does not get 

An envy-free allocation that is not Pareto-optimal

Envy-freeness vs. Pareto-optimality

goods

agents

				
	**$3**	$0	**$5**	**$12**
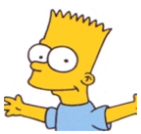	$0	**$8**	$8	$4

	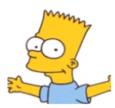	**PO**	**EF**
		**?**	**?**

Envy-freeness vs. Pareto-optimality

goods

agents

				
	**$3**	$0	**$5**	**$12**
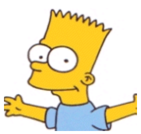	$0	**$8**	$8	$4

	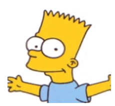	**PO**	**EF**
		**YES**	**NO**

Envy-freeness vs. Pareto-optimality

goods

agents

				
	**$3**	$0	**$5**	$12
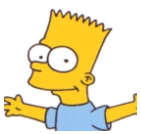	$0	**$8**	$8	**$4**

	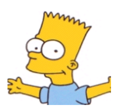	**PO**	**EF**
		**YES**	**NO**
		**?**	**?**

Envy-freeness vs. Pareto-optimality

goods

agents

				
	**$3**	$0	**$5**	$12
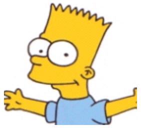	$0	**$8**	$8	**$4**

	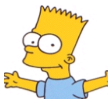	**PO**	**EF**
  		YES	NO
 	 	NO	NO

Envy-freeness vs. Pareto-optimality

goods

agents

				
	**$3**	$0	$5	**$12**
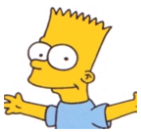	$0	**$8**	**$8**	$4

	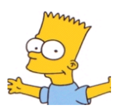	**PO**	**EF**
		**YES**	**NO**
		**NO**	**NO**
		**?**	**?**

Envy-freeness vs. Pareto-optimality

goods

agents

				
	**$3**	$0	$5	**$12**
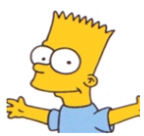	$0	**$8**	**$8**	$4

	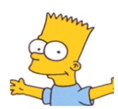	**PO**	**EF**
		YES	NO
		NO	NO
		YES	YES

Envy-freeness vs. Pareto-optimality

goods

agents				
	**$3**	$0	$5	$12
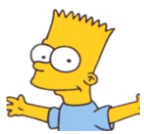	$0	**$8**	**$8**	**$4**

	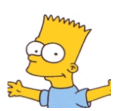	**PO**	**EF**
		YES	NO
		NO	NO
		YES	YES
		**?**	**?**

Envy-freeness vs. Pareto-optimality

goods

agents

				
	**$3**	$0	$5	$12
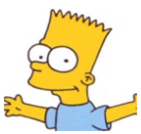	$0	**$8**	**$8**	**$4**

	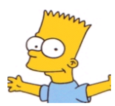	**PO**	**EF**
		YES	NO
		NO	NO
		YES	YES
		YES	NO

Envy-freeness vs. Pareto-optimality

- **Theorem**: Consider an allocation instance with 2 agents that has at least one EF allocation. Then, **there is an EF allocation that is simultaneously PO**.

Envy-freeness vs. Pareto-optimality

- **Theorem**: Consider an allocation instance with 2 agents that has at least one EF allocation. Then, **there is an EF allocation that is simultaneously PO**.
- **Proof**. Sort the EF allocations in lexicographic order of agents' valuations. The first allocation in this order is clearly PO.

Envy-freeness vs. Pareto-optimality

- **Theorem**: Consider an allocation instance with 2 agents that has at least one EF allocation. Then, **there is an EF allocation that is simultaneously PO**.
- **Proof**. Sort the EF allocations in lexicographic order of agents' valuations. The first allocation in this order is clearly PO.
- **Question**: What about 3-agent instances?
- **Question**: What about Proportionality vs PO?
  - See Bouveret & Lemaitre (2016)

Social welfare

- **Social welfare** is a measure of global value of an allocation $A = (A_1, \ldots, A_n)$
- **Utilitarian social welfare** of an allocation A:
  - the total utility of the agents for the goods allocated to them in A, i.e., $\mathrm{uSW(A)} = \sum_{\mathrm{i \in N}} \mathrm{v_i(A_i)}$
- **Egalitarian social welfare**: $\mathrm{eSW(A)} = \min_{\mathrm{i \in N}} \mathrm{v_i(A_i)}$
- **Nash social welfare**: $\mathrm{nSW(A)} = \prod_{\mathrm{i \in N}} \mathrm{v_i(A_i)}$

An example

- SW-maximizing allocations?

goods

agents

				
	15	0	40	45
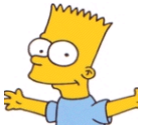	0	30	30	40

An example

- SW-maximizing allocations?

goods

agents

				
	15	0	40	45
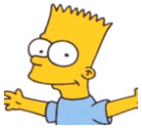	0	30	30	40

		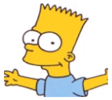
uSW	?	?
eSW	?	?
nSW	?	?

An example

- SW-maximizing allocations?

good

agents

				
	**15**	0	**40**	**45**
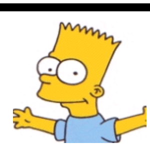	0	**30**	30	40

		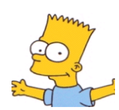
uSW		
eSW	**?**	**?**
nSW	**?**	**?**

Give each good to the agent who values it the most
uSW=130

An example

- SW-maximizing allocations?

goods

agents

				
	**15**	0	40	**45**
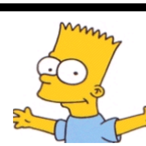	0	**30**	**30**	40

		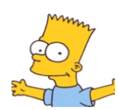
uSW		
eSW		
nSW	**?**	**?**

eSW=60

An example

- SW-maximizing allocations?

goods

agents

				
	**15**	0	**40**	45
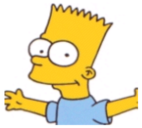	0	**30**	30	**40**

		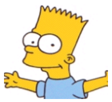
uSW		
eSW		
nSW		

nSW=3850

An example

- SW-maximizing allocations?

goods

agents

				
	15	0	40	45
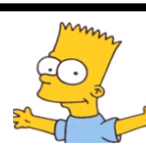	0	30	30	40

		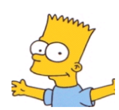	**EF**
uSW			**?**
eSW			**?**
nSW			**?**

An example

- SW-maximizing allocations?

goods

agents

				
	15	0	40	45
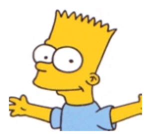	0	30	30	40

		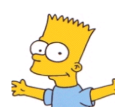	**EF**
uSW			**NO**
eSW			**YES**
nSW			**YES**

Price of fairness

- **Price of fairness** (in general)
  - how far from its maximum value can the social welfare of the best fair allocation be?
- More specifically:
  - Which definition of social welfare to use?
  - Which fairness notion to use?
- Answer:
  - **Any combination of them**

Price of fairness

- How **large** the social welfare of a **fair** allocation can be?
  - C., Kaklamanis, Kanellopoulos, and Kyropoulou (2012)

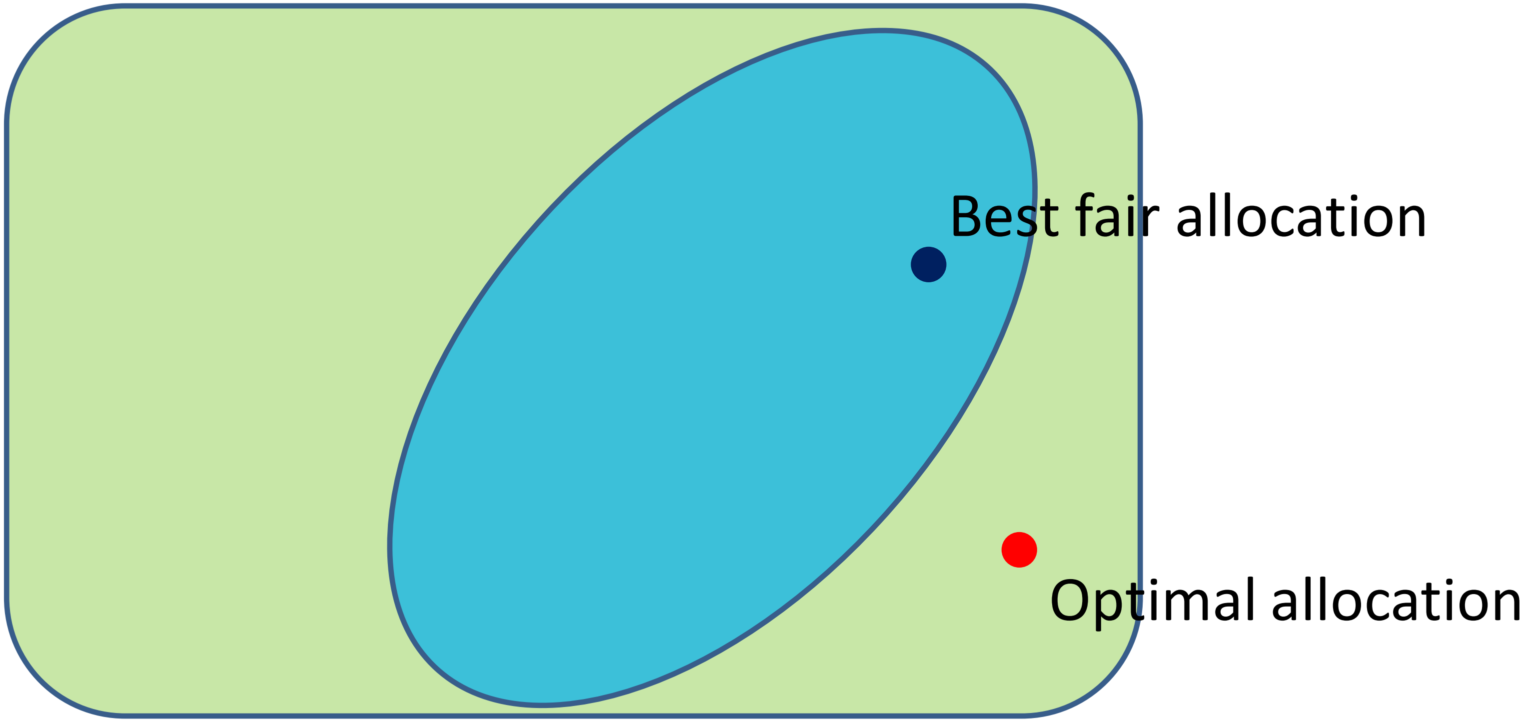

Price of fairness

- How **large** the social welfare of a **fair** allocation can be?
  - C., Kaklamanis, Kanellopoulos, and Kyropoulou (2012)

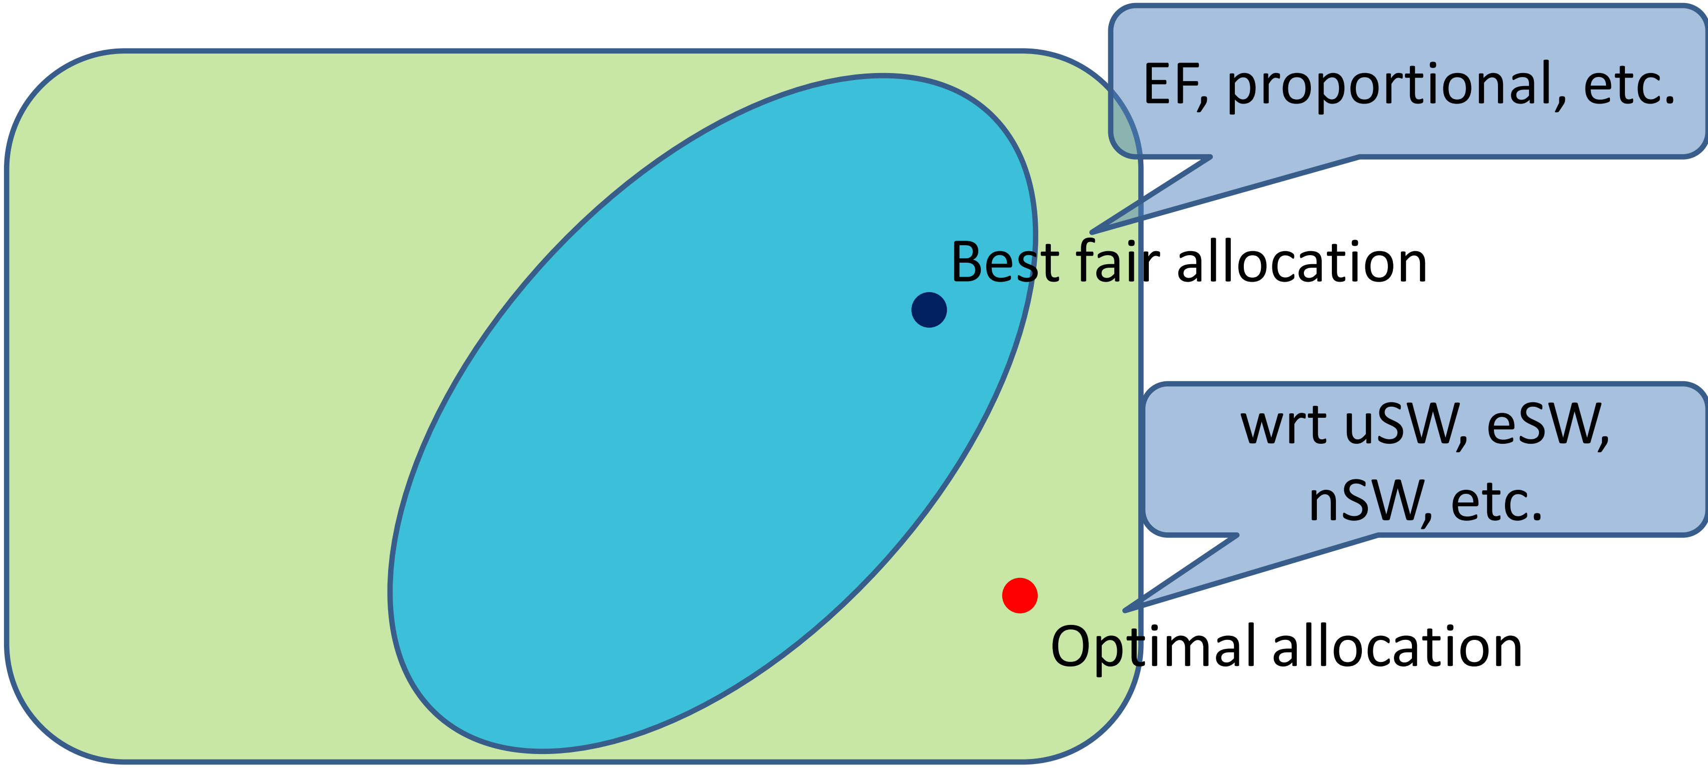

PoP & uSW for 2 agents

- **Theorem**: The price of proportionality with respect to the utilitarian social welfare for 2-agent instances is 3/2 (**tight bound**)

PoP & uSW for 2 agents

- **Theorem**: The price of proportionality with respect to the utilitarian social welfare for 2-agent instances is **at least 3/2**.

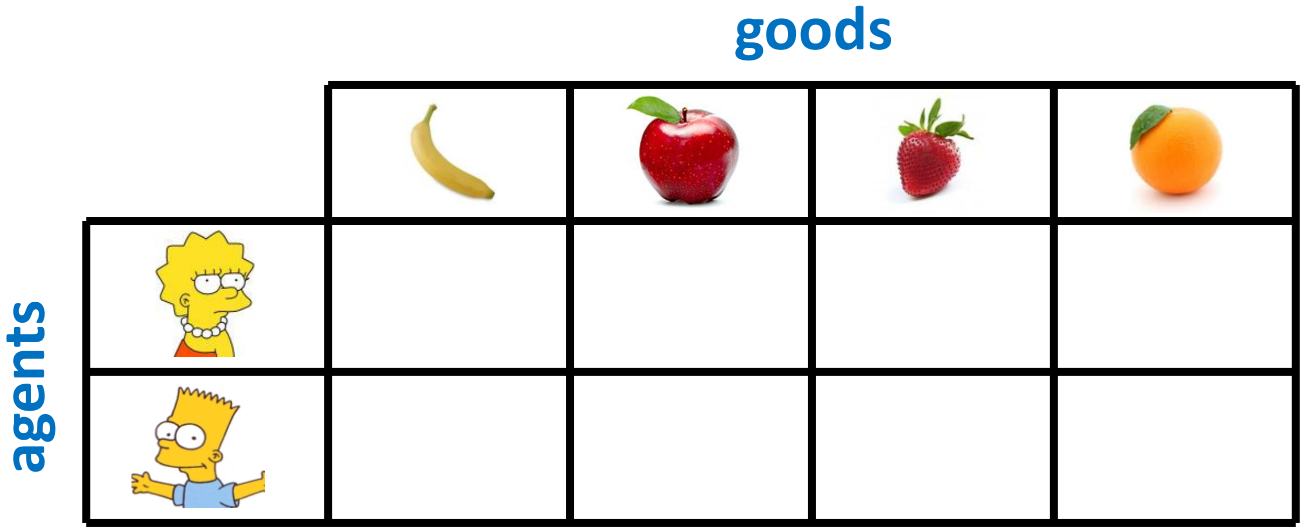

PoP & uSW for 2 agents

- **Theorem**: The price of proportionality with respect to the utilitarian social welfare for 2-agent instances is **at least 3/2**.

goods

agents

				
	0.5-ε	0.5-ε	ε	ε
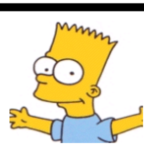	0.25+ε	0.25+ε	0.25-ε	0.25-ε

PoP & uSW for 2 agents

- **Theorem**: The price of proportionality with respect to the utilitarian social welfare for 2-agent instances is **at least 3/2**.

goods

agents				
	**0.5-ε**	**0.5-ε**	ε	ε
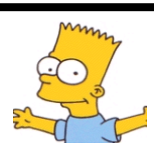	0.25+ε	0.25+ε	**0.25-ε**	**0.25-ε**

- Optimal allocation (uSW ≈ 1.5)

PoP & uSW for 2 agents

- **Theorem**: The price of proportionality with respect to the utilitarian social welfare for 2-agent instances is **at least 3/2**.

goods

agents

				
	0.5-ε	0.5-ε	ε	ε
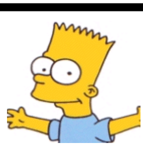	0.25+ε	0.25+ε	0.25-ε	0.25-ε

- Optimal allocation (uSW ≈ 1.5)

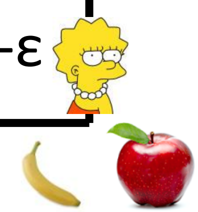

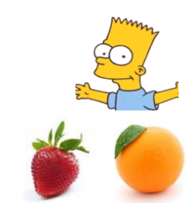

- Best proportional allocation

? ?

PoP & uSW for 2 agents

- **Theorem**: The price of proportionality with respect to the utilitarian social welfare for 2-agent instances is **at least 3/2**.

goods

agents

				
	**0.5-ε**	0.5-ε	ε	**ε**
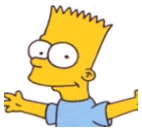	0.25+ε	**0.25+ε**	**0.25-ε**	0.25-ε

- Optimal allocation (uSW ≈ 1.5)
- Any prop. allocation has uSW ≈ 1

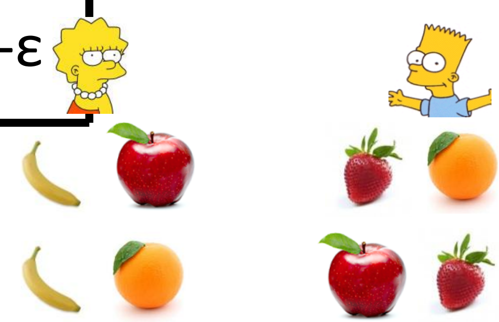

PoP & uSW for 2 agents

- **Theorem**: The price of proportionality with respect to the utilitarian social welfare for 2-agent instances is **at most 3/2**.

PoP & uSW for 2 agents

- **Theorem**: The price of proportionality with respect to the utilitarian social welfare for 2-agent instances is **at most 3/2**.
- **Proof**: If the uSW-maximizing allocation is proportional, then PoP=1.

PoP & uSW for 2 agents

- **Theorem**: The price of proportionality with respect to the utilitarian social welfare for 2-agent instances is **at most 3/2**.
- **Proof**: If the uSW-maximizing allocation is proportional, then PoP=1. So, assume otherwise. Then, some agent has utility less than 1/2 for a total of at most 3/2. In any proportional allocation, uSW=1.

PoP & uSW for 2 agents

- **Theorem**: The price of proportionality with respect to the utilitarian social welfare for 2-agent instances is **at most 3/2**.
- **Proof**: If the uSW-maximizing allocation is proportional, then PoP=1. So, assume otherwise. Then, some agent has utility less than 1/2 for a total of at most 3/2. In any proportional allocation, uSW=1.
- **Question:** PoP/PoEF wrt uSW for many agents?

Computational (in)efficiency

- Computing a proportional/EF allocation is **NP-hard**
- Reduction from **Partition**:
  - Partition instance: given items with weights w_1, w_2, ..., w_m, decide whether they can be partitioned into two sets with equal total weight
  - Proportionality/EF instance: A good for each item; 2 agents with identical valuation of w_i for good i

Summary

- What have we covered today?
  - Problem definition
  - Properties of EF/proportionality
  - Pareto-optimality
  - Social welfare
  - Fairness vs efficiency

Further reading