

Fair Division: Criteria and Algorithms for Allocating Indivisible Resources

Vangelis Markakis

Athens University of Economics and Business (AUEB)

Dept. of Informatics

Outline

- **Part 1:**
 - Basic models, and definitions of fairness criteria
- **Part 2:**
 - Algorithms for relaxed versions of envy-freeness
- **Part 3:**
 - Relations between (approximate) relaxations of envy-freeness

Cake-cutting problems

Input:

- A set of resources
- A set of agents, with possibly different preferences



likes everything equally



is allergic to strawberry



loves chocolate



www.shutterstock.com · 114352723

Cake-cutting problems

Input:

- A set of resources
- A set of agents, with possibly different preferences

Goal: Divide the resources among the agents in a **fair** manner

Empirically: since ancient times

Mathematical formulations: Initiated by
[Steinhaus, Banach, Knaster '48]



www.shutterstock.com · 114352723

Application domains

- Computer Science:

- Resource allocation in computer systems
- Job scheduling in operating systems
- Task assignment in robot teams

- Operations research:

- Load balancing problems, e.g. flight assignments to aircrafts

- Microeconomics:

- Mechanism design for allocations of goods

- Political sciences:

- Design of voting rules (for committee elections, seat allocations)

Available implementations: <http://www.spliddit.org> [Jonathan Goldman, Ariel Procaccia, CMU]

- Algorithms for various problems (rent division, division of goods, etc)

Modeling fair division problems: the discrete setting

- This talk: focus on indivisible items
- Resources = a set of **indivisible** goods $M = \{1, 2, \dots, m\}$
- Set of agents: $N = \{1, 2, \dots, n\}$
- An allocation of M is a partition $S = (S_1, S_2, \dots, S_n), S_i \subseteq M$
 - $\bigcup_i S_i = M$ and $S_i \cap S_j = \emptyset$
- Preferences: modeled by a valuation function for each agent



Valuation functions

All valuations we consider satisfy:

- $v_i(\emptyset) = 0$ (normalization)
- $v_i(S) \leq v_i(T)$, for any $S \subseteq T$ (monotonicity)

Special cases of interest:

- Additive: $v_i(S \cup T) = v_i(S) + v_i(T)$, for any disjoint sets S, T
 - Assumed in the majority of the literature
 - Suffices to specify v_{ij} for any good j : $v_i(S) = \sum_{j \in S} v_{ij}$, for any $S \subseteq M$
- Additive with identical rankings on the value of the goods
- Identical agents: Same valuation function for everyone
- Subadditive: $v_i(S \cup T) \leq v_i(S) + v_i(T)$, for any $S, T \subseteq M$
- Submodular: $v_i(S \cup \{j\}) - v_i(S) \geq v_i(T \cup \{j\}) - v_i(T)$, for any $S \subseteq T$, and $j \notin T$

The discrete setting

Example with additive valuations

						
Charlie		35	5	25	0	35
Franklin		30	40	35	5	40
Marcie		30	20	40	30	0

Part 1: solution concepts for indivisible items

Solution Concepts

1. Proportionality

An allocation (S_1, S_2, \dots, S_n) is **proportional**, if for every agent i ,

$$v_i(S_i) \geq 1/n \cdot v_i(M)$$

Historically, the first concept studied in the literature

[Steinhaus, Banach, Knaster '48]

Solution Concepts

2. Envy-freeness

An allocation (S_1, S_2, \dots, S_n) is **envy-free**, if $v_i(S_i) \geq v_i(S_j)$ for any pair of players i and j

- Suggested as a math puzzle in [Gamow, Stern '58]
- More formally discussed in [Foley '67, Varian '74]

A stronger concept than proportionality (as long as valuations are subadditive)

Some issues

- These definitions are “too strong” for indivisible goods
- No guarantee of existence
- No guarantees of reasonable approximations either
- More appropriate for the continuous setting (existence is always guaranteed there)
- Need to explore **relaxed versions of fairness**

Solution Concepts

3. Envy-freeness up to 1 good (EF1)

An allocation (S_1, S_2, \dots, S_n) satisfies **EF1**, if for any pair of agents i, j , **there exists** a good $g \in S_j$ such that $v_i(S_i) \geq v_i(S_j \setminus \{g\})$

- i.e., for any player who may envy agent j , there exists an item to remove from S_j and eliminate envy
- Defined by **[Budish '11]**

Solution Concepts

4. Envy-freeness up to any good (EFX)

An allocation (S_1, S_2, \dots, S_n) satisfies **EFX**, if for any pair of agents i, j , and **any** good $g \in S_j$, we have $v_i(S_i) \geq v_i(S_j \setminus \{g\})$

- Removing any item from each player's bundle eliminates envy from other players
- Defined by [Caragiannis et al. '16]

Fact: Envy-freeness \Rightarrow EFX \Rightarrow EF1

Solution Concepts

5. Maximin Share Allocations (MMS)

A thought experiment:

- Suppose we run the cut-and-choose protocol for n agents.
- Say agent i is given the chance to suggest a partition of the goods into n bundles
- The rest of the agents then choose a bundle and i chooses last
- **Worst case for i :** he is left with his least desirable bundle

Solution Concepts

- Given n agents and $S \subseteq M$, the **n -maximin share of i w.r.t. M** is

$$\mu_i(n, S) = \max_{T \in \Pi(S)} \min_{T_j \in T} v_i(T_j)$$

- max is over all possible partitions of S into n bundles
 - min is over all bundles of a partition $T = (T_1, T_2, \dots, T_n)$
- For $S = M$, $\mu_i := \mu_i(n, M)$ = the **maximin share of i**

Solution Concepts

An allocation (S_1, S_2, \dots, S_n) is a **maximin share (MMS) allocation** if for every agent i , $v_i(S_i) \geq \mu_i$

- Introduced by [Budish '11] (motivated by fairness properties of approximate *competitive equilibrium of equal incomes*)
- **Fact:** Proportionality \Rightarrow MMS

Maximin shares



35

5

25

0

35

$$\mu_1 = 30$$



30

40

35

5

40

$$\mu_2 = 40$$



30

20

40

30

0

$$\mu_3 = 30$$

MMS vs EF1 (and vs EFX)

How do MMS allocations compare to EF1 and EFX?

- They are generally incomparable
- There exist EFX allocations that are not MMS allocations
- There exist MMS allocations that do not satisfy EF1 (hence not EFX either)

MMS vs EF1 (and vs EFX)

						
	35	5	25	0	35	$\mu_1 = 30$
	30	40	35	5	40	$\mu_2 = 40$
	30	20	40	30	0	$\mu_3 = 30$

A MMS allocation that does not satisfy EF1

- Charlie envies Franklin even after removing any item from Franklin's bundle

Solution Concepts

6. Pairwise Maximin Share Allocations (PMMS)

- Motivated by the definition of MMS allocations but moving towards envy-freeness
- Request an MMS-type guarantee but for pairs of agents
- Consider an allocation $S = (S_1, S_2, \dots, S_n)$
- We look at every pair of agents, say i, j
- Collectively they receive $S_i \cup S_j$ under S
- We demand that both i and j receive at least their MMS share for the items in $S_i \cup S_j$

Solution Concepts

An allocation (S_1, S_2, \dots, S_n) is a **pairwise maximin share (PMMS) allocation** if for every pair of agents i, j , $v_i(S_i) \geq \mu_i(2, S_i \cup S_j)$

- Introduced by [Caragiannis et al. '16]
- **Fact 1:** PMMS \Rightarrow EFX
- **Fact 2:** PMMS and MMS are incomparable
 - There exist PMMS allocations that are not MMS allocations and vice versa

Solution Concepts

7. Groupwise Maximin Share Allocations (GMMS)

- Even further: A generalization of PMMS allocations [Barman et al. '18]
- Apply the idea of PMMS, not only for pairs but for any subset of agents
- An allocation (S_1, S_2, \dots, S_n) is a **groupwise maximin share (GMMS) allocation** if for every i , and every $R \subseteq N$ with $i \in R$

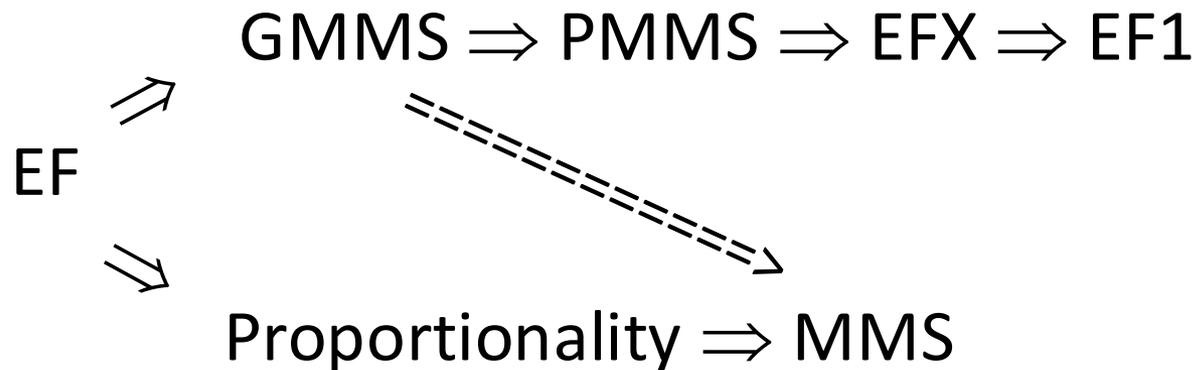
$$v_i(S_i) \geq \mu_i(|R|, \cup_{j \in R} S_j)$$

- **Facts:** GMMS \Rightarrow PMMS, MMS, and EFX
 - The strongest of all relaxations

Relations between fairness criteria

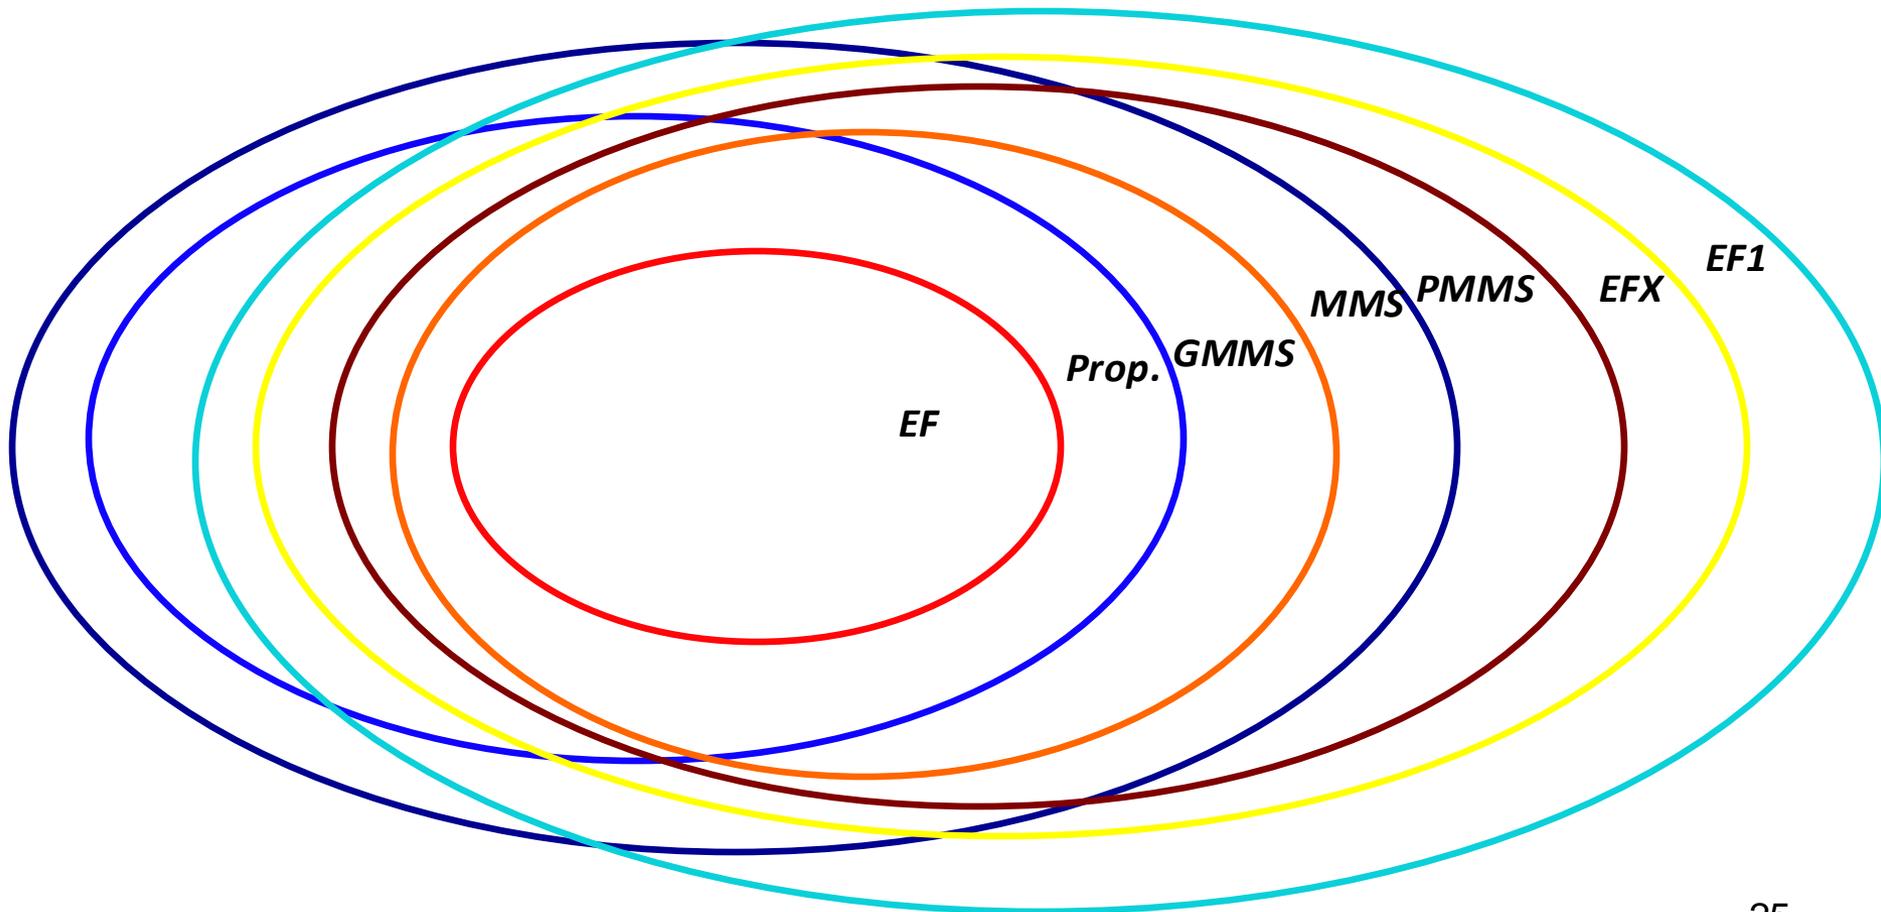
For additive valuation functions

- Upper part holds for general monotone valuations



Relations between fairness criteria

Pictorially:



Approximate versions of solution concepts

- We are also interested in approximate versions of all the relaxations
- How can we define approximation here?
 - We focus on multiplicative approximations
- **Example 1:** An allocation (S_1, S_2, \dots, S_n) is an α -EFX allocation, for $\alpha \in [0, 1]$, if for any players i and j , and **any** good $g \in S_j$, we have

$$v_i(S_i) \geq \alpha v_i(S_j \setminus \{g\})$$

- Similarly we can define α -EF1
- **Example 2:** An allocation (S_1, S_2, \dots, S_n) is an α -MMS allocation if for every agent i , $v_i(S_i) \geq \alpha \mu_i$
 - Similarly we can define α -PMMS and α -GMMS

Part 2: Existence and computation

EF1

Existence of EF1 allocations?

Theorem: For monotone valuation functions, EF1 allocations always exist and can be computed in polynomial time

- For additive valuations, a simple greedy round-robin algorithm suffices:
 - Select an arbitrary ordering of the agents
 - Each agent in his turn picks his favorite good among the currently available ones
- For non-additive valuations, more insightful to look at a graph-theoretic representation

EF1 for General Valuations

- Let S be an allocation (not necessarily of the whole set M)
- The **envy-graph** of S :
 - Nodes = agents
 - Directed edge (i, j) if i envies j under S
- How does this help?
 - Agents that nobody envies: **in-degree = 0**
 - Agents that at least one other person envies them: **in-degree ≥ 1**

EF1 for General Valuations

- An iterative algorithm till we reach a complete allocation
 - Suppose we have built a partial allocation that is EF1
 - If there exists a node with in-degree 0: give to this agent one of the currently unallocated goods
 - If this is not the case:
 - The graph has cycles
 - Start removing them one by one, by exchanging bundles, as dictated by each cycle
 - each node in the cycle gets the bundle of the neighbor he is pointing to
 - Until we have a node with in-degree 0

EF1 for General Valuations

Algorithm 2 – The Cycle Elimination Algorithm

[Lipton, Markakis, Mossel, Saberi '04]

- Fix an arbitrary ordering of the goods, say, $1, 2, \dots, m$
- At iteration i :
 - Find a node j with in-degree 0 (by possibly eliminating cycles from the envy-graph)
 - Give good i to agent j

Proof of correctness:

- Removing cycles terminates fast
 - Number of edges decreases after each cycle is gone
- At every step, we create envy only for the last item
- The allocation remains EF1 throughout the algorithm

Conclusion: EF1 is an “easy” concept to satisfy

EFX

Existence of EFX allocations?

- for $n = 2$
 - YES (for general valuations)
 - Variant of cut-and-choose
- for $n=3$
 - YES [Chaudhury, Garg, Melhorn '20]
- for $n \geq 4$
 - Great open problem!
 - Guaranteed to exist only for agents with identical valuations or identical rankings over the goods
- Existence of approximately EFX allocations?
 - Recent positive results, best possible approximation not yet resolved

Algorithmic and complexity results for EFX allocations

[Plaut, Roughgarden '18]:

Separation between general and additive valuations

Theorem:

1. exponential lower bound even for 2 agents with **identical submodular** valuations
2. Polynomial time algorithm for 2 agents and **arbitrary additive** valuations
3. Polynomial time algorithm for any n , and **additive valuations with identical rankings**
 - All agents have the same ordering on the value of the goods

Algorithmic and complexity results for EFX allocations

Algorithm for **additive valuations with identical rankings**:

- Cycle elimination comes to rescue
- Run the cycle elimination algorithm
 - but first order the goods in **decreasing order** of value
- At every step of the algorithm we allocate the next item to an agent no-one envies
- Envy we create is only for the item at the current iteration
- But this has lower value than all the previous goods for all agents, since they have identical rankings
- Hence the allocation remains EFX throughout the algorithm

Approximate EFX allocations

[Plaut, Roughgarden '18]:

- Existence of $\frac{1}{2}$ -EFX for arbitrary additive valuations
 - Extends also to subadditive valuations
- Again making use of the cycle elimination algorithm
- But with the option of de-allocating some items that were given to a player in previous steps
- Worst case exponential complexity

More recently:

[Chan et al. '19]:

- $\frac{1}{2}$ -EFX in polynomial time

Approximate EFX allocations

[Amanatidis, Markakis, Ntokos '20]:

- 0.618-EFX Poly-time algorithm for $(0.618 = 1/\phi)$
- A careful combination of:
 - Greedy round-robin
 - The cycle elimination algorithm
- It is known that both algorithms produce EF1 allocations (but are bad w.r.t. EFX)
- Currently best known approximation for EFX
- It satisfies other fairness criteria as well (to be discussed later)

Approximate EFX allocations

A sketch of our algorithm:

1. Find a “good” ordering of the players
 - Simulate step 2: If an agent envies a predecessor by a large factor, he is moved to a high priority position in the ordering (these are the **privileged agents**)
2. Run one round of Greedy Round-Robin using the above ordering
3. Run a partial round of Greedy round-robin with the reverse ordering
 - Only for the non-privileged agents
4. Run the cycle elimination algorithm for the remaining goods

Summary on EFX allocations

- Severe hardness results for non-additive valuations
- Existence results for identical agents or agents with identical rankings
- 0.618-approximation for additive valuations

Main open questions:

- Do EFX allocations exist for arbitrary **additive** valuations?
 - [Chaudhury, Garg, Melhorn '20]: YES for **n=3** (but pseudo-polynomial time algorithm)
- What is the best approximation we can have in polynomial time?

Some experimental evidence

[M.Sc. thesis of A. Ntokos '19]

- Series of experiments ran on synthetic data
- Data only for additive valuations
- Various distributions utilized for generating the value of every agent for each item

Some conclusions:

- EFX allocations existed in all generated instances
- As the number of items grows, the percentage of EFX allocations decreases
- As the number of items grows, $\Pr[\text{an EFX allocation is also EF}]$ increases
- We probably need more subtle constructions for finding counterexamples to existence (if they exist at all!)

MMS allocations

Existence?

- for $n = 2$
 - YES (via a discrete version of cut-and-choose)
- for $n \geq 3$
 - NO [Procaccia, Wang '14]
 - Known counterexamples build on sophisticated constructions
- How often do they exist for $n \geq 3$?
 - Actually extremely often
 - Extensive simulations [Bouveret, Lemaitre '14] with randomly generated data did not reveal negative examples

Progress on α -MMS allocations

Established approximation guarantees for additive valuations:

- [Procaccia, Wang '14]
 - $2/3$ -ratio, exponential dependence on n
- [Amanatidis, Markakis, Nikzad, Saberi '15]
 - $(2/3-\epsilon)$ -ratio for any $\epsilon>0$, poly-time for any n and m
- [Barman, Murty '17]
 - $2/3$ -ratio, simpler algorithm, more intricate analysis
- [Garg, McGlaughlin, Taki '19]
 - $2/3$ -ratio, simpler algorithm, simpler analysis
- [Ghodsi, Hajiaghayi, Seddighin, Seddighin, Yami '18]
 - $3/4$ -ratio, poly-time but more complicated algorithm

2/3-approximation algorithms

The algorithm of [Barman, Murty '17]

Lemma 1: It suffices to establish the approximation ratio for additive valuations with identical rankings

For agents with identical rankings:

- Run the cycle elimination algorithm
 - but first order the goods in decreasing order of value
- We have seen that this is EFX

Lemma 2: This EFX allocation for additive valuations with identical rankings achieves also a 2/3-approximation for MMS

- **Attention:** not true in general for any EFX allocation

PMMS and GMMS allocations

Existence?

- for $n = 2$
 - YES (both PMMS and GMMS are identical with MMS)
- for $n \geq 3$
 - Still an open problem
 - Recall that PMMS implies EFX
 - If we prove existence for either GMMS or PMMS, we have resolved the same question for EFX too
- How often do they exist for $n \geq 3$?
 - Very often (as with MMS and EFX)
 - Simulations by [Barman et al. '18] did not reveal negative examples

Approximate PMMS and GMMS allocations

[Kurokawa '17]:

- A 0.781-PMMS algorithm
- Currently best known

[Barman et al. '18]:

- A $\frac{1}{2}$ -GMMS algorithm
- Yet another variant of cycle elimination:
 - Run the cycle elimination algorithm with no a priori ordering of the goods
 - in every step, determine an un-envied agent, as before
 - But now we give her the good she values the most out of the available ones
- This particular EF1 allocation is also a $\frac{1}{2}$ -GMMS allocation

Approximate GMMS allocations

Back to our work [Amanatidis, Markakis, Ntokos '20]:

Theorem 1: Our algorithm produces an allocation that is

- EF1
- 0.618-EFX
- 2/3-PMMS
- 0.553-GMMS (and same for MMS)

Theorem 2: We can tweak the initial algorithm to produce a 4/7-GMMS allocation in polynomial time

- 4/7-GMMS also achieved in [Chaudhury, Telikepalli, Mehlhorn, Sgouritsa, '20]
 - But in exponential time

Part 3: Relations between approximate versions of fairness

Approximately fair allocations

- We have already seen how the exact versions of fairness concepts are related to each other
- Does the picture change if we focus on the approximate versions of these concepts?
- **Q:** Does an allocation that is fair or approximately fair according to one notion implies any approximation guarantee regarding a different notion?
 - i.e. what does an α -MMS allocation imply about the other notions?

Approximately fair allocations

Why should we care?

- If we have an algorithm for approximating one notion, this may imply an approximation algorithm for another notion too
- Also technically interesting to understand which implications for exact versions carry over to the approximate versions
 - E.g., we know that a PMMS allocation is also EFX
 - Is an α -PMMS allocation also α -EFX?

Approximately fair allocations

- PMMS \Rightarrow EFX



I found a PMMS allocation.

This is *awesome*,
because it is also EFX

Approximately fair allocations

α -PMMS \Rightarrow β -EFX ???



I found a 0.9-PMMS allocation.

This is *awesome*,
because I think it is also 0.9-EFX

Approximately fair allocations

α -PMMS $\not\Rightarrow$ β -EFX



I found a 0.9-PMMS allocation.

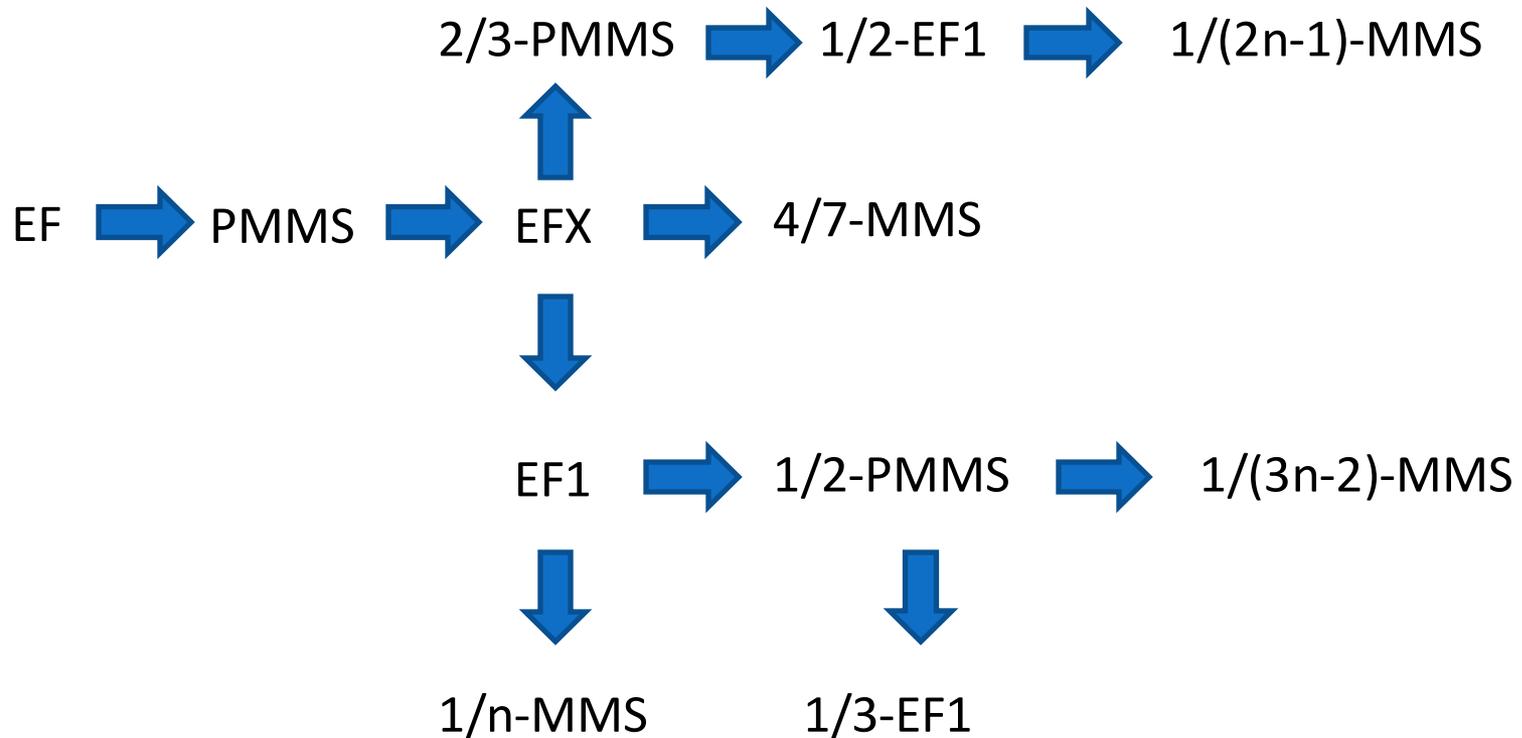
This is awesome,
because I think it is also 0.9-EFX

WRONG!!!

[Amanatidis, Birmpas, Markakis '18]:

A close to complete charting of the relationships between exact and approximate versions of EF1, EFX, MMS, PMMS for additive valuations

An Indicative Picture of the Landscape



Highlights

Some representative results

- **Theorem 1:** Any EFX or PMMS allocation is also a $4/7 (=0.57)$ -MMS allocation. On the other hand, an EFX allocation is not necessarily an α -MMS allocation for $\alpha > 0.59$
- **Theorem 2:** Any α -EF1 allocation is also an $\alpha/(1+\alpha)$ -PMMS allocation for any $\alpha \in (0,1]$, and this is tight
- **Theorem 3:** An α -PMMS allocation is not necessarily a β -EFX allocation for any $\alpha, \beta \in (0, 1)$ [But PMMS \Rightarrow EFX!]
- **Theorem 4:** An α -MMS allocation is not necessarily a β -PMMS, β -EF1, β -EFX allocation for any $\alpha, \beta \in (0, 1)$

Some implications

- By **Theorem 2**, $EF1 \Rightarrow \frac{1}{2}$ -PMMS
 - **Corollary**: Greedy Round-Robin produces a $\frac{1}{2}$ -PMMS allocation
 - Although we know better approximations, this is a much simpler algorithm
- By **Theorem 3**, approximation algorithms for PMMS do not necessarily yield any approximation for EFX
 - MMS approximations also do not yield guarantees for EFX
- EFX approximations seem harder to get

Open problems and future directions

A rich area with several challenging ways to go

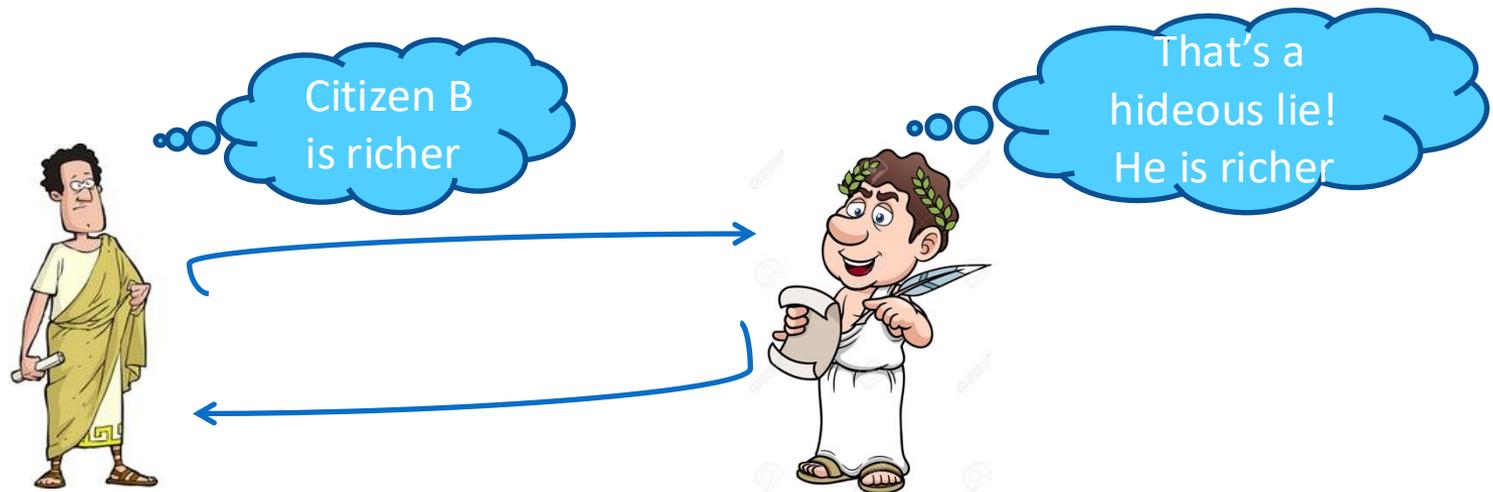
- Conceptual
 - Define or investigate further new notions and models
 - Fairness under mixed manna (desirable and undesirable goods)
 - Fairness in the presence of a social graph: evaluate fairness w.r.t. your neighbors
[Chevalleyre, Endriss, Maudet '07, '17, Abebe, Kleinberg, Parkes '17, Bei, Qiao, Zhang '17, Aziz et al. '18]
- Algorithmic
 - Best approximation algorithm for the relaxed notions?
 - The continuous setting of fair division also has its own algorithmic challenges
- Game-theoretic
 - Mechanism design aspects?
 - So far progress mostly for $n=2$ [Amanatidis et al. '17]

Cycle elimination in ancient Greece:

Αντιδοσις

A mechanism used for the sponsorship of theatrical performances

- The city of Athens would assign a wealthy citizen to cover the costs (for the chorus, costumes, instruments, etc)
- Suppose citizen A was chosen
- In case he wanted to appeal, he had to point to another wealthy citizen



- If this happened, A and B had to exchange all their belongings!

