Fairness with Indivisible Goods: Solution Concepts and Algorithms

Vangelis Markakis

Athens University of Economics and Business (AUEB)
Dept. of Informatics

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]

Some early references

- Ancient Egypt:
  - ➢ Land division around Nile (i.e., of the most fertile land)
- Ancient Greece:
  - ➢ Sponsorships of theatrical performances
    - Undertaken by most wealthy citizens
    - Mechanism used was giving incentives so that wealthier citizens could not avoid becoming sponsors
- First references of the cut-and-choose protocol
  - ➢ Theogony (Hesiod, 8^{th} century B.C.): run between Prometheus and Zeus
  - ➢ Bible: run between Abraham and Lot

Available implementations

- http://www.spliddit.org
  - Jonathan Goldman, Ariel Procaccia
  - Algorithms for various classes of problems (rent division, division of goods, etc)
- http://www.nyu.edu/projects/adjustedwinner/
  - Steven Brams, Alan Taylor
  - Implementation of the “adjusted winner” algorithm for 2 players
- https://www.math.hmc.edu/~su/fairdivision/calc/
  - Francis Su
  - Implementation of algorithms for allocating goods with any number of players

Modeling Fair Division Problems

Preferences:

- Modeled by a valuation function for each agent
- $v_i(S)$ = value of agent i for obtaining a subset S

Type of resources:

1. Continuous models
   - Infinitely divisible resources (usually just the interval [0, 1])
   - Valuation functions: defined on subsets of [0, 1]
2. Discrete models
   - Set of indivisible goods
   - Valuation functions: defined on subsets of the goods

The discrete setting

For this talk:

- Resources = a set of indivisible goods M = {1, 2, …, m}
- Set of agents: N = {1, 2, …, n}
- An allocation of M is a partition $S = (S_1, S_2, \ldots, S_n)$, $S_i \subseteq M$
  - $\cup_i S_i = M$ and $S_i \cap S_j = \varnothing$

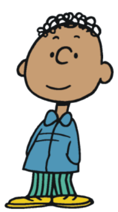



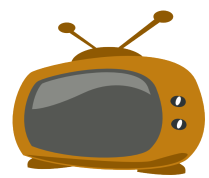


Valuation functions

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

Valuation functions

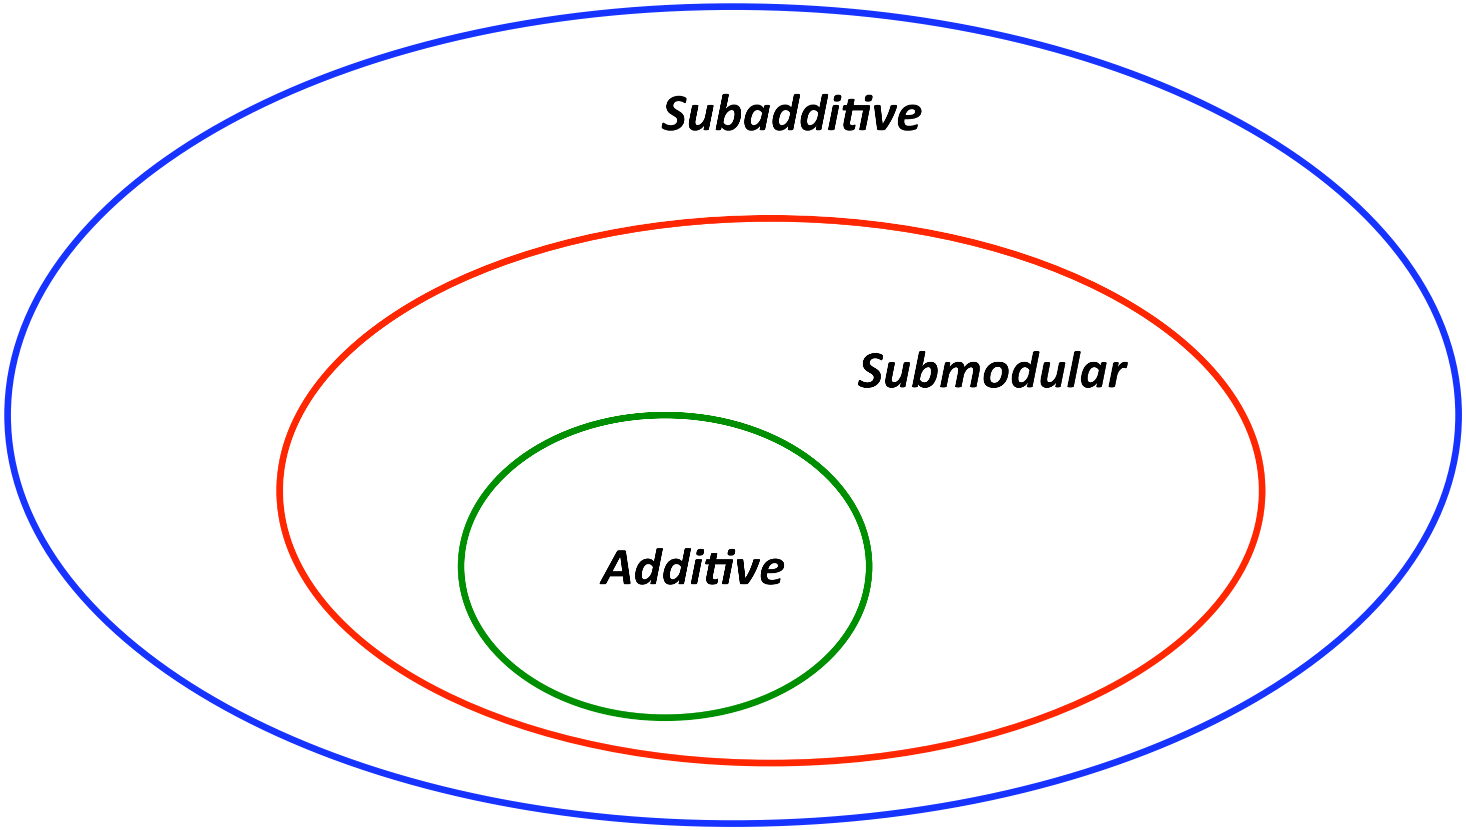

The discrete setting

Example with additive valuations

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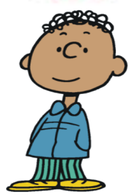	30	40	35	5	40
Marcie		30	20	40	30	0

Part 1: A hierarchy of some solution concepts in fair division

Solution Concepts

1. Proportionality

An allocation $(S_1, S_2,..., S_n)$ 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 $(S_1, S_2, \ldots, S_n)$ 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):

Envy-freeness $\Rightarrow n \cdot v_i(S_i) \geq v_i(M) \Rightarrow$ Proportionality

The discrete setting

In our example:

					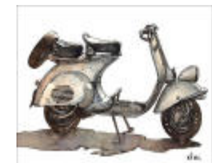
	35	5	25	0	35
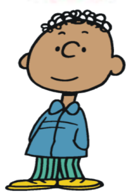	30	40	35	5	40
	30	20	40	30	0

A proportional and envy-free allocation

The discrete setting

In our example:

					
	35	5	25	0	35
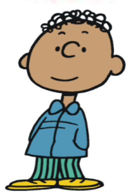	30	40	35	5	40
	30	20	40	30	0

A proportional but not envy-free allocation

Solution Concepts

3. Competitive Equilibrium from Equal Incomes (CEEI)

Suppose each agent is given the same (virtual) budget to buy goods.

A CEEI consists of

- An allocation $S = (S_1, S_2, ..., S_n)$
- A pricing on the goods $p = (p_1, p_2, ..., p_m)$

such that $v_i(S_i)$ is maximized subject to the budget constraint

An allocation $S = (S_1, S_2, ..., S_n)$ is called a CEEI allocation if it admits a pricing $p = (p_1, ..., p_m)$, such that (S, p) is a CEEI

Solution Concepts

- A well established notion in economics [Foley '67, Varian '74]
- Combining fairness and efficiency
- Quote from[Arnsperger '94]: "To many economists, CEEI is the description of perfect justice"

Claim: A CEEI allocation is

- envy-free (due to equal budgets)
- Pareto-efficient in the continuous setting
- Pareto-efficient in the discrete setting when valuations are strict (no 2 bundles have the same value)

Containment Relations in the space of allocations

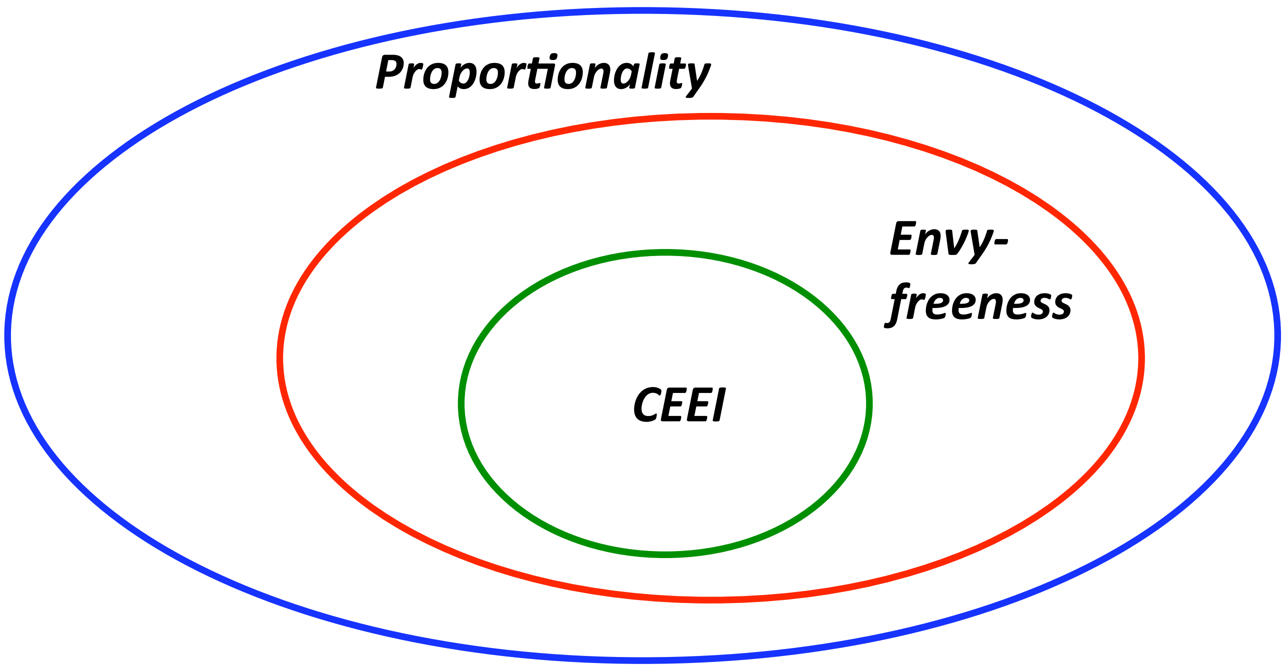

Some issues

- All 3 definitions are “too strong” for indivisible goods
- No guarantee of existence
- More appropriate for the continuous setting (existence is always guaranteed)
- Need to explore relaxed versions of fairness

Solution Concepts

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

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

- 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

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

An allocation $(S_1, S_2, \ldots, S_n)$ satisfies EFX, if for any players *i* and *j*, and any good $a \in S_j$, we have $v_i(S_i) \geq v_i(S_j \setminus \{a\})$

- 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

6. 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

$$\mu_i := \mu_i(n, M) = \max_{S \in \Pi_n(M)} \min_{S_j \in S} v_i(S_j)$$

- max is over all possible partitions of M

- min is over all bundles of a partition $S = (S_1, S_2, \ldots, S_n)$

Introduced by [Budish '11]

Solution Concepts

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

Fact: Proportionality $\Rightarrow$ MMS

Maximin shares

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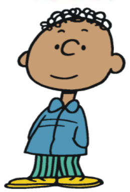	30	40	35	5	40	$\mu_2 = 40$
	30	20	40	30	0	$\mu_3 = 30$

MMS vs EF1 (and vs EFX)

How do MMS allocations compare to EF1 and EFX?

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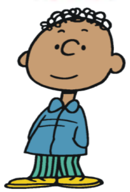	30	40	35	5	40	$\mu_2 = 40$
	30	20	40	30	0	$\mu_3 = 30$

A MMS allocation that does not satisfy EF1

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

Relations between fairness criteria

For subadditive valuation functions

- Upper part holds for general monotone valuations

$$\begin{array}{lll} & & \text{EFX} \Rightarrow \text{EF1} \\ & \nearrow & \\ \text{CEEI} \Rightarrow \text{EF} & & \\ & \searrow & \\ & & \text{Proportionality} \Rightarrow \text{MMS} \end{array}$$

Relations between fairness criteria

Pictorially:

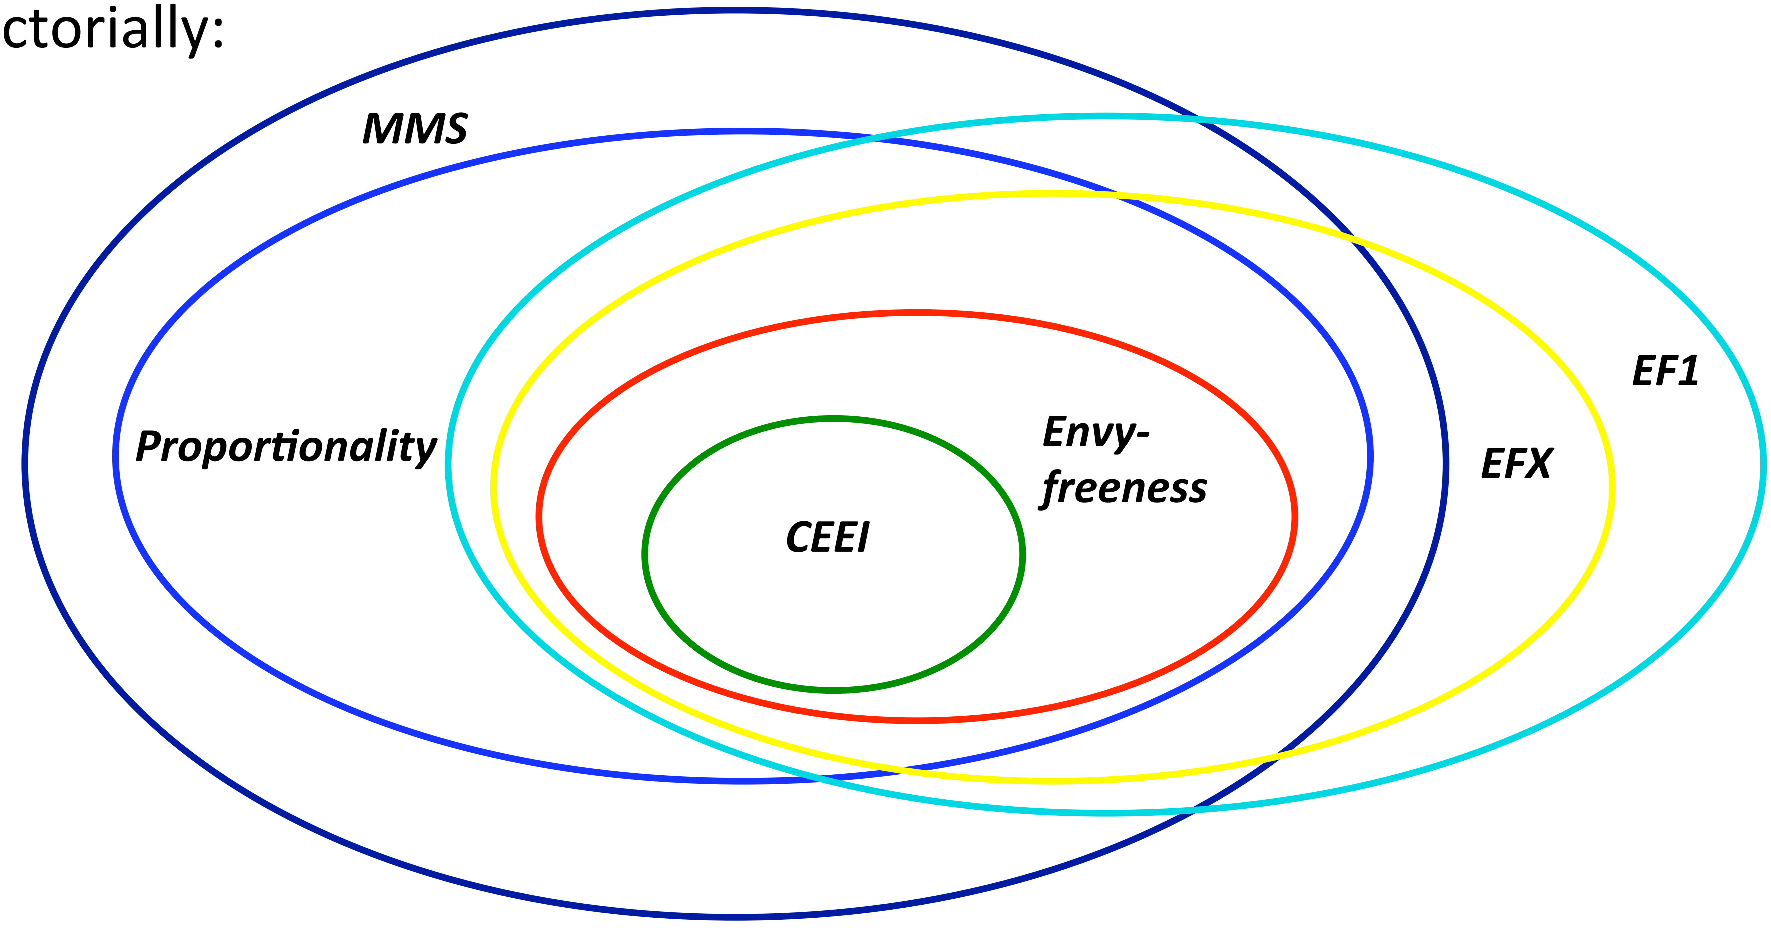

Part 2: Existence and Computation

Envy-freeness and Proportionality

Mostly bad news:

- No guarantee of existence for either proportionality or envy-freeness
- NP-hard to decide existence even for n=2 (equivalent to makespan for 2 identical processors)
- NP-hard to compute decent approximations
  - E.g. For approximating the minimum envy allocation [Lipton, Markakis, Mossel, Saberi '04]
- Still open to understand if there exist subclasses that admit good approximations
- On the positive side: Existence with high prob. on random instances, when n = O(m/logm) [Dickerson et al. '14]

Part 2a:
EF1 and EFX

EF1

Existence of EF1 allocations?

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

EF1 for Additive Valuations

Existence established through an algorithm

Algorithm 1 - Greedy Round-Robin

- Fix an ordering of the agents
- While there exist unallocated items
  - Let *i* be the next agent in the round-robin order
  - Ask *i* to pick his most desirable item among the unallocated ones

Algorithm 1 works for additive valuations

Proof: Throughout the algorithm, each player may have an advantage only by 1 item w.r.t. other players $\Longrightarrow$ EF1

EF1 for General Valuations

- For non-additive valuations, more insightful to look at a graph-theoretic representation
- 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?

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 by exchanging bundles, as dictated by each cycle
    - Until we have a node with in-degree 0

EF1 for General Valuations

Algorithm 2 – The Cycle Elimination Algorithm

- Fix an ordering of the goods, say, 1, 2, ..., 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

EFX

Existence of EFX allocations?

- for n = 2
  - YES (for general valuations)
- for n ≥ 3
  - Great open problem!
  - Guaranteed to exist only for agents with identical valuations

A detour: the leximin solution

[Rawls '71]

The leximin solution is the allocation that

- Maximizes the minimum value attained by an agent
- If there are multiple such allocations, pick the one maximizing the 2nd minimum value
- Then maximize the 3rd minimum value
- And so on...

- This induces a total ordering over allocations
- Leximin is a global maximum under this ordering

Existence results for EFX allocations

[Plaut, Roughgarden '18]: a slightly different version

A leximin++ allocation

- Maximizes the minimum value attained by an agent
- Maximizes the bundle size of the agent with the minimum value
- Then maximizes the 2nd minimum value
- Followed by maximizing the bundle size of the 2nd minimum value
- And so on...

Theorem: For general but identical agents, the leximin++ solution is EFX

Algorithmic results

[Plaut, Roughgarden '18]:

Separation between general and additive valuations

Theorem:

1. exponential lower bound on query complexity
   - 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 results

Algorithm for additive valuations with identical rankings:

Run the cycle elimination algorithm, by ordering 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
- Hence the allocation remains EFX throughout the algorithm

Algorithmic results

Algorithm for 2 agents and arbitrary additive valuations

Variation of cut and choose

- Agent 1 runs the previous algorithm with 2 copies of herself
- Agent 2 picks her favorite out of the 2 bundles created
- Agent 1 picks the left over bundle

Part 2b:
MMS allocations

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

Computation

Approximate MMS allocations

Q: What is the best α for which we can compute an allocation $(S_1, S_2, \ldots, S_n)$ satisfying $v_i(S_i) \geq \alpha\, \mu_i$ for every *i*?

We will again start with additive valuations

Approximation Algorithms for Additive Valuations

For n=2

- NP-hard to even compute the quantity μ_i for agent i
- Existence proof of MMS allocations yields an exponential algorithm
    1. Let player 1 compute a partition that guarantees μ_1 to him
        - i.e., a partition that is as balanced as possible
    2. Player 2 picks the best out of the 2 bundles
- Convert Step 1 to poly-time by losing ε, e.g. using the PTAS of [Woeginger '97]

Corollary: For n=2, we can compute in poly-time a (1-ε)-MMS allocation

Approximation Algorithms for Additive Valuations

For $n \geq 3$

- Start with an additive approximation
- Recall the greedy round-robin algorithm (Algorithm 1)

Theorem:

Greedy Round-Robin produces an allocation $(S_1, S_2, \ldots, S_n)$ such that

$$v_i(S_i) \geq \mu_i - v_{max}, \quad \text{where } v_{max} = \max v_{ij}$$

Approximation algorithms for additive valuations

When does Greedy Round-Robin perform badly?

- In the presence of goods with very high value
- BUT: each such good can satisfy some agent
- Suggested algorithm: Get rid of the most valuable goods before running Greedy Round-Robin

Algorithm 3:

- Let S := M, and $\alpha_i := v_i(S)/n$
- While ∃ i, j, such that $v_{ij} \geq \alpha_i/2$,
    - allocate j to i
    - n := n-1, S := S \ {j}, recompute the α_i's
- Run Greedy Round-Robin on remaining instance

A ½-approximation for Additive Valuations

All we need is to ensure a monotonicity property

Lemma:

If we assign a good *j* to some agent, then for any other agent $i \neq$ j:

$$\mu_i(n-1, M\setminus\{j\}) \geq \mu_i(n, M)$$

Theorem:

Algorithm 2 produces an allocation $(S_1, S_2, \ldots, S_n)$ such that for every agent *i*:

$$v_i(S_i) \geq 1/2\ \mu_i(n, M) = 1/2\ \mu_i$$

Beyond 1/2...

- Algorithm 2 is tight
- What if we change the definition of “valuable” by considering $v_{ij} \geq 2\alpha_i/3$ instead of $\alpha_i/2$?
- Not clear how to adjust Greedy Round-Robin for phase 2
- Beating 1/2 needs different approaches

Beyond 1/2...

2/3-approximation guarantees:

- [Procaccia, Wang '14]
  - 2/3-ratio, exponential dependence on n
- [Amanatidis, Markakis, Nikzad, Saberi '15]
  - (2/3-ε)-ratio for any ε>0, poly-time for any n and m
- [Barman, Murty '17]
  - 2/3-ratio, poly-time for any n and m

2/3-approximation algorithms

Recursive algorithms of

[Procaccia, Wang '14], [Amanatidis, Markakis, Nikzad, Saberi '15]

Based on:

- Exploiting certain monotonicity properties of $\mu_i(\cdot, \cdot)$
  - To be able to move to reduced instances
- Results from job scheduling
  - To be able to compute approximate MMS partitions from the perspective of each agent
- Matching arguments (perfect matchings + finding counterexamples to Hall's theorem when no perfect matchings exist)
  - To be able to decide which agents to satisfy within each iteration

2/3-approximation algorithms

Recursive algorithms of

[Procaccia, Wang '14], [Amanatidis, Markakis, Nikzad, Saberi '15]

High level description:

- Each iteration takes care of ≥ 1 person, until no-one left
- During each iteration,

  Let {1, 2, ..., k} = still active agents

  1. Ask one of the agents, say agent 1, to produce a MMS partition with k bundles according to his valuation function
  2. Find a subset of agents such that:
     a) they can be satisfied by some of these bundles
     b) the remaining goods have "enough" value for the remaining agents

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

Lemma 2: For additive valuations with identical rankings, the cycle elimination algorithm (after ordering the goods in decreasing order of value) achieves a 2/3-approximation

The case of n = 3 agents

- An intriguing case...
- For n=2, MMS allocations always exist
- The problems start at n=3!
- Still unclear if there exists a PTAS

Progress achieved so far:

Algorithms	Approx. ratio
[Procaccia, Wang '14]	3/4
[Amanatidis, Markakis, Nikzad, Saberi '15]	7/8
[Gourves, Monnot '17]	8/9

Non-additive valuations

- None of the algorithms go through with non-additive valuations
- No positive results known for arbitrary valuations

Theorem [Barman, Murty '17]: For agents with submodular valuations, there exists a polynomial time 1/10-approximation algorithm

And some more recent progress

[Ghodsi, Hajiaghayi, Seddighin, Seddighin, Yami '17]:

Positive results for various classes of valuation functions:

- Additive: Polynomial time ¾-approximation
- Submodular: Polynomial time 1/3-approximation
- Subadditive: Existence of O(logm)-approximation

Part 3: Related open problems and other research directions

Other fairness notions

Can we think of alternative relaxations to envy-freeness and/or proportionality?

[Caragiannis et al. '16]:

- Pairwise MMS allocations
  - Consider an allocation $S = (S_1, S_2, \ldots, S_n)$, and a pair of players, i, j
  - Let B:= all partitions of $S_i \cup S_j$ into two sets (B_1, B_2)
  - Fairness requirement for every pair i, j:

$$v_i(S_i) \geq \max_{B=(B_1,B_2)} \min\{v_i(B_1), v_i(B_2)\}$$

- A stronger criterion than EFX
- Related but incomparable to MMS allocations
- Existence of φ-approximation (golden ratio)
  - Open problem whether pairwise MMS allocations always exist

Other fairness notions

Can we think of alternative relaxations to envy-freeness and/or proportionality?

Fairness in the presence of a social graph

[Chevaleyre, Endriss, Maudet '17, Abebe, Kleinberg, Parkes '17, Bei, Qiao, Zhang '17]

- Evaluate fairness with regard to your neighbors
  - Most definitions easy to adapt
  - E.g., graph envy-freeness: suffices to not envy your neighbors

[Caragiannis et al. '18]:

- More extensions, without completely ignoring the goods allocated to non-neighbors

Mechanism design aspects

- So far we assumed agents are not strategic
- Can we design truthful mechanisms?
- [Amanatidis, Birmpas, Christodoulou, Markakis ’17]:
  - Mechanism design without money
  - Tight results for 2 players through a characterization of truthful mechanisms
  - Best truthful approximation for MMS: O(1/m)
  - Truthful mechanisms for EF1: only if $m \leq 4$
- Characterization results for ≥3 players?

The continuous setting

- Cake: M = [0, 1]
- Set of agents: N = {1, 2, …, n}
- Valuation functions:
  - Given by a non-atomic probability measure v_i on [0, 1], for each i
- Access to the valuation functions:
  - Value queries: ask an agent for her value of a given piece
  - Cut queries: ask an agent to produce a piece of a given value

Envy-free allocations in the continuous setting

- Envy-free (and hence proportional) allocations always exist

Computation?

- n=2: cut-and-choose (2 queries)
- n=3: [Selfridge, Conway circa 60s] (less than 15 queries)
- n=4: [Aziz, Mckenzie '16a] (close to 600 queries)
- General n:
  - [Brams, Taylor '95]: Finite procedure but with no upper bound on number of queries
  - [Aziz, Mackenzie '16b]: First bounded algorithm but with exceptionally high complexity

#queries ≤ $n^{n^{n^{n^{n^{n}}}}}$

Envy-free allocations in the continuous setting

Lower bounds

- Contiguous pieces: there can be no finite protocol that produces envy-free allocation
- Non-contiguous pieces: $\Omega(n^2)$ [Procaccia '09]
  - Separating envy-freeness from proportionality
- Can we do shorten the gap between the upper and lower bound?

Summarizing...

A rich area with several challenging ways to go

- Conceptual
  - Define or investigate further new notions
- Algorithmic
  - Best approximation for MMS allocations?
  - EFX for arbitrary additive valuations?
  - Algorithms for the continuous setting?
- Game-theoretic
  - Mechanism design aspects?


Thank you