COMSOC Fair Division Lecture II: Indivisible Goods

Dominik Peters

2025-02-19

Recap: Lecture 1

- Cake cutting
  - Additive, continuous, normalized valuations over $[0, 1]$.
  - Protocols in the Robertson-Webb model (eval, cut queries)
  - Cut and choose
  - Steinhaus protocol for proportionality, $n = 3$
  - Moving knife protocol for proportionality, $n \geq 2$, $O(n^2)$ queries
  - Even–Paz divide-and-conquer protocol for proportionality, $O(n \log n)$ queries. Asymptotically optimal.
  - Envy-freeness for $n = 3$ with Selfridge–Conway protocol (trimmings etc.)

Recap: Lecture 1

- Rent division
  - Need to assign n roommates to n rooms.
  - Quasi-linear utilities.
  - Envy-free allocations exist (LP duality)
  - Lemmas: prices combine with optimal allocations (bipartite matching) but doesn't matter which optimal matching
  - Maximin, lexislack
- (A taste of) Homogeneous divisible goods
  - Additive linear valuations.
  - Pareto-optimality, proportionality, envy-freeness, strategyproofness.
  - Trivial rule is proportional, envy-free, strategyproof
  - Serial dictatorship is Pareto-optimal, strategyproof
  - Nash ($\max \prod_{i \in N} u_i(A)$) is Pareto-optimal, envy-free (no proof)

Today

- Fair allocation of indivisible items
  - Model, definitions
  - Computational complexity
  - Envy-freeness up to 1 good
- Random assignment
  - Probabilistic serial mechanism

Allocation of indivisible items

- $N = \{1, \ldots, n\}$ is a set of agents.
- $O = \{o_1, \ldots, o_m\}$ is a set of items/objects/goods.
- An allocation is a list $A = (A_1, \ldots, A_n)$,
  where $A_i \subseteq O$ is a bundle of items assigned to agent i.
    - Bundles must be pairwise disjoint.
    - We also must have $A_1 \cup \cdots \cup A_n = O$; if not, it is a partial allocation.
- Each agent i has a valuation function $v_i : 2^O \to \mathbb{R}_{\geq 0}$ that is monotonic:
  $B_1 \subseteq B_2 \implies v_i(B_1) \leq v_i(B_2)$. (items are goods)
- A valuation function is additive if $v_i(B) = \sum_{o \in B} v_i(\{o\})$ for all $B \subseteq O$.
    - In this case, we also write $v_i(o) := v_i(\{o\})$.
    - What are some examples of non-additive valuation functions?

Georgios Amanatidis, Haris Aziz, Georgios Birmpas, Aris Filos-Ratsikas, Bo Li, Hervé Moulin, Alexandros A. Voudouris, and Xiaowei Wu. "Fair Division of Indivisible Goods: Recent Progress and Open Questions". In: *Artificial Intelligence* (2023), p. 103965

Example

	🏠	🧸	💎	🍭	⚽
🐕	15	3	2	2	6
🐬	7	5	5	5	7
🐦	20	3	3	3	3

Proportionality and envy-freeness

Let A be an allocation.

- A is proportional if $v_i(A_i) \geq \frac{1}{n} v_i(O)$ for every $i \in N$.
- A is envy-free if $v_i(A_i) \geq v_i(A_j)$ for all $i, j \in N$

Question: are there examples where no envy-free allocation exists? no proportional allocation?

Proportionality and envy-freeness

Let A be an allocation.

- A is proportional if $v_i(A_i) \geq \frac{1}{n} v_i(O)$ for every $i \in N$.
- A is envy-free if $v_i(A_i) \geq v_i(A_j)$ for all $i, j \in N$

Question: are there examples where no envy-free allocation exists? no proportional allocation?

Yes. $N = \{1, 2\}$, $O = \{o_1\}$, $v_1(o_1) = v_2(o_1) = 1$.

- For the allocation $(\{o_1\}, \emptyset)$, 2 envies 1 and doesn't get proportional share.
- For the allocation $(\emptyset, \{o_1\})$, 1 envies 2 and doesn't get proportional share.

Deciding existence

Consider the following decision problem [and variant]:
EXISTENCE OF PROPORTIONAL [ENVY-FREE] ALLOCATION

- **Input**: Additive valuations $(v_i(o))_{i \in N, o \in O}$.
- **Question:** Does there exist a (complete) allocation A that is proportional? [that is envy-free?]

This problem is NP-complete.

Obvious reduction from PARTITION, works even for $n = 2$ agents.

- **Input**: List of numbers $(x_1, \dots, x_m)$
- **Question:** Does there exist a partition (S_1, S_2) of $\{1, \dots, m\}$ such that $\sum_{i \in S_1} x_i = \sum_{i \in S_2} x_i$?

Exercise: This only shows *weak* NP-hardness (binary encoding of numbers). Show the problem is strongly NP-hard (unrestricted n).

Some allocation rules

- Maximize utilitarian social welfare: Pick an allocation A that maximizes $\sum_{i \in N} v_i(A_i)$.
- Maximize egalitarian social welfare: Pick an allocation A that maximizes $\min_{i \in N} v_i(A_i)$.
- Maximize Nash social welfare: Pick an allocation A that maximizes $\prod_{i \in N} v_i(A_i)$.
    - This is the same as maximizing $\sum_{i \in N} \log v_i(A_i)$.
    - This is scale-free: multiplying the valuations of an agent by any factor does not change the optimal allocation.
    - It lies "between" utilitarian and egalitarian social welfare: $\min_{i \in N} v_i(A_i) \leq \sqrt[n]{\prod_{i \in N} v_i(A_i)} \leq \frac{1}{n} \sum_{i \in N} v_i(A_i)$. (AM-GM inequality)

Question: What is the computational complexity of computing optimal allocations for these objectives?

Envy-freeness up to 1 good (EF1)

An allocation is envy-free up to 1 good (EF1) if for all $i, j \in N$,

$$\text{either } v_i(A_i) \geq v_i(A_j) \text{ or there is } o \in A_j \text{ with } v_i(A_i) \geq v_i(A_j \setminus \{o\}).$$

Theorem: An EF1 allocation always exists.

Eric Budish. "The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes". In: *Journal of Political Economy* 119.6 (2011), pp. 1061–1103

Example

	🏠	🧸	💎	🍭	⚽
🐕	**15**	3	2	2	**6**
🐬	7	**5**	5	5	7
🦤	20	3	**3**	**3**	3

Example

	🏠	🧸	💎	🍭	⚽
🐕	15	3	2	2	6
🐬	7	5	5	5	7
🐦	20	3	3	3	3

Example

	🏠	🧸	💎	🍭	⚽
🐕	15	3	2	2	6
🐬	7	5	5	5	7
🐦	20	3	3	3	3

Example

	🏠	🧸	💎	🍭	⚽
🐕	15	3	2	2	6
🐬	7	5	5	5	7
🐦	20	3	3	3	3

Round robin rule

Consider the following procedure:

Repeatedly go through the agents in order

$$1\ 2\ 3\ \dots n\ 1\ 2\ 3\ \dots n\ 1\ 2\ 3\ 4$$

and on each turn, let the agent pick an unpicked good that is most valuable to them.

- Clearly, this is EF1 for agent 1 (in fact, for 1 it is envy-free).
- But it is also EF1 for everyone else. Consider for example agent 3. Let him ignore the first item that agent 1 picked, and the first item that agent 2 picked. With these ignored, no envy remains.

Question: what are some other agent orderings that guarantee EF1? what are some that don't?

Question: does this algorithm work for non-additive valuations?

Question: what are some interesting properties of this rule?

Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. "The unreasonable fairness of maximum Nash welfare". In: *ACM Transactions on Economics and Computation (TEAC)* 7.3 (2019), pp. 1–32

Envy graph, cycle elimination

Given an allocation A, its envy graph is the directed graph with 1 vertex for each agent, and an arc from i to j if i envies j.

Consider some allocation A. Suppose the envy graph has a cycle 1-2-3-4-5-1, meaning that

$$
\begin{aligned}
v_1(A_1) &< v_1(A_2) \\
v_2(A_2) &< v_2(A_3) \\
v_3(A_3) &< v_3(A_4) \\
v_4(A_4) &< v_4(A_5) \\
v_5(A_5) &< v_5(A_1).
\end{aligned}
$$

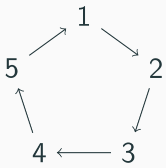

Then we can eliminate the cycle by giving A_2 to A_1, A_3 to A_2, etc. The resulting allocation has fewer envy edges (and it is a Pareto improvement). If A was EF1, then same is true after.

Envy graph algorithm

1. Start with the empty (partial) allocation A.
2. While not all items are allocated:
   - Compute the envy graph for A, and update A by eliminating any cycles.
   - Now the envy graph has no cycles.
   - Pick an agent i who is a source in the envy graph, i.e. is not envied by anybody.
   - Add i's favorite unallocated item o to A_i.

Theorem: This algorithm always terminates with an EF1 allocation.

Proof: The partial allocation is EF1 throughout: Let A be allocation before adding o, and B the allocation afterwards. The only possible new EF1 violation is towards i, but

$$v_j(B_j) = v_j(A_j) \overset{i \text{ source}}{\geq} v_j(A_i) = v_j(B_i \setminus \{o\}).$$

Question: Does this algorithm work for non-additive valuations?

Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. "On approximately fair allocations of indivisible goods". In: *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*. 2004, pp. 125–131

Pareto-optimality

An allocation A is Pareto-optimal if there is no other allocation B such that $v_i(B_i) \geq v_i(A_i)$ for all $i \in N$ and $v_i(B_i) > v_i(A_i)$ for some $i \in N$.

Questions: Which rules are Pareto-optimal? Is round robin? Is envy graph?

Pareto-optimality

An allocation A is Pareto-optimal if there is no other allocation B such that $v_i(B_i) \geq v_i(A_i)$ for all $i \in N$ and $v_i(B_i) > v_i(A_i)$ for some $i \in N$.

Questions: Which rules are Pareto-optimal? Is round robin? Is envy graph?

Question: Does there always exist a Pareto-optimal EF1 allocation?

Maximizing Nash Welfare is PO and EF1

The MNW (Max Nash Welfare) rule selects an allocation maximizing $\prod_{i \in N} v_i(A_i)$.

Clearly, this rule is PO.*

Proved in 2016: it also satisfies EF1.

> Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. "The unreasonable fairness of maximum Nash welfare". In: *ACM Transactions on Economics and Computation (TEAC)* 7.3 (2019), pp. 1–32

Nice proof due to Nisarg Shah:

- Fix any agents $i, j \in N$, and consider moving object $o \in A_j$ from A_j to A_i.
- $v_i(A_i \cup \{o\}) \cdot v_j(A_j \setminus \{o\}) \leq v_i(A_i) \cdot v_j(A_j)$.
- $\Rightarrow$: $1 - v_j(o)/v_j(A_j) \leq 1 - v_i(o)/(v_i(A_i) + v_i(o))$.
- $\Rightarrow$: $v_j(o)/v_j(A_j) \geq v_i(o)/(v_i(A_i) + v_i(o^*))$ for $o^* \in \arg\max_{o' \in A_j} v_i(o')$.
- Sum over all $o \in A_j$.

Maximizing Nash Welfare

- Used on Spliddit
- NP-hard (consider two agents with identical valuations)
- Can calculate with ILP.

Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. "The unreasonable fairness of maximum Nash welfare". In: *ACM Transactions on Economics and Computation (TEAC)* 7.3 (2019), pp. 1–32

- `https://pref.tools/nash-indivisible/`
- There is a pseudo-polynomial algorithm achieving PO + EF1 (i.e., polynomial in n, m, $\max_{i,o} v_i(o)$).

Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. "Finding fair and efficient allocations". In: *Proceedings of the 2018 ACM Conference on Economics and Computation (EC)*. 2018, pp. 557–574

- Nash is the unique welfarist rule (one maximizing some function of the utilities of the agents) that satisfies EF1.

Sheung Man Yuen and Warut Suksompong. "Extending the characterization of maximum Nash welfare". In: *Economics Letters* 224 (2023), p. 111030

An algorithm for PO + EF1

Two natural strategies for designing an algorithm getting PO + EF1:

- Start with EF1 allocation, and repeatedly Pareto-improve it.
- Start with a PO allocation, and make it fairer until it is EF1.

An algorithm for PO + EF1

Fact 1: If $w_1, \ldots, w_n > 0$ are positive weights and A maximizes weighted welfare $\sum_{i \in N} w_i u_i(A_i)$, then A is Pareto optimal.

Fact 2: Allocation A maximizes weighted welfare if and only if each object o is allocated to an agent i with maximum weighted utility $w_i u_i(o)$ for o.

Fact 3: If in a weighted welfare maximizing allocation, we have for all $i, j \in N$ that $w_i u_i(A_i) \geq w_j u_j(A_j \setminus \{o\})$ for some o, then A is EF1.

High-level idea of algorithm:

- Start with $w_1 = \cdots = w_n = 1$, and let A be a utilitarian allocation.
- Try to move objects from higher-utility agents to lower-utility agents as much as possible.
- Increase weight of lowest-utility agents until they become eligible to get additional items.

Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. "Finding fair and efficient allocations". In: *Proceedings of the 2018 ACM Conference on Economics and Computation (EC)*. 2018, pp. 557–574

An algorithm for PO + EF1

The algorithm actually gets something stronger due to:

Fact 4: If $w_1, \dots, w_n > 0$ are positive weights and A maximizes $\sum_{i \in N} w_i u_i(A_i)$, then A is fractionally Pareto optimal (fPO), i.e., is not even dominated by a fractional allocation.

- fPO is easy to check by linear programming
- when fPO is compatible with another property, there is hope for an algorithm
- but fPO is very restrictive:

	🚗 car	🎈 balloon	🧦 socks
👨 A	**100**	2	1
👩 B	100	**1**	**2**

dominated by

👨: 0.98 🚗 + 🎈
👩: 0.02 🚗 + 🧦

Is EF1 enough?

	car	balloon	socks
A	100	2	1
B	100	1	2

	car	balloon	socks
A	100	2	1
B	100	1	2

Envy-freeness up to any good (EFX)

Definition: An allocation A satisfies EFX if for all $i, j \in N$, and for any good $o \in A_j$, we have

$$v_i(A_i) \geq v_i(A_j \setminus \{o\})$$

- Open: Does there always exist an EFX allocation?
- Open: For positive valuation, does there always exist a PO + EFX allocation?
- Known: exists for identical valuations and for ordered valuations.
- Known: exists for two agents (easy), exists for three agents (very hard)

Bhaskar Ray Chaudhury, Jugal Garg, and Kurt Mehlhorn. "EFX exists for three agents". In: *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*. 2020, pp. 1–19

Bhaskar Ray Chaudhury, Jugal Garg, Kurt Mehlhorn, Ruta Mehta, and Pranabendu Misra. "Improving EFX guarantees through rainbow cycle number". In: *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*. 2021, pp. 310–311

- Recent work about finding partial allocations that are EFX.

Ben Berger, Avi Cohen, Michal Feldman, and Amos Fiat. "Almost full EFX exists for four agents". In: *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*. vol. 36. 5. 2022, pp. 4826–4833

EFX: Identical valuations

Suppose agents have identical valuations: $v_i(o) = v_j(o)$ for all $o \in O$ and $i, j \in N$.

If we further assume positive valuations ($v_i(o) > 0$), then a leximin allocation satisfies EFX – this is an allocation that maximizes the utility of the worst-off agent, and subject to this, maximizes the utility of the second-worst-off agent, etc.

Proof: Let A be a leximin allocation. If A fails EFX, there are i, j such that $v(A_i) < v(A_j \setminus \{o\})$. If we move o from A_j to A_i, then i becomes strictly better off but is still worse off than j. Thus the move yields an allocation that is leximin-better than A, contradiction.

Without assuming positive valuations, can use leximin$++$ which maximizes lowest utility, then the size of the bundle with lowest utility, then the second-lowest utility, then the size, ...

> Benjamin Plaut and Tim Roughgarden. "Almost envy-freeness with general valuations". In: *SIAM Journal on Discrete Mathematics* 34.2 (2020), pp. 1039–1068

- PO, works beyond additive valuations

EFX: Ordered valuations

🏠 ≻ 🚗 ≻ 💻 ≻ 🍭

Generalization of identical valuations: label $O = \{o_1, \ldots, o_m\}$. An instance has ordered valuations if for all $i \in N$, we have $v_i(o_1) \geq v_i(o_2) \geq \cdots \geq v_i(o_m)$.

Envy graph algorithm gives an EFX allocation for ordered valuations.

Reason (intuitively): when an object o is added during the algorithm, all agents agree that o is the worst object among allocated objects. Since the object is added to an unenvied agent i, any new envy towards i can be removed by ignoring o.

> Benjamin Plaut and Tim Roughgarden. "Almost envy-freeness with general valuations". In: *SIAM Journal on Discrete Mathematics* 34.2 (2020), pp. 1039–1068

EFX: Two agents

Divide and choose

Take agent 1's valuation v_1, and consider a virtual instance with two copies of v_1, and take an EFX allocation (A_1, A_2) using one of the previous algorithms.

Let agent 2 choose their preferred of the two bundles; the other bundle goes to 1.

> Benjamin Plaut and Tim Roughgarden. "Almost envy-freeness with general valuations". In: *SIAM Journal on Discrete Mathematics* 34.2 (2020), pp. 1039–1068

EFX: Approximation

Envy graph algorithm selects an allocation that satisfies $\frac{1}{2}$-EFX.

For all $i, j \in N$: $v_i(A_i) \geq \frac{1}{2} \cdot v_i(A_j \setminus \{o\})$ for all $o \in A_j$.

> Benjamin Plaut and Tim Roughgarden. "Almost envy-freeness with general valuations". In: *SIAM Journal on Discrete Mathematics* 34.2 (2020), pp. 1039–1068

Non-additive valuations?

A valuation function $v_i : 2^O \to \mathbb{R}$ is submodular if for all $A \subseteq B$ and all $x \in O \setminus B$,

$$v_i(B \cup \{o\}) - v_i(B) \leq v_i(A \cup \{o\}) - v_i(A).$$

Example: course allocation.

An EF1 allocation always exists for submodular violation.

- Open: does a PO + EF1 allocation always exist?
  Nash is not EF1.

What about chores?

A chore for agent i is an item with $v_i(o) < 0$.

We can define EF1 for mixed instances as follows:

> *An allocation A is EF1 if for all $i, j \in N$, either $v_i(A_i) \geq v_i(A_j)$ or there is some object $o \in A_i$ such that*
> $$v_i(A_i \setminus \{o\}) \geq v_i(A_j)$$

Haris Aziz, Ioannis Caragiannis, Ayumi Igarashi, and Toby Walsh. "Fair allocation of indivisible goods and chores". In: *Autonomous Agents and Multi-Agent Systems* 36 (2022), pp. 1–21

- EF1 exists via round robin, but need to be careful with envy graph.
- For 2 agents, can do PO + EF1 (adjusted winner).
- Open: can we do PO + EF1 for 3+ agents?
- Open: does EFX exist?

Things people have studied I

- Maximin fair share: similar to proportionality, but approximations exist.

> David Kurokawa, Ariel D Procaccia, and Junxing Wang. “Fair enough: Guaranteeing approximate maximin shares”. In: *Journal of the ACM (JACM)* 65.2 (2018), pp. 1–27

- PROP1, PROPX, EQ1, EQX.
- Items are arranged in a line; each bundle needs to be an interval.

> Vittorio Bilò, Ioannis Caragiannis, Michele Flammini, Ayumi Igarashi, Gianpiero Monaco, Dominik Peters, Cosimo Vinci, and William S Zwicker. “Almost envy-free allocations with connected bundles”. In: *Games and Economic Behavior* 131 (2022), pp. 197–221

Things people have studied II

- Best of both worlds: A lottery over EF1 allocations that is envy-free in expectation.

Haris Aziz, Rupert Freeman, Nisarg Shah, and Rohit Vaish. "Best of both worlds: Ex ante and ex post fairness in resource allocation". In: *Operations Research* (2023)

- Weighted version: each agent deserves goods in proportion to their weight.

Mithun Chakraborty, Ayumi Igarashi, Warut Suksompong, and Yair Zick. "Weighted envy-freeness in indivisible item allocation". In: *ACM Transactions on Economics and Computation (TEAC)* 9.3 (2021), pp. 1–39

- Some items are divisible.

Xiaohui Bei, Zihao Li, Jinyan Liu, Shengxin Liu, and Xinhang Lu. "Fair division of mixed divisible and indivisible goods". In: *Artificial Intelligence* 293 (2021), p. 103436

- We can give agents some money to stop envy.

Daniel Halpern and Nisarg Shah. "Fair division with subsidy". In: *Proceedings of the 12th International Symposium on Algorithmic Game Theory (SAGT)*. Springer. 2019, pp. 374–389