FAIRLY ALLOCATING INDIVISIBLE GOODS TO STRATEGIC AGENTS

Georgios Birmpas

Workshop on the Foundations of Modern AI

Αρχιμήδης

4 July 2024

Setting

- Set M of m indivisible goods
- Set N of n agents
- Each agent has an additive valuation function
  - v_{ij} = value derived by agent i for obtaining good j
  - $v_i(S) = \Sigma_{j \in S} v_{ij}$, for a set S of goods
- An allocation is a partition $S = (S_1, S_2, ..., S_n)$ of the set of goods

Focus

- We are interested in fair allocations
  - Each agent should think she got a fair share according to her own valuation function
  - Several fairness notions have been proposed
    - Our focus will be EF1

Solution Concepts

Envy-freeness (EF) [Foley '67, Varian '74]

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

Some issues

- This notion is "too strong" for indivisible goods

- No guarantee of existence
  - Consider instances with only one good

- Need to explore relaxations

Solution Concepts

Envy-freeness up to one 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 $g \in S_j$, such that
$v_i(S_i) \geq v_i(S_j \setminus \{g\})$

- i.e., for any agent who may envy agent *j*, there exists a good to remove from S_j and eliminate envy
  - Introduced as a concept by [Lipton et al. '04]
  - Formally defined by [Budish '11]

Algorithmic Setting

- Most of the related literature regards the algorithmic version of the problem
  - The agents are non-strategic
  - Given the true values of the agents, the goal is to design an algorithm that will produce fair outcomes

Algorithm Design

- In a nutshell
  - The agents submit their true values for the goods
  - An algorithm takes these values as an input
  - The algorithm outputs an allocation of the goods to the agents
  - The produced allocation needs to be fair according to the desired criterion

Results in the Algorithmic Setting

- EF1 allocations always exist
- They can be computed in polynomial time
- Easily achievable by very simple algorithms
  - Round Robin
- Ongoing research for other fairness notions, with many questions currently being open

Strategic Setting

- Setting introduction and,
- 4 of our results

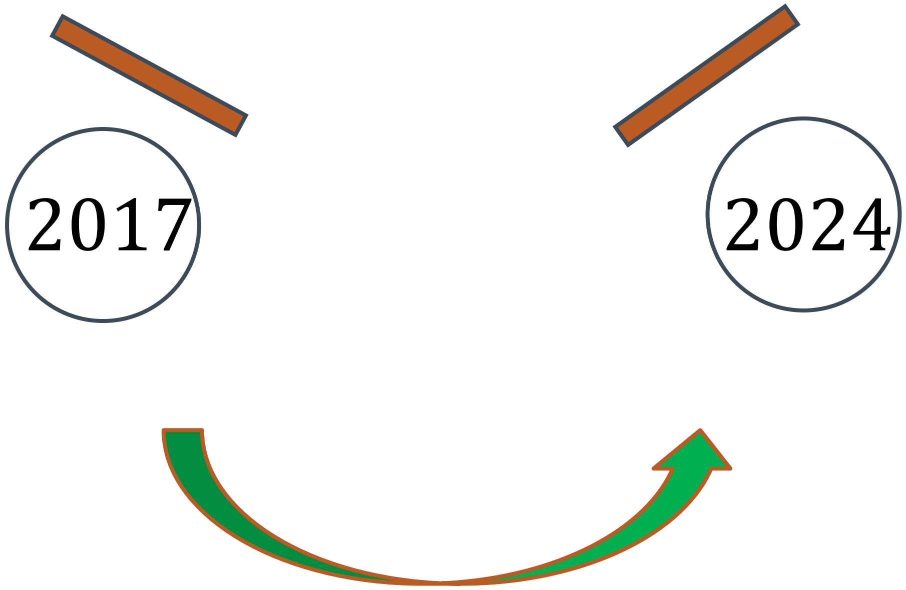

Strategic Setting

- The agents are strategic
- The utility of an agent i for a bundle $S \subseteq M$ of goods, is defined as the value that she has for this bundle

$$u_i(S) = v_i(S)$$

- There are no payments in this setting

Strategic Setting

- The agents are strategic
  - The goal of each agent is to maximize her own utility

  - An agent may misreport how she truly values the goods, if by doing so she ends up with a better bundle of goods

Strategic Setting

- The agents are strategic
  - The problem of producing fair allocations becomes even more challenging
    - The reports of the agents might not be true
    - At the same time, we desire the produced allocations to be fair according to the true values of the agents

Goal

- Design algorithms-mechanisms that are fair in their stable outcomes

Truthful Mechanism Design

- *The first obvious direction*: Design of Truthful mechanisms
- No agent can increase her utility by misreporting her true values
  - The latter is true regardless of the behavior of the other agents

Truthful Mechanism Design

- The first obvious direction: Design of Truthful mechanisms

- No agent can increase her utility by misreporting her true values
  - The latter is true regardless of the behavior of the other agents

- The obstacle of not having the true values as input is removed

Truthfulness and Fairness

- Can we have truthful mechanisms that are also fair?

Result 1

[Amanatidis, B., Christodoulou, Markakis] EC 17

- Unfortunately: The answer is no!

- Truthfulness and fairness are incompatible
  - There is no truthful mechanism that produces fair allocations under any meaningful fairness notion

PNE and Fairness

- The next natural question:
  - Is it possible to have non-truthful mechanisms whose equilibria define fair allocations?

PNE and Fairness

- We are interested in mechanisms that
  - Have PNE for every instance
  - Provide fairness guarantees at the allocations that correspond to these PNEs
    - According to the true values of the agents

Mechanisms and PNE

- Let $\boldsymbol{b_i} = (b_{i1}, b_{i2}, .., b_{im})$ to be the bidding vector of agent i for the goods in M
- Let A be an allocation mechanism and $\mathbf{b} = (\boldsymbol{b_1}, \dots, \boldsymbol{b_n})$ be a bidding profile of the agents in N
- We say that $\mathbf{b}$ is a PNE of A, if for every agent i we have that

$$v_i\big(A(\boldsymbol{b})\big) \geq v_i\big(A(\boldsymbol{b'_i}, \boldsymbol{b_{-i}})\big)$$

Result 2

[Amanatidis, B., Fusco, Lazos, Leonardi, Reiffenhauser] WINE 21, MOR 23

- Round Robin (produces EF1 allocations under non-strategic agents)
  - Has PNE for every valuation instance
    - [Aziz et al., 2017]
    - Our work
  - All of its PNE are also EF1 with respect to the true values of the agents

Round Robin

- The agents declare their bids for the goods
- Round Robin
  - Order the agents in an arbitrary way
  - For $i = 1$ to n give to each agent her favorite good
    - According to what she declared
  - Repeat step 2 until there are no more goods

Example

				
	5	4	3	2
	2	4	1	3

Example

				
	5	4	3	2
	2	4	1	3

Example

				
	5	4	3	2
	2	4	1	3

Example

				
	5	4	3	2
	2	4	1	3

Example

				
	5	4	3	2
	2	4	1	3

Example

				
	5	4	3	2
	2	4	1	3

Agent 1 gets a utility of 8
Agent 2 gets a utility of 7

Round Robin and EF1

- Property: If an agent declares her true values for the goods, then the produced allocation is EF1 for her (EF if she is agent 1)

Round Robin and EF1

- Property: If an agent declares her true values for the goods, then the produced allocation is EF1 for her (EF if she is agent 1)
- However, Round Robin is not truthful!

Truthful Reporting

				
	5	4	3	2
	2	4	1	3

Agent 1 gets a utility of 8
Agent 2 gets a utility of 7

Agent 1 Deviates

				
	5 4	4 5	3 3	2 2
	2	4	1	3

True

Reported

Agent 1 Deviates

				
	4	5	3	2
	2	4	1	3

Agent 1 Deviates

				
	4	5	3	2
	2	4	1	3

Agent 1 Deviates

				
	4	5	3	2
	2	4	1	3

Agent 1 Deviates

				
	4	5	3	2
	2	4	1	3

Utilities after the Deviation

				
	5	4	3	2
	2	4	1	3

Agent 1 gets a value of 9>8
Agent 2 gets a value of 4

Result 2

- Round Robin (produces EF1 allocations under non-strategic agents)
  - Has PNE for every valuation instance
    - [Aziz et al., 2017]
    - Our work
  - All of its PNE are also EF1 with respect to the true values of the agents

Warm Up: The Case of 2 Agents

- An allocation $(S_1, S_2, \ldots, S_n)$ is proportional, if for every agent i,

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

- *Fact:* EF ≡ Proportionality, when there are only 2 agents

Warm Up: The Case of 2 Agents

- Suppose that we have a PNE where an agent is not EF1

Warm Up: The Case of 2 Agents

- Suppose that we have a PNE where an agent is not EF1
- Say that this is agent 1
  - The argument for agent 2 is similar

Warm Up: The Case of 2 Agents

- If agent 1 is not EF1, she is also not EF

Warm Up: The Case of 2 Agents

- If agent 1 is not EF1, she is also not EF
- Fact: If she is not EF, she is also not proportional

Warm Up: The Case of 2 Agents

- If agent 1 is not EF1, she is also not EF
- Fact: If she is not EF, she is also not proportional
- Property: By declaring her true values agent 1 achieves EF

Warm Up: The Case of 2 Agents

- If agent 1 is not EF1, she is also not EF
- Fact: If she is not EF, she is also not proportional
- Property: By declaring her true values agent 1 achieves EF
  - Thus agent 1 achieves proportionality
  - This implies a higher value

Warm Up: The Case of 2 Agents

- If agent 1 is not EF1, she is also not EF
- Fact: If she is not EF, she is also not proportional
- Property: By declaring her true values agent 1 achieves EF
  - Thus agent 1 achieves proportionality
  - This implies a higher value
- Contradiction, this is not a PNE

The Case of n Agents

- Proportionality and EF are no longer identical
- We cannot use the same argument
- The problem becomes much more difficult

The Case of n Agents: Intuition

- The proof reduces to showing that the first agent views the final allocation as EF, when she bids a best response to other agents' bids
  - A PNE is a collection of best responses
  - Every agent can be seen as "agent 1" in the set of goods $M \setminus B$, where B is the set of goods lost in the first round, by the agents that precede this agent
    - This implies the EF1 guarantee for every agent

Round Robin Beyond Additive Agents

- What if we consider agents with more complex valuation functions?

Round Robin Beyond Additive Agents

- What if we consider agents with more complex valuation functions?
  - E.g., Submodular Valuation Functions
    - $f(S \cup \{j\}) - f(S) \geq f(T \cup \{j\}) - f(T)$
      - for any $S \subseteq T$, and $j \notin T$

Result 3

[Amanatidis, B., Lazos, Leonardi, Reiffenhauser] EC 23

- Surprisingly, the (approximate) PNE of Round Robin still provide fairness guarantees
  - Let $0 \leq \alpha \leq 1$. Any α-approximate PNE of Round Robin under submodular agents, corresponds to an $\alpha/3$-approximate EF1 allocation, according to the true values of the agents

Result 3

[Amanatidis, B., Lazos, Leonardi, Reiffenhauser] EC 23

- Surprisingly, the (approximate) PNE of Round Robin still provide fairness guarantees
  - Let $0 \leq \alpha \leq 1$. Any α-approximate PNE of Round Robin under submodular agents, corresponds to an $\alpha/3$-approximate EF1 allocation, according to the true values of the agents
    - This result is almost tight
    - We construct a 0.5-approximate PNE that guarantees 0.5-approximate EF1 Fairness

Approximate PNE

- Let $\boldsymbol{b_i} = (b_{i1}, b_{i2}, \ldots, b_{im})$ to be the bidding vector of agent i for the goods in M
- Let A be an allocation mechanism and $\mathbf{b} = (\boldsymbol{b_1}, \ldots, \boldsymbol{b_n})$ be a bidding profile of the agents in N
- We say that **b** is an α-approximate PNE of A, if for every agent i we have that

$$v_i\big(A(\boldsymbol{b})\big) \geq a \cdot v_i\big(A(\boldsymbol{b'_i}, \boldsymbol{b_{-i}})\big)$$

Approximate EF1

Envy-freeness up to one good (EF1)

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

Is Round Robin Perfect Then?

- Unfortunately, there are instances where no exact PNE exists. In particular…
  - For agents with submodular valuation functions, there are instances where no (¾+ε)-approximate PNE exists

Result 4

[Amanatidis, B., Lazos, Leonardi, Reiffenhauser] Preprint 24

- Although the notion of approximate PNE seems weak
  - In general, computing a strategy that provides even a 1+ε mutliplicative improvement cannot be done in polynomial time
    - This also applies to the 0.5-approximate PNE that we present

Future Directions

- Do other fair division algorithms (viewed as mechanisms) always have PNE?
- Is it possible to achieve stronger fairness guarantees in the strategic setting?

The End!