Fair Allocation of Indivisible Public Goods

BRANDON FAIN. etc

Definition of terms

- **Public goods** are those which can be enjoyed by multiple agents simultaneously.
- **Indivisible** means unable to be divided or separated.
  - Example of indivisible public goods: public roads, museums…

Example of public resource allocation

Suppose that the next time you vote, you see that there are four referendums for your consideration on the ballot, all of which concern the allocation of various public goods in your city: A = a new school, B = enlarging the public library, C = renovating the community college, and D = improving a museum. And you could only vote for only 2 of the four projects.

Hypothetical results of the example

A little more than half of the population voted for (A, B), a little less than half voted for (C, D), and every other combination received a small number of votes. Which projects should be funded? (Any thoughts?)

Problem definition

If we naïvely tally the votes, we would fund A and B, and ignore the preferences of a very large minority. In contrast, funding A and C seems like a reasonable compromise. It's true that it is impossible to satisfy all voters. But given a wide enough range of possible outcomes, perhaps we can find one that fairly reflects the preferences of large subsets of the population. (The main purpose of this paper)

Approach to solve the problem

- This paper models the public goods as elements with feasibility constraints on what subsets of elements can be chosen, and assume that agents have additive utilities across elements.

- This paper studies a groupwise fairness notion called the core, which generalizes well-studied notions of proportionality and Pareto efficiency, and requires that each subset of agents must receive an outcome that is fair relative to its size.

Public Goods Model

- A set of voters (or agents) N = [n]
- Sets of Public Goods: W, and m = |W|
- Set of feasible outcomes: $F \subseteq 2^W$
- Outcome c: $c \in F$
- The utility of agent i for element j : u_{ij}
- the utility of agent i under outcome $c \in F$ $u_i(\mathbf{c}) = \sum_{j \in \mathbf{c}} u_{ij}$
    - Assuming agents have additive utilities.

Constraints

Results of the model differ by the feasibility constraints imposed on the outcome. There are three types of constraints used in the model.

- Matroid Constraints.
  - In this setting, we are given a matroid M over the ground set W , and the feasibility constraint is that the chosen elements must form a basis of M
- Matching Constraints.
  - In this setting, the elements are edges of an undirected graph G(V, E), and the feasibility constraint is that the subset of edges chosen must form a matching.

Constraints Cont.

- Packing Constraints.
  - In this setting, we impose a set of packing constraints $Ax \leq b$, where $x_j \in \{0, 1\}$ is the indicator denoting whether element j is chosen in the outcome.

Remark:

Packing constraints capture the general Knapsack setting, in which there is a set of m items, each item j has an associated size s_j , and a set of items of total size at most B must be selected.

More definitions

- Pareto optimality
  - An outcome c satisfies Pareto optimality if there is no other outcome c′ such that u_i (c′) ≥ u_i (c) for all agents i ∈ N, and at least one inequality is strict.
- Proportional share
  - The proportional share of an agent i ∈ N is $Prop_i := V_i/n$. For β ∈ (0, 1], we say that an outcome c satisfies β-proportionality if u_i (c) ≥ β · $Prop_i$ for all agents i ∈ N. If β = 1, we simply say that c satisfies proportionality. (V_i is the maximal utility of an agent under all outcomes.)
- proportionality up to one issue
  - an outcome c of a public decision making problem satisfies proportionality up to one issue if for all agents i ∈ N, there exists an outcome c ′ that differs from c only on a single issue and u_i (c ′) ≥ Prop i

More definitions Cont.

- core outcome
  - Given an outcome c, we say that a set of agents $S \subseteq N$ form a blocking coalition if there exists an outcome c' such that $(|S|/n) \cdot u_i(c') \geq u_i(c)$ for all $i \in S$ and at least one inequality is strict. We say that an outcome c is a core outcome if it admits no blocking coalitions.

Remark: non-existence of blocking coalitions of size 1 is equivalent to proportionality, and non-existence of blocking coalitions of size n is equivalent to Pareto optimality. Therefore, a core outcome is both proportional and Pareto optimal.

More definitions Cont.

- (δ, α)-core outcome
  - For δ,α ≥ 0, an outcome c is a (δ, α)-core outcome if there exists no set of agents S ⊆ N and outcome c ′ such that for all ∈ S, and at least one inequality is strict.

$$\frac{|S|}{n} \cdot u_i(\mathbf{c}') \geq (1+\delta) \cdot u_i(\mathbf{c}) + \alpha$$

  - Remark: (0, 0)-core outcome is simply a core outcome. (δ, 0)-core outcome satisfies δ-proportionality.

Main focus on solving the problem.

- Finding integer allocations that approximate the core.
  - How to find a core?
    - Using Nash social welfare, which is the product of agent utilities.
  - Why finding the core.
    - Core is a generalization of proportionality and Pareto efficiency. And approximations of the core provide reasonable fairness guarantees in allocating public goods.

Integer Nash Welfare

- The integer Max Nash Welfare (MNW) solution is an outcome c that maximizes sum of $\ln u_i(c)$ for $i \in N$. More technically, if every integer allocation gives zero utility to at least one agent, the MNW solution first chooses a largest set S of agents that can be given non-zero utility simultaneously, and maximizes the product of utilities to agents in S.

Can INW approximate the core?

Consider an instance of public decision making with m issues and two alternatives per issue. Specifically, each issue t has two alternatives $\{a_1^t, a_2^t\}$, and exactly one of them needs to be chosen. There are two sets of agents X = {1, . . . ,m} and Y = {m + 1, . . . , 2m}. Every agent i ∈ X has $u_i^i(a_1^i) = 1$, and utility 0 for other alternatives. Every agent i ∈ Y has $u_i^t(a_1^t) = 1/m$ and $u_i^t(a_2^t) = 1$ for all issues i ∈ {1, 2, . . . ,m}.

Can INW approximate the core? Cont.

	a_1^1	a_2^1	a_1^2	a_2^2	$\dots$	a_1^m	a_2^m
$u_{1\in X}$	1	0	0	0	$\dots$	0	0
$u_{2\in X}$	0	0	1	0	$\dots$	0	0
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$		$\vdots$	$\vdots$
$u_{m\in X}$	0	0	0	0	$\dots$	1	0
$u_{i\in Y}$	$1/m$	1	$1/m$	1	$\dots$	$1/m$	1

Can INW approximate the core? Cont.

- The integer MNW outcome is c = $(a_1^1, a_1^2, \ldots, a_1^m)$ because any other outcome gives zero utility 1 to at least one agent. However, coalition Y can deviate to outcome c ′ = $(a_2^1, a_2^2, \ldots, a_2^m)$, and achieve utility m for each agent in Y. For c to be a (δ,α)-core outcome, we need

$$\exists i \in Y : (1+\delta) \cdot u_i(\mathbf{c}) + \alpha \geq \frac{|Y|}{|Y|+|X|} \cdot u_i(\mathbf{c}') \quad \Rightarrow \quad 1 + \delta + \alpha \geq \frac{m}{2}$$

Hence, c is not a (δ,α)-core outcome for any δ = o(m) and α = o(m). Therefore INW may not have no integral outcome in the core.

How to approximate the core?

- Using the constraints introduced before.
- Having a smooth Nash welfare objective function.
- Types of constraints:
  - Matroid constraints.
  - Matching constraints.
  - Packing constraints.

Smooth Nash Welfare

- Why Smooth Nash Welfare?
  - One issue with the Nash welfare objective is that it is sensitive to agents receiving zero utility. INW solution first chooses a largest set S of agents that can be given non-zero utility simultaneously, and maximizes the product of utilities to agents in S.
- What is Smooth Nash Welfare?

$$F(\mathbf{c}) := \sum_{i \in N} \ln\left(\ell + u_i(\mathbf{c})\right)$$

where *l* is >=0 as a parameter.

Remark: Showed in other papers, the objective function is NP-hard to optimize.

Influence of constraints

- Matroid constraints
  - The paper shows that with the correct choice of *l*, the that local search procedures for the smooth Nash welfare objective yield to a (0, 2)-core outcome for matroid constraints.
- Matching constraints
  - The paper shows that with the correct choice of *l*, the that local search procedures for the smooth Nash welfare objective yield to a (δ, k/δ)-core outcome for matroid constraints. For real number k greater than 0.

Influence of constraints Cont.

- Packing constraints
  - Optimizing any fixed smooth Nash welfare objective with packing constraints cannot guarantee a good approximation to the core

Conclusion

- The paper is about solving the fair allocation of public goods. The main way this paper address this problem is by calculating the approximate core, which provides reasonable fairness in the allocation of public goods.

- To calculate the approximate core, given that the current methods could not guarantee an integral outcome in the core, this paper presents algorithms with 3 types of constraints that could efficiently produce integral outcomes that are constant or near-constant approximations of the core.