Fair Division and Wagering: Two Equivalent Mechanism Classes

September 3, 2020

Denizalp Goktas

1 Fair Division

Fair Division is the problem of dividing one or several goods amongst two or more agents in a way that satisfies a suitable fairness criterion. We first describe a basic fair division model and then give a precise definition of a fair division mechanism.

1.1 Model

The fair division model consists of a set of agents $[n] = \{1, \dots, n\}$ and a set of goods/items $[m] = \{1, \dots, m\}$ with unit supply, i.e., there is only one unit of each good in the market. Each agent $i \in [n]$ is characterized by a weight e_i that determines the priority or entitlement of the agent i in the division problem (generally, a higher weight corresponds to being allocated a more desirable bundle) and a utility function $u_i : \mathbb{R}^m \to \mathbb{R}$ that gives the utility that an agent derives from a bundle of goods. We denote $\boldsymbol{e} = (e_1, e_2, ..., e_n)^T$

A divisible fair division mechanism, which we use when working with divisible goods, is a function $\mathcal{A} : (\mathbb{R}^m \to \mathbb{R})^n \times \mathbb{R}^n \to [0,1]^{n \times m}$ that takes as input the reported utility functions of the agents, $\boldsymbol{u}$ and weights $\boldsymbol{e}$, and outputs a matrix of **allocations** $\boldsymbol{X}$ s.t. $x_{ij} \geq 0$ denotes the fraction of the good j that agent i is allocated.

An indivisible fair division mechanism, which we use when working with indivisible goods, is a function $\mathcal{A} : (\mathbb{R}^m \to \mathbb{R})^n \times \mathbb{R}^n \to \{0,1\}^{n \times m}$ that takes as input the reported utility functions of the agents, $\boldsymbol{u}$ and weights $\boldsymbol{e}$, and outputs a matrix of **allocations** $\boldsymbol{X}$ s.t. $x_{ij} \in \{0,1\}$ denote whether or not agent i is allocated good j.

We denote the allocation of good j to agent i by the fair division mechanism $\mathcal{A}$ as $\mathcal{A}_{ij}(\boldsymbol{u}, \boldsymbol{e})$. We denote the bundle of goods allocated to agent i by the fair division mechanism $\mathcal{A}$ as a vector written as $\mathcal{A}_i(\boldsymbol{u}, \boldsymbol{e})$. We say that a fair division mechanism is valid iff $\mathcal{A}_{ij}(\boldsymbol{u}, \boldsymbol{e}) \geq 0$. That is a fair divison mechanism is valid if we do not "steal" non-existing good from the agents!

1.2 Design Goals

Fair division mechanisms can be designed with many goals in mind. When designing a fair division mechanism we can choose to satisfy certain properties and/or maximize certain welfare metrics.

1.2.1 Properties of Fair Division Mechanisms

In this section, we will go over a few terms to describe the properties of budget aggregation mechanisms.

Feasible: A fair division mechanism is feasible if $\forall j \in [m], \ \sum_{i \in i} \mathcal{A}_{ij}(\boldsymbol{u}, \boldsymbol{e}) \leq 1$. This means that a good is not allocated more than its available quantity.

Pareto-Dominance: A fair division mechanism $\mathcal{A}$ Pareto-dominates $\mathcal{A}'$ if $\forall i \in [n],\ \ u_i(\mathcal{A}_i(\boldsymbol{u}, \boldsymbol{e})) \geq u_i(\mathcal{A}_i'(\boldsymbol{u}, \boldsymbol{e}))$ and $\exists k \in [n],\ \ u_k(\mathcal{A}_k(\boldsymbol{u}, \boldsymbol{e})) > u_k(\mathcal{A}_k'(\boldsymbol{u}, \boldsymbol{e}))$. In other words, no agent is worse-off with mechanism $\mathcal{A}$, but some agent is better-off.

Pareto-Efficiency/Pareto-Optimality: A fair division mechanism is Pareto-Efficient/Pareto-Optimal, if there exists no feasible mechanism that Pareto-dominates it.

Proportionality: A fair division mechanism $\mathcal{A}$ is proportional iff it guarantees each agent will get at least $\frac{1}{n}$ of the utility they would had gotten if they were assigned all of the goods. That is, $\forall i \in [n],\ \ u_i(\mathcal{A}_i(\boldsymbol{u}, \boldsymbol{e})) \geq \frac{1}{n} u_i(\mathbf{1}_m)$

Envy-Free: A fair division mechanism $\mathcal{A}$ is envy free iff no agent desires the output bundle of another agent over their own bundle. That is, $\forall i, k \in [n],\ \ u_i(\mathcal{A}_i(\boldsymbol{u}, \boldsymbol{e})) \geq u_i(\mathcal{A}_k(\boldsymbol{u}, \boldsymbol{e}))$

Incentive-Compatibility: A fair division mechanism $\mathcal{A}$ is incentive compatible iff agents are better-off reporting their true utilities over any other utilities. Formally, let u_i be the true utility function of buyer i and let u_i' be any other utility function, then $\forall i \in [n],\ \ u_i(\mathcal{A}_i(\boldsymbol{u}, \boldsymbol{e})) \geq u_i(\mathcal{A}_i((\boldsymbol{u}_{-i}, u_i'), \boldsymbol{e}))$.

1.2.2 Welfare

One way to design a fair division mechanism is by taking a **welfarist** approach. That is, instead of looking at the allocation of agents, we compare only the utility vectors $\boldsymbol{u} \in \mathbb{R}^n$ associated with allocations. Note that in the preceding sentence, $\boldsymbol{u}$ is a vector of *utility values*, as opposed to our previous definition of it being a vector of *utility functions*. The welfarist approach can then allow us to come up with an objective function to maximize for our mechanism. [1]

Social Welfare Orderings: A Social Welfare Ordering (SWO) is a binary relation $\succeq$ over the space $\mathbb{R}^n$ of all utility vectors. This binary relationship must be transitive, reflexive, and complete. A social welfare ordering allows us to compare different utilities.

Most social welfare orderings can be represented using collective utility functions.

Collective Utility Functions: A Collective Utility Function (CUF) is a function $SW : \mathbb{R}^n \to \mathbb{R}$. For any vector $\boldsymbol{u} \in \mathbb{R}^n$, we call the value $SW(\boldsymbol{u})$, the social welfare of $\boldsymbol{u}$. Every CUF SW represents a SWO $\succeq$. In particular, if $SW(\boldsymbol{u}) \geq SW(\boldsymbol{u}')$ then $\boldsymbol{u} \succeq \boldsymbol{u}'$.

Note that any CUF can be represented as a SWO, but the converse is not true. All CUFs that we will be working with are **anonymous** (i.e., the output of the function SW is not dependent on the order of the utility values) and **unanimous** (i.e., if $\forall i \in [n],\ \ u_i > v_i$ then $SW(\boldsymbol{u}) > SW(\boldsymbol{v})$). We now present a few important CUFs. When designing a fair division mechanism we can pick any of the below CUFs as an objective to be maximized by our mechanism in order to achieve a desired outcome:

Utilitarian Welfare maps utilities to the weighted sum of inidividual utilities:

$$SW_{util}(\boldsymbol{u}) = \sum_{i \in [n]} e_i u_i \tag{1}$$

An agreement with maximal utilitarian social welfare is an agreement that maximises average utility, which explains why this may be considered an attractive social criterion. On the other hand, this definition of social welfare completely ignore fairness considerations:an allocation giving utility 101 to one agent and 0 to another would be considered socially superior to an allocation giving both of them utility 50.

Egalitarian Welfare maps to the minimum weighted utility value.

$$SW_{egal}(\boldsymbol{u}) = \min_{i \in [n]} \{e_i u_i\} \tag{2}$$

That is, maximising egalitarian social welfare amounts to raising the utility of the worst-offmember of society (whoever that may end up being) as much as possible

Nash Social Welfare maps to the weighted geometric mean of the utilities

$$SW_{nash}(\boldsymbol{u}) = \left(\prod_{i \in [n]} u_i^{e_i} \right)^{\frac{1}{\sum_{i \in [n]} e_i}} \tag{3}$$

The Nash CUF combines efficiency and fairness considerations. Like the utilitarian CUF, it favours high total utility. But at the same time it also encourages inequality-reducing transfers of utility. For example, the utilitarian CUF cannot distinguish between $\boldsymbol{u} = (4, 4), \boldsymbol{e} = (1, 1)$ and $\boldsymbol{u} = (2, 6), \boldsymbol{e} = (1, 1)$, while the Nash CUF will favour the former. Note that the above CUFs can be considered to be special cases of the CES family of CUFs, which take a parameter ρ:

$$SW_{CES}(\boldsymbol{u}) = \left(\sum_{i \in [n]} e_i u_i^{\rho} \right)^{\frac{1}{\rho}} \tag{4}$$

1. When $\rho = 1$ then the CES welfare function is exactly the utilitarian welfare function.

2. When $\rho \to -\infty$, the CES welfare function is exactly the egalitarian welfare function.

3. When $\rho \to 0$, the CES welfare function is exactly the Nash Social Welfare function.

Elitist Social Welfare maps to the maximum weighted utility value.

$$SW_{elit}(\boldsymbol{u}) = \max_{i \in [n]} \{e_i u_i\} \tag{5}$$

Elitist social welfare prioritizes one individual over all others.

1.3 Examples

One example of an indivisible fair division mechanism is **the random serial dictatorship mechanism**. The random serial dictatorship mechanism picks a random permutation of the n agents and then lets the agents successively choose an object in that order (so the first agent in the ordering gets first pick and so on). Once an item is picked by an agent, it is removed from the list of items that the rest of the agents can pick.

One example of a divisible fair division mechanism is the **market equilibrium mechanism**. The market equilibrium mechanism allocates each agent $i \in [n]$ with e_i units of currency and simulates a competitive equilibrium for the Fisher market $(\boldsymbol{u}, \boldsymbol{e})$ in which each good has a price and all agents spend their entire budget on goods that maximize their utility per price ratio.

2 Wagering

Wagering is the problem of eliciting the subjective beliefs of a group of agents. Given a group of agents that has differing beliefs about certain events, wagering mechanisms allow us to aggregate the beliefs of the agents about the outcome of these events. A wagering mechanism does this by collecting the reported beliefs and wagers of these agents and redistributing the total amount of money wagered across the agents that participated in the mechanism. One reason to create this type of mechanism is to truthfully elicit the agents' beliefs in order to predict the outcome of future event.

2.1 Model

The wagering model consists of a random variable X which takes values in $[m] = \{1, 2, ..., m\}$, and a set of n agents indexed by $[n] = \{1, 2, ..., n\}$. Every agent $i \in [n]$ has a vector of private, subjective beliefs, $\boldsymbol{p}_i = (p_{i1}, p_{i2}, ..., p_{im}) \in \mathbb{R}^m$, s.t. $p_{ij} \geq 0$ denotes the belief of agent i that event j will occur, and a wager $w_i \in \mathbb{R}_+$, which denotes the amount of money that agent i is betting. We denote $\boldsymbol{w} = (w_1, w_2, ..., w_n)^T$, and $W = \sum_{i\in[n]} w_i$. Each agent reports a vector of beliefs to the mechanism, which we denote by $\hat{\boldsymbol{p}}_i = (\hat{p}_{i1}, \hat{p}_{i2}, ..., \hat{p}_{im}) \in \mathbb{R}^m_+$. We denote all agents' reports as $\hat{\boldsymbol{P}} = (\hat{\boldsymbol{p}}_1, \ldots, \hat{\boldsymbol{p}}_n) \in \mathbb{R}^{n\times m}$.

A **wagering mechanism** $\mathcal{B} : \mathbb{R}^{n\times m} \times \mathbb{R}^n$ takes as inputs the reported beliefs of the agents $\hat{\boldsymbol{P}}$ and the wagers of the agents $\boldsymbol{w}$, along with the outcome x of the random variable X, and it outputs non-negative payments for each agent. We denote the payment of the mechanism $\mathcal{B}$ for agent i by $\mathcal{B}_i(\hat{\boldsymbol{P}}, \boldsymbol{w}, x) \in \mathbb{R}$ using mechanism $\mathcal{B}$. The **net profit** of agent i is defined as $\mathcal{B}_i(\hat{\boldsymbol{P}}, \boldsymbol{w}, x) - w_i$.

A wagering mechanism is said to be **valid** iff $\forall i \in [n], j \in [m], \hat{\boldsymbol{P}}, \boldsymbol{w}, \mathcal{B}_i(\hat{\boldsymbol{P}}, \boldsymbol{w}, j) \geq 0$. That is, a wagering mechanism is valid iff bettors do not lose more than their wagers

2.2 Design Goals

When designing wagering mechanisms, we want our mechanism to fulfill certain properties. We define some of the most common properties present in the literature.

Budget-balanced: $\mathcal{B}$ is budget-balanced if the market generates no profit or loss, or equivalently, for all $\hat{\boldsymbol{P}}, \boldsymbol{w}$, and x,

$$\sum_{i\in[n]} \mathcal{B}_i(\hat{\boldsymbol{P}}, \boldsymbol{w}, x) = \sum_{i\in[n]} w_i \tag{6}$$

$\mathcal{B}$ is **weakly budget balanced** if the above equal sign is a less than or equal sign, which means that the market generates no loss, but could generate a profit for the center/bookeeper.

Anonymous: $\mathcal{B}$ is if the payouts do not depend on the identities of the agents. Equivalently, for any permutation σ of the agents, and any x,

$$\mathcal{B}_i(\hat{\boldsymbol{P}}, \boldsymbol{w}, x) = \mathcal{B}_i([\hat{\boldsymbol{p}}_i]_{i\in\sigma}, [w_i]_{i\in\sigma}, x) \tag{7}$$

Incentive-compatible: $\mathcal{B}$ is incentive-compatible if agents maximize their expected payoff by reporting truthfully. In other words, for all $i \in [n]$, $\hat{\boldsymbol{P}}$ and **w**,

$$\mathbb{E}_{x\sim F_i}[\mathcal{B}_i((\boldsymbol{p}_i, \hat{\boldsymbol{p}}_{-i}), \boldsymbol{w}, x)] \geq \mathbb{E}_{x\sim F_i}[\mathcal{B}_i(\hat{\boldsymbol{P}}, \boldsymbol{w}, x)] \tag{8}$$

where F_i denotes the distribution of outcomes of the event X from the world view of agent i, that is a multinomial (more specifically m-nomial) distribution parametrized by the vector of probabilities $\frac{1}{\sum_{j\in[m]} p_{ij}}\hat{\boldsymbol{p}}_i$. $\mathcal{B}$ is **strictly incentive-compatible** if the above inequality is strict for all $\hat{\boldsymbol{p}}_i \neq \boldsymbol{p}_i$

Individually rational: $\mathcal{B}$ is individually rational if for any realization of the event X, each agent i has some report for which they have a non-negative expected net profit. In other words, for all $i \in [n], \boldsymbol{w}, \hat{\boldsymbol{p}}_{-i}$, and realizations of X, there exists some $\boldsymbol{p}_i^*$ such that

$$\mathbb{E}_{x\sim F_i}[\mathcal{B}_i((\boldsymbol{p}_i^*, \hat{\boldsymbol{p}}_{-i}, \boldsymbol{w}, x)] \geq w_i \tag{9}$$

where F_i denotes the distribution of outcomes of the event X from the world view of agent i, that is a multinomial (more specifically m-nomial) distribution parametrix by the vector of probabilities $\frac{1}{\sum_{j\in[m]} p_{ij}}\hat{\boldsymbol{p}}_i$. That is, an individually rational mechanism is one that ensures that from the perspective of the agent, there exists some report that makes the bidder lose no money in expectation. It is noteworthy that this definition is perhaps weaker than definition of individual rationality within the fair division settings.

2.3 Examples

Before we give an example of a wagering mechanism, we introduce the concept of a **proper scoring rule**. A proper scoring rule $s : \mathbb{R}^m \times [m] \to \mathbb{R}$ is a function that scores the reports of an agent such that:

$$\forall i \in [n], \hat{\boldsymbol{p}} \in \mathbb{R}^m \qquad \sum_{j\in[m]} p_{ij} s(\boldsymbol{p}_i, j) \geq \sum_{j\in[m]} p_{ij} s(\hat{\boldsymbol{p}}, j) \tag{10}$$

The **Weighted Score Wagering Mechanism** (WSWM) takes a strictly proper scoring rule s, and it defines the payout of agent i as:

$$\mathcal{B}(\hat{\boldsymbol{P}}, \boldsymbol{w}, x) = w_i \left(1 + s(\hat{\boldsymbol{p}}_i, x) - \frac{\sum_{k\in[n]} s(\hat{\boldsymbol{p}}_k, j) w_k}{\sum_{k\in i} w_k}\right) \tag{11}$$

A strictly proper scoring rule is any function from $\hat{\boldsymbol{p}}_i$ and x to a real number such that its expected value is maximized when agent i reports her subjective belief $\boldsymbol{p}_i$. Since WSWM's use a strictly proper scoring rule, they are always strictly incentive compatible. They are also budget-balanced, since the payouts will always add up to the sum of the weights. WSWM's are also anonymous, and they are individually rational.

The **Parimutuel wagering mechanism** is another example of a wagering mechanism. In this mechanism, agents report their beliefs $\hat{\boldsymbol{P}}$ and wagers $\boldsymbol{w}$ to the mechanism which runs the Eisenberg-Gale program to calculate an equilibrium $(\boldsymbol{\beta}, \boldsymbol{\pi})$ of the parimutuel market given by $(\hat{\boldsymbol{P}}, \boldsymbol{w})$. The mechanism then outputs for a given realization x of the event X the following payoffs:

$$\mathcal{B}_i(\hat{\boldsymbol{P}}, \boldsymbol{w}, x) = \frac{\beta_{ix}}{\pi_x} \tag{12}$$

3 Equivalence of the Two Classes

We now introduce an equivalence relationship that exists between fair division and wagering. This equivalene relationship allows us to derive wagering mechanisms using already existing fair division mechanisms and vice-versa.

We say that a divisible fair division mechanism $\mathcal{A}$ and a wagering mechanism $\mathcal{B}$ are **corresponding mechanisms** if:

$$\forall i \in [n], j \in [m], \boldsymbol{u}, \hat{\boldsymbol{P}}, \qquad \mathcal{B}_i(\hat{\boldsymbol{P}}, \boldsymbol{w}, j) = \mathcal{A}_{ij}(\boldsymbol{u}, \boldsymbol{w}) \sum_{i \in [n]} w_i \tag{13}$$

or equivalently

$$\forall i \in [n], \boldsymbol{u}, \hat{\boldsymbol{P}}, \qquad \mathcal{A}_i(\boldsymbol{u}, \boldsymbol{w}) = \frac{\mathcal{B}_i(\hat{\boldsymbol{P}}, \boldsymbol{w}, j)}{\sum_{i \in i} e_i} \tag{14}$$

Theorem 3.1. *A fair division mechanism $\mathcal{A}$ is valid iff its corresponding mechanism is also valid.*

This result can be seen by combining the definition of a valid mechanism for fair division and wagering with the definition of corresponding mechanisms.

References

[1] Ulle Endriss. Lecture notes on fair division. *CoRR*, abs/1806.04234, 2018.