

IEOR8100: Economics, AI, and Optimization

Lecture Note 9: Intro to Fair Division and Market Equilibrium

Christian Kroer*

March 6, 2020

1 Fair Division Intro

In this lecture note we start the study of *fair division* as well as *mechanism design*. In fair division, we have one or more goods that we wish to allocate to a set of agents. The goal will be to allocate the items in a manner that is efficient, while attempting to satisfy various notions of fairness towards each individual agent. Fair division has many applications such as assigning course seats to students, pilot-to-plane assignment for airlines, dividing estates, chores, or rent, and fair recommender systems.

In this note we study fair division problems with the following setup: we have a set of m infinitely-divisible goods that we wish to divide among n agents. Without loss of generality we may assume that each good has supply 1. We will denote the bundle of goods given to agent i as x_i , where x_{ij} is the amount of good j that is allocated to agent i . Each agent has some utility function $u_i(x_i) \in \mathbb{R}_+$ denoting how much they like the bundle x_i . We shall use x to denote an assignment of goods to agents. We will later study the indivisible setting.

Given all of the above, we would like to choose a “good” assignment x of goods to agents. However, “good” turns out to be very complicated in the setting of fair division, as there are many possible desiderata we may wish to account for.

First, we would like the allocation to somehow be efficient, meaning that it should lead to high utilities for the agents. One option would be to try to maximize the *social welfare* $\sum_i u_i(x_i)$, the sum of agent utilities. However, this turns out to be incompatible with the fairness notions that we will introduce later. An easy criticism of social welfare in the context of fair division is that it favors *welfare monsters*: agents with much greater capacity for utility are given more goods¹. Instead, we shall focus on the much less satisfying notion of *Pareto optimality*: we wish to find an allocation x such that for *every* other allocation x' , if some agent i' is better off under x' , then some other agent is strictly worse off. In other words, x should be such that no other allocation weakly improves all agent’s utilities, unless all utilities stay the same.

We will consider the following measures of how fair an allocation x is:

- *No envy*: x has no envy if for every pair of agents i, i' , $u_i(x_i) \geq u_i(x_{i'})$. In other words, every agent likes their own bundle at least as much as that of anyone else

*Department of Industrial Engineering and Operations Research, Columbia University. Email: christian.kroer@columbia.edu.

¹See also <https://existentialcomics.com/comic/8>

- *Proportionality*: x satisfies proportionality if $u_i(x_i) \geq u_i\left(\vec{1} \cdot \frac{1}{n}\right)$. That is, every agent likes their bundle x_i at least as well as the bundle where they receive $\frac{1}{n}$ of every item.

We begin our study of fair division mechanisms with a classic: *competitive equilibrium from equal incomes* (CEEI). In CEEI, we construct a mechanism for fair division by giving each agent a unit budget of fake currency (or *funny money*), computing what is called a competitive equilibrium (also known as *Walrasian equilibrium* or *market equilibrium*; we will use the latter terminology) under this new market, and using the corresponding allocation as our fair division. The fake currency is then thrown away, since it had no purpose except to define a market.

To understand this mechanism, we first introduce *market equilibrium*. In a market equilibrium, we wish to find a set of prices $p \in \mathbb{R}_+^m$ for each of the m goods, as well as an allocation x of goods to agents such that everybody is assigned an optimal allocation given the prices and their budget. Formally, the *demand set* of an agent i with budget B_i is

$$D(p) = \operatorname{argmax}_{x_i \geq 0} u_i(x_i) \text{ s.t. } \langle p, x_i \rangle \leq B_i$$

A market equilibrium is an allocation-price pair (x, p) s.t. $x_i \in D(p)$ for all agents i , and $\sum_i x_{ij} = 1$.

CEEI is a perfect solution to our desiderata that we asked for. It is Pareto optimal (every market equilibrium is Pareto optimal by the *first welfare theorem*). It has no envy: since each agent has the same budget $B_i = 1$ in CEEI and every agent is buying something in their demand set, no envy must be satisfied, since they can afford the bundle of any other agent. Finally, proportionality is satisfied, since each agent can afford the bundle where they get $\frac{1}{n}$ of each good (convince yourself why).

Market-equilibrium-based allocation for divisible items has applications in large-scale Internet markets. First, it can be applied in *fair recommender systems*. As an example, consider a job recommendations site. It's a two-sided market. On one side are the users, whom view job ads. On the other side are the companies creating job ads. Naively, a system might try to simply maximize the number of job ads that users click on, or apply to. This can lead to extremely imbalanced allocations, where a few job ads get a huge number of views and applicants, which is bad both for users and the companies. Instead, the system may wish to fairly distribute user views across the many different job ads. In that case, CEEI can be used. In this setting the agents are the job ads, and the items are slots in the ranked list of job ads shown to the user. Secondly, there are strong connections between market equilibrium and the allocation of ads in large-scale Internet ad markets. This connection will be explored in detail in a later note.

Motivated by the application to fair division, we will now cover market equilibrium, both when they exist and how to find one when they do.

2 Fisher Market

We first study market equilibrium in the *Fisher market* setting. We have a set of m infinitely-divisible goods that we wish to divide among n buyers. Without loss of generality we may assume that each good has supply 1. We will denote the bundle of goods given to buyer i as x_i , where x_{ij} is the amount of good j that is allocated to buyer i . Each buyer has some utility function $u_i(x_i) \in \mathbb{R}_+$ denoting how much they like the bundle x_i . We shall use x to denote an assignment of goods to buyers. Each buyer is endowed with a budget B_i of currency.

2.1 Linear Utilities

We start by studying the simplest setting, where the utility of each buyer is linear. This means that every buyer i has some valuation vector $v_i \in \mathbb{R}^m$, and $u_i(x_i) = \langle v_i, x_i \rangle$.

Amazingly, there is a nice convex program for computing a market equilibrium. Before giving the convex program, let's consider some properties that we would like. First, if we are going to find a feasible allocation, we would obviously like the supply constraints to be respected, i.e.

$$\sum_i x_{ij} \leq 1, \forall j.$$

Secondly, since a buyer's demand does not change even if we rescale their valuation by a constant, we would like the optimal solution to our convex program to also remain unchanged. Similarly, splitting the budget of a buyer into two separate buyers with the same valuation function should leave the allocation unchanged. These conditions are satisfied by the budget-weighted geometric mean of the utilities:

$$\left(\prod_i u_i(x_i)^{B_i} \right)^{1/\sum_i B_i}.$$

Since taking roots does not affect optimality, and taking the log of the whole expression, this is equivalent to optimizing

$$\begin{array}{ll} \max_{x \geq 0} & \sum_i B_i \log \langle v_i, x_i \rangle & \text{Dual variables} \\ \text{s.t.} & \sum_i x_{ij} \leq 1, \quad \forall j = 1, \dots, m, & \left| \begin{array}{l} p_j \end{array} \right. \end{array} \quad (\text{EG})$$

On the right are the dual variables associated to each constraint. It is easy to see that this is a convex program. First, the feasible set is defined by linear inequalities. Second, we are taking a max of a sum of concave functions composed with linear maps. Since taking a sum and composing with a linear map both preserve concavity we get that the objective is concave.

The solution to the primal problem x along with the vector of dual variables p yields a market equilibrium. Here we assume that for every item j there exists i such that $v_{ij} > 0$, and every buyer values at least one good above 0.

Theorem 1. *The pair of allocations x and dual variables p from EG forms a market equilibrium.*

Proof. To see this, we need to look at the KKT conditions of the primal and dual variables. Writing the Lagrangian relaxation gives

$$\begin{aligned} & \min_{p \geq 0} \max_{x \geq 0} \sum_i B_i \langle v_i, x_i \rangle + \sum_j p_j \left(1 - \sum_i x_{ij} \right) \\ &= \min_{p \geq 0} \max_{x \geq 0} \sum_i [B_i \langle v_i, x_i \rangle - \langle p, x_i \rangle] + \sum_j p_j \end{aligned} \quad (1)$$

Looking at optimality conditions we get

1. $p_j > 0 \Rightarrow \sum_i x_{ij} = 1$
2. $\frac{B_i}{\langle v_i, x_i \rangle} \leq \frac{p_j}{v_{ij}}$
3. $x_{ij} > 0 \Rightarrow \frac{B_i}{\langle v_i, x_i \rangle} = \frac{p_j}{v_{ij}}$

The first condition shows that every item is fully allocated, since for every j there is some buyer i with non-zero value and by the second condition $p_j \geq \frac{v_{ij}B_i}{\langle v_i, x_i \rangle} > 0$.

The second condition for market equilibrium is that every buyer is assigned a bundle from their demand set. We will use $\beta_i = \frac{B_i}{\langle v_i, x_i \rangle} = \frac{B_i}{u_i(x_i)}$ to denote the *utility price* that buyer i pays. First off, by the second condition we have that the utility price that buyer i gets satisfies

$$\beta_i \leq \frac{p_j}{v_{ij}}.$$

By the third condition, we have that if $x_{ij} > 0$ then for all other items j' we have

$$\frac{p_j}{v_{ij}} = \beta_i \leq \frac{p_{j'}}{v_{ij'}}.$$

Thus, any item j that buyer i is assigned has at least as good utility price as any other item j' . In other words, they only buy items that have the best bang-per-buck among all the items. Thus we get that they only purchase optimal items, it remains to show that they spent their whole budget. Multiplying the third condition by x_{ij} and rearranging gives

$$x_{ij}v_{ij} \frac{B_i}{\langle v_i, x_i \rangle} = p_j x_{ij},$$

for any j such that $x_{ij} > 0$. Summing across all such j yields

$$\sum_j p_j x_{ij} = \sum_j x_{ij} v_{ij} \frac{B_i}{\langle v_i, x_i \rangle} = \langle v_i, x_i \rangle \frac{B_i}{\langle v_i, x_i \rangle} = B_i.$$

□

EG gives us an immediate proof of existence for the linear Fisher market setting: the feasible set is clearly non-empty, and the max is guaranteed to be achieved.

In a previous lecture note we referenced Pareto optimality as a property of market equilibrium. It is now trivial to see that Pareto optimality holds in Fisher-market equilibrium: since it is a solution to EG, it must be. Otherwise we construct a solution with strictly better objective!

From the EG formulation we can also see that the equilibrium utilities and prices are in fact unique. First note that any market equilibrium allocation would satisfy the optimality conditions of EG, and thus be an optimal solution. But if there were more than one set of utility vectors that were equilibria, then by the strong concavity of the log we would get that there is a strictly better solution, which is a contradiction. That equilibrium prices are unique now follows from the third optimality condition, since all terms except the utilities are constants.

3 More General Utilities

It turns out that EG can be applied to a broader class of utilities. This class is the set of utilities that are concave, homogeneous, and continuous.

In that case we get an optimization problem of the form

$$\begin{array}{ll} \max_{x \geq 0} & \sum_i B_i \log u_i(x_i) & \text{Dual variables} \\ \text{s.t.} & \sum_i x_{ij} \leq 1, \quad \forall j = 1, \dots, m, & \left| \begin{array}{l} p_j \end{array} \right. \end{array} \quad \text{(EG)}$$

This is still a convex optimization problem, since composing a concave and nondecreasing function (the log) with a concave function (u_i) yields a concave function.

Beyond linear utilities, the most famous classes of utilities that fall under this category is:

1. Cobb-Douglas utilities: $u_i(x_i) = \prod_j (x_{ij})^{a_{ij}}$, where $\sum_j a_{ij} = 1$
2. Leontief utilities: $u_i(x_i) = \min_j \frac{x_{ij}}{a_{ij}}$
3. The family of constant elasticity of substitution (CES) utilities: $u_i(x_i) = \left(\sum_j a_{ij} x_{ij}^\rho\right)^{1/\rho}$, where a_{ij} are the utility parameters of a buyer, and ρ parameterizes the family, with $-\infty < \rho \leq 1$ and $\rho \neq 0$

CES utilities turn out to generalize all the other utilities we have seen so far: Leontief utilities are obtained as ρ approaches $-\infty$, Cobb-Douglas utilities as ρ approaches 0, and linear utilities when $\rho = 1$. More generally, $\rho < 0$ means that items are complements, whereas $\rho > 0$ means that items are substitutes.

If u_i is continuously differentiable then the proof that EG computes a market equilibrium in this more general setting essentially follows that of the linear case. The only non-trivial change is that when we derive optimality conditions by taking the derivative of the Lagrangian with respect to x_i we get

1. $\frac{B_i}{u_i(x_i)} \leq \frac{p_j}{\partial u_i(x_i)/\partial x_{ij}}$
2. $x_{ij} > 0 \Rightarrow \frac{B_i}{u_i(x_i)} = \frac{p_j}{\partial u_i(x_i)/\partial x_{ij}}$

In order to prove that buyers spend their budget exactly in this setting we can apply Euler's homogeneous function theorem $u_i(x_i) = \sum_j x_{ij} \frac{\partial u_i(x_i)}{\partial x_{ij}}$ to get

$$\sum_j x_{ij} p_j = \sum_j x_{ij} \frac{\partial u_i(x_i)}{\partial x_{ij}} \frac{B_i}{u_i(x_i)} = B_i.$$

4 Computing Market Equilibrium

So now we know how to write a market equilibrium problem as a convex program. How should we solve it? One option is to build the EG convex program explicitly using mathematical programming software. However, most contemporary software is not very good at handling this kind of objective function (formally this falls under exponential cone programming, which is still relatively new). As of 2019, my experience with open-source solvers was that they fail at 150 items and 150 buyers, with randomly-generated valuations. The Mosek solver [15] is currently the only industry-grade solver that supports exponential cone programming (support for exponential cones was only just added in 2018). It fares much better, and scales to a few thousand buyers and items. For problems of moderate-to-large size, this is the most effective approach. However, for very large instances, the iterations of the interior-point solver used in Mosek become too slow.

Instead, for extremely large problems we may invoke some of our earlier results on saddle-point problems. In particular, the formulation (1) is amenable to online mirror descent and the folk-theorem based approach for solving saddle-point problems. In that framework, we can interpret the repeated game as being played between a pricer trying to minimize over prices p , and the set of buyers choosing allocations x .

As an exercise, convince yourself that the OMD/folk theorem approach works. Pay particular attention to the assumptions needed for online mirror descent.

5 Historical Notes

The original Eisenberg-Gale convex program was given for linear utilities in [10]. [9] later extended it to utilities that are concave, continuous, and homogeneous.

Fairly assigning course seats to students via market equilibrium was studied by Budish [4]. Goldman and Procaccia [12] introduce an online service `spliddit.org` which has a user-friendly interface for fairly dividing many things such as estates, rent, fares, and others. The motivating example of fair recommender systems, where we fairly divide impressions among content creators via CEEI was suggested in [14] and [13]. Similar models, but where money has real value, were considered for ad auctions with budget constraints by several authors [3, 7, 8]

There is a rich literature on various iterative approaches to computing market equilibrium in Fisher markets. One can apply first-order methods or regret-minimization approaches to the saddle-point formulation (1) directly, which was done in Kroer et al. [14] and Gao et al. [11]. There is a large literature on interpreting first-order methods through the lens of dynamics between a pricer who increases and decreases prices as goods become oversubscribed and undersubscribed, and buyers report their preferred bundles, or make gradient-steps in the direction of their preference [6, 2, 5]. There is also a literature deriving auction-like algorithms, which can similarly sometimes be viewed as instantiations of gradient descent and related algorithms [1, 16].

References

- [1] Xiaohui Bei, Jugal Garg, and Martin Hoefer. Ascending-price algorithms for unknown markets. *ACM Transactions on Algorithms (TALG)*, 15(3):1–33, 2019.
- [2] Benjamin Birnbaum, Nikhil R Devanur, and Lin Xiao. Distributed algorithms via gradient descent for fisher markets. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 127–136. ACM, 2011.
- [3] Christian Borgs, Jennifer Chayes, Nicole Immorlica, Kamal Jain, Omid Etesami, and Mohammad Mahdian. Dynamics of bid optimization in online advertisement auctions. In *Proceedings of the 16th international conference on World Wide Web*, pages 531–540, 2007.
- [4] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- [5] Yun Kuen Cheung, Richard Cole, and Nikhil R Devanur. Tatonnement beyond gross substitutes? gradient descent to the rescue. *Games and Economic Behavior*, 2019.
- [6] Richard Cole and Lisa Fleischer. Fast-converging tatonnement algorithms for one-time and ongoing market problems. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 315–324, 2008.
- [7] Vincent Conitzer, Christian Kroer, Eric Sodomka, and Nicolás E Stier-Moses. Multiplicative pacing equilibria in auction markets. In *International Conference on Web and Internet Economics*, 2018.
- [8] Vincent Conitzer, Christian Kroer, Debmalya Panigrahi, Okke Schrijvers, Eric Sodomka, Nicolas E Stier-Moses, and Chris Wilkens. Pacing equilibrium in first-price auction markets. In *Proceedings of the 2019 ACM Conference on Economics and Computation*. ACM, 2019.
- [9] Edmund Eisenberg. Aggregation of utility functions. *Management Science*, 7(4):337–350, 1961.

- [10] Edmund Eisenberg and David Gale. Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics*, 30(1):165–168, 1959.
- [11] Yuan Gao, Christian Kroer, and Donald Goldfarb. Increasing iterate averaging for solving saddle-point problems. *arXiv preprint arXiv:1903.10646*, 2019.
- [12] Jonathan Goldman and Ariel D Procaccia. Spliddit: Unleashing fair division algorithms. *ACM SIGecom Exchanges*, 13(2):41–46, 2015.
- [13] Christian Kroer and Alexander Peysakhovich. Scalable fair division for ‘at most one’ preferences. *arXiv preprint arXiv:1909.10925*, 2019.
- [14] Christian Kroer, Alexander Peysakhovich, Eric Sodomka, and Nicolas E Stier-Moses. Computing large market equilibria using abstractions. *arXiv preprint arXiv:1901.06230*, 2019.
- [15] APS Mosek. The mosek optimization software. *Online at <http://www.mosek.com>*, 54(2-1):5, 2010.
- [16] Yurii Nesterov and Vladimir Shikhman. Computation of fisher-gale equilibrium by auction. *Journal of the Operations Research Society of China*, 6(3):349–389, 2018.