CS 290 Market Design | **Fall 2020**

Lecture 2

Lecturer: Vijay Vazirani | *Scribe: Karthik Gajulapalli, Will Overman*

2.1 Markets with ordinal preferences and no initial endowment

In this section markets where each agent has an ordinal preference list and no initial endowment. We introduce some new definitions we will need later.

Definition 2.1. ***Stochastic dominance**: Let i be some agent, with a preference list $\pi(G)$ over goods G. Consider two allocation vectors x, y over goods ordered by $\pi(G)$ then x dominates y for agent i, if $\forall k \in \{1, ..., n\}$ we have :*

$$\sum_{j=1}^{k} x_j \geq \sum_{j=1}^{k} y_j$$

Definition 2.2. ***Envy-Free (Cardinal):** For an agent i, let $(u_{i1}, u_{i2}, ..., u_{in})$ be its utility vectors over all goods $g_1, ...g_n$. Let μ be an allocation where each agent i gets an allocation vector x_i over the goods. Then μ is envy-free if $\forall\ k\ \neq i$:*

$$\sum_{j=1}^{n} u_{ij} x_{ij} \ \geq \ \sum_{j=1}^{n} u_{ij} x_{kj}$$

Definition 2.3. ***Envy-Free (Ordinal):** An allocation μ is envy-free if for all agents i and their allocation x_i, x_i stochastically dominates allocations of all other agents, x_j, under i's preference list.*

The first mechanism we consider for this setting is *priority rule*. The mechanism picks an arbitrary ordering of agents, and according to that ordering assigns the agents its favorite good not assigned yet.

PRIORITY RULE:
Input: $(A, G, \succ)$
Output: Perfect Matching μ

1. $L = A$
2. Choose a fixed ordering of the agents $1, ..., n$
3. For $i = 1$ to n:
    (a) Find o_i, the most preferred object in L with respect to agent i's preference list
    (b) Set $\mu(i) = o_i$
    (c) $L \leftarrow L \setminus \{o_i\}$
4. Return μ.

Lemma 2.4. *Priority Rule is Pareto Optimal.*

Proof. Let $\pi = g_1 ... g_n$ be the permutation on the set of goods that is returned by the priority mechanism, that is, $\mu(i) = g_i$. Now assume for the sake of contradiction that there is another allocation μ' with corresponding permutation π' such that $\forall i, \mu'(i) \succeq_i \mu(i)$ and $\exists i, \mu'(i) \succ_i \mu(i)$. Consider the first i from 1 to n for which this is the case. Then π and π' are the same up to this point, so then $\mu'(i)$ must have been available when it was i's turn in the run of the algorithm that gave μ, which contradicts the fact that i chooses his most preferred available object. □

Lemma 2.5. *Priority Rule is strategy-proof, but not envy-free .*

Proof. Strategy-proof follows from the fact that each agent chooses the best object available to her and misreporting her preferences would only lead to choosing a less desirable object. To see that it is not envy-free note the case where two agents desire the same object, the agent higher in the ordering will get the object and the agent lower will be envious of that agents allocation. □

RANDOM PRIORITY:
Input: $(A, G, \succ)$
Output: Fractional Perfect Matching μ

1. Choose a fixed ordering of the agents $1, ..., n$
2. Initialize an array μ of size $n \times n$ with all values set to 0
3. For each bijection $\sigma : \{1, ..., n\} \to \{1, ..., n\}$
   (a) Run **Priority Rule** with this bijection as the chosen ordering, obtaining assignment μ_σ
   (b) For $i = 1$ to n, set $\mu[i, \mu_\sigma(i)] \leftarrow \mu[i, \mu_\sigma(i)] + \frac{1}{n!}$
4. Return μ.

Now we can use stochastic dominance to discuss a notion of Pareto optimality in this setting.

Definition 2.6. *We say that an allocation x is **Pareto optimal** if for all other allocations y:*

$$\forall\, i : y^i \succeq_i x^i \Rightarrow y = x.$$

where x^i and y^i represent the fractional allocation of i, and $\succeq_i$ is stochastic domination.

Remark 2.7. *Two allocations can be incomparable if some agents prefer one allocation, and other agents prefer the other. However if an allocation is pareto, no other allocation can dominate it.*

The final algorithm for this setting is known as *Probabilistic Serial*. The algorithm simulates the following process: The n agents have buckets of 1 unit volume, and the "items" are kept as a liquid in n individual containers of 1 unit volume. At the start of the algorithm, the agents go to the container of their most preferred item, and the liquid of that item flows at a rate of 1 unit per hour to each of the agents at that container. At the end of the hour, the agents then go to their next preferred item's container until that one also runs out, and so on. Once all of the liquid runs out each agent will be left with a full bucket of 1 unit volume where the fractions of each liquid in the bucket corresponds to the fraction of each item they receive.

PROBABILISTIC SERIAL:
Input: $(A, G, \succ)$
Output: Fractional Perfect Matching μ

1. $G_Q \rightarrow$ Priority Queue of Goods ordered by goods most desired by agents.
2. While $G_Q \neq \emptyset$
   (a) $g^* = G_Q$.pop() (remove most desired good).
   (b) $A_{g^*} = |\{i \in A \mid i\text{'s favorite remaining good is } g^*\}|$
   (c) $c(g)$ = remaining quantity left of good g for all g.
   (d) Each agent is allocated $\frac{1}{A_{g^*}} c(g)$ of good g it most desires in μ.
   (e) Remove all goods from G_Q where $c(g) = 0$.
3. Return μ.

We provide an example where Random Priority and Probabilistic Serial are Pareto Incomparable in the sense that different agents can prefer the allocation given by one of the algorithms more than the other.

Example 2.8. *Let the preferences of the agents be as follows:*

$$\begin{array}{lcccc} 1: & a & b & c & d \\ 2: & a & b & d & c \\ 3: & b & a & c & d \\ 4: & c & d & a & b \end{array}$$

Running Random Priority will return the following allocations

Agent	a	b	c	d
1	1/2	1/6	1/12	1/4
2	1/2	1/6	0	1/3
3	0	2/3	1/12	1/4
4	0	0	5/6	1/6

And running Probabilistic Serial will return the following allocations

Agent	a	b	c	d
1	1/2	1/6	1/9	2/9
2	1/2	1/6	0	1/3
3	0	2/3	1/9	2/9
4	0	0	7/9	2/9

Agents 1 and 3 prefer the Probabilistic Serial allocation, while agent 4 prefers the Random Priority Allocation. Agent 2 is indifferent to both.

We now show an example where Random-Priority is not envy-free.

Example 2.9. *Let the preferences of the agents be as follows:*

$$\begin{array}{llll} 1: & a & b & c \\ 2: & b & a & c \\ 3: & b & c & a \end{array}$$

Running Random Priority will return the following allocations

Agent	a	b	c
1	5/6	0	1/6
2	1/6	3/6	2/6
3	0	3/6	3/6

And running Probabilistic Serial will return the following allocations

Agent	a	b	c
1	3/4	0	1/4
2	1/4	1/2	1/4
3	0	1/2	1/2

The allocation over goods a, b in the Random Priority allocation for agents 1 and 2 is 5/6 and 4/6, respectively. Agent 2 is thus envious of agent's 1 allocation with respect to agent 2's preference list. One can verify that for this example, however, that Probabilistic-Serial is envy-free.

Lemma 2.10. *PS is envy-free*

Proof. At every timestep an agent gets the good he most desires. Since this good is being distributed evenly with everyone else who desires that good, no other agent can get more of this good during this timestep, and so the allocation is envy-free. □

We now show an example where RP is not Pareto-Optimal

Example 2.11. *Let the preferences of the agents be as follows:*

$$\begin{array}{lllll} 1: & a & b & c & d \\ 2: & a & b & c & d \\ 3: & b & a & d & c \\ 4: & b & a & d & c \end{array}$$

Running RP, we get the following allocation:

Agent	a	b	c	d
1	5/12	1/12	5/12	1/12
2	5/12	1/12	5/12	1/12
3	1/12	5/12	1/12	5/12
4	1/12	5/12	1/12	5/12

But the following allocation is better for all particpants:

Agent	a	b	c	d
1	1/2	0	1/2	0
2	1/2	0	1/2	0
3	0	1/2	0	1/2
4	0	1/2	0	1/2

Lemma 2.12. *PS is Pareto optimal*

Proof. For each time step, i, let the net fractional allocation at that time be t_i. Then for any time step j, there is no allocation of goods where each agent gets exactly t_j goods and is Pareto better than the allocation of Probabilistic Serial at time j.
We show this by induction. As a base case at time 0, the empty allocation is Pareto. Let us consider the allocation at time $k+1$, let A_{k+1} be the allocation by PS, and by contradiction let us assume that there is a Pareto better allocation P_{k+1}. Look at the fractional allocation $\alpha_k = t_{k+1} - t_k$. By the inductive hypothesis A_{k+1} and P_{k+1} are the same upto t_k. For any agent A_{k+1} gives each agent α_k of their most desirable good. If P_{k+1} is better it must be giving some agent a good better good than its most desirable good, a contradiction. □

Lemma 2.13. *RS is strategy proof*

Proof. From Lemma 2.5, we know that at every run it must be strategy-proof. So lying at any run will only result in a worse allocation.

□

We show an example where PS is not strategy-proof

Example 2.14. *Let the preferences of the agents be as follows:*

$$\begin{array}{llll} 1: & a & b & c \\ 2: & a & c & b \\ 3: & b & a & c \end{array}$$

Then the PS will return the following allocation:

Agent	a	b	c
1	1/2	1/4	1/4
2	1/2	0	1/2
3	0	3/4	1/4

However if Agent 3 lies, and reports its preference list as a, b,c, the following allocation ensues:

Agent	a	b	c
1	1/3	1/2	1/6
2	1/3	0	4/6
3	1/3	1/2	1/6

By lying Agent 3 gets 5/6 of his favorite two goods, while he only gets 3/4 when is being honest.

Definition 2.15. *A mechanism is weakly strategy-proof if no agent can lie about her true preferences in order to obtain an allocation that stochastically dominates their true allocation.*

Lemma 2.16. *PS is weakly strategy-proof.*

Proof. Assume agent i is able to lie about her preferences to obtain a better allocation y^i than the allocation x^i he obtains in the true run. Consider the vectors x^i and y^i and let k be the first coordinate at which they differ. Then according to PS agent i will obtain the largest amount of good k possible, so the only way in which x and y can differ is for y to obtain less of good k, which implies we cannot have that y stochastically dominates x. $\square$

2.2 Exercises

Consider a setting where instead of using stochastic dominance to compare allocations we compare them lexicographically. An allocation $x^i = (x_1, ...x_n)$ for an agent i is lexicographically better than an allocation $y^i = (y_1, ..., y_n)$ if on ordering the goods by i's preference list the first k where the allocation differs $x_k > y_k$. Show that PS is strategy-proof for lexicographic preferences.

(Hint: Show and use the fact that for any false preference list given by agent i there is some good b that is *sacrificed* in the sense that i obtains less of b in the false run than in the true run.)