

CSC2556

Lecture 8

Leximin
Rent Division

Leximin (DRF)

Computational Resources

- **Resources:** Homogeneous divisible resources like CPU, RAM, or network bandwidth
- **Valuations:** Each player wants the resources in a **fixed proportion** (Leontief preferences)
- **Example:**
 - Player 1 requires (2 CPU, 1 RAM) for each copy of task
 - Indifferent between (4,2) and (5,2), but prefers (5,2.5)
 - “fractional” copies are allowed


Model

- Set of **players** $N = \{1, \dots, n\}$
- Set of **resources** R , $|R| = m$
- **Demand** of player i is $d_i = (d_{i1}, \dots, d_{im})$
 - $0 < d_{ir} \leq 1$ for every r , $d_{ir} = 1$ for some r
 - “For every 1% of the total available CPU you give me, I need 0.5% of the total available RAM”
- **Allocation**: $A_i = (A_{i1}, \dots, A_{im})$ where A_{ir} is the fraction of available resource r allocated to i
 - Utility to player i : $u_i(A_i) = \min_{r \in R} A_{ir} / d_{ir}$.
 - We’ll assume a **non-wasteful** allocation
 - Allocates resources proportionally to the demand.



Dominant Resource Fairness

- **Dominant resource** of i is r such that $d_{ir} = 1$
- **Dominant share** of i is A_{ir} , where $r =$ dominant resource of i
- Dominant Resource Fairness (**DRF**) Mechanism
 - Allocate maximal resources while maintaining equal dominant shares.

DRF animated



Total



Properties of DRF

- **Envy-free:** $u_i(A_i) \geq u_i(A_j), \forall i, j$
 - Why? [Note: EF no longer implies proportionality.]
- **Proportionality:** $u_i(A_i) \geq 1/n, \forall i$
 - Why?
- **Pareto optimality** (Why?)
- **Group strategyproofness:**
 - If a group of players manipulate, it can't be that none of them lose, and at least one of them gains.
 - We'll skip this proof.

The Leximin Mechanism

- Generalizes the DRF Mechanism
- Mechanism:
 - Choose an allocation A that
 - Maximizes $\min_i u_i(A_i)$
 - Among all minimizers, breaks ties in favor of higher second minimum utility.
 - Among all minimizers, breaks ties in favor of higher third minimum utility.
 - And so on...
- Maximizes the egalitarian welfare

The Leximin Mechanism

- DRF is the leximin mechanism
 - In the previous illustration, we didn't need tie-breaking because we assumed $d_{ir} > 0$ for every $i \in N, r \in R$.
 - In practice, not all the players need all the resources.
 - When $d_{ir} = 0$ is allowed, we need to continue allocating even after some agents are saturated.
 - Not all agents have equal dominant shares in the end.
- **Theorem [Parkes, Procaccia, S '12]:**
 - When $d_{ir} = 0$ is allowed, the leximin mechanism still retains all four properties (proportionality, envy-freeness, Pareto optimality, group strategyproofness).

A Note on Dynamic Settings

- We assumed that all agents are present from the start, and we want a one-shot allocation.
- Real-life environments are dynamic. Agents arrive and depart, and their demands change over time.
- **Theorem [Kash, Procaccia, S '14]:**
 - A dynamic version of the leximin mechanism satisfies proportionality, Pareto optimality, and strategyproofness along with a relaxed version of envy-freeness when agents arrive one-by-one.

A Note on Dynamic Settings

- Dynamic mechanism design
 - Designing fair, efficient, and game-theoretic mechanisms in dynamic environments is a relatively new research area, and we do not know much.
 - E.g., what if agents can depart, demands can change over time, or agents can submit and withdraw multiple jobs over time?
 - Lots of open questions!

Leximin (Dichotomous Matching)

Matching + Dichotomous Prefs

- Recall the stable matching setting of **matching n men to n women**.
 - We assumed ranked preferences, and showed that the Gale-Shapley algorithm produces a stable matching.
 - What if agent preferences weren't ranked?
- Suppose the men and women have **dichotomous preferences** over each other.
 - Each man finds a subset of women “acceptable” (utility 1), and the rest “unacceptable” (utility 0).
 - Same for women's preferences over men.

Matching + Dichotomous Prefs

- Dichotomous preferences induce a bipartite graph between men and women.
 - If a perfect matching exists, it's awesome.
 - What if there is no perfect matching?
 - Any deterministic matching unfairly gives 0 utility to some agents.
 - Solution: randomize!
- Under a random matching, utility to an agent = probability of being matched to an acceptable partner.

Matching + Dichotomous Prefs

- (Integral) Matching:
 - “Select” or “not select” each edge such that the number of selected edges incident on each vertex is at most 1.
- Fractional Matchings:
 - “Put a weight” on each edge such that the total weight of edges incident on each vertex is at most 1.
- Birkoff von-Neumann Theorem:
 - Every fractional matching can be “implemented” as a probability distribution over integral matchings.

Matching + Dichotomous Prefs

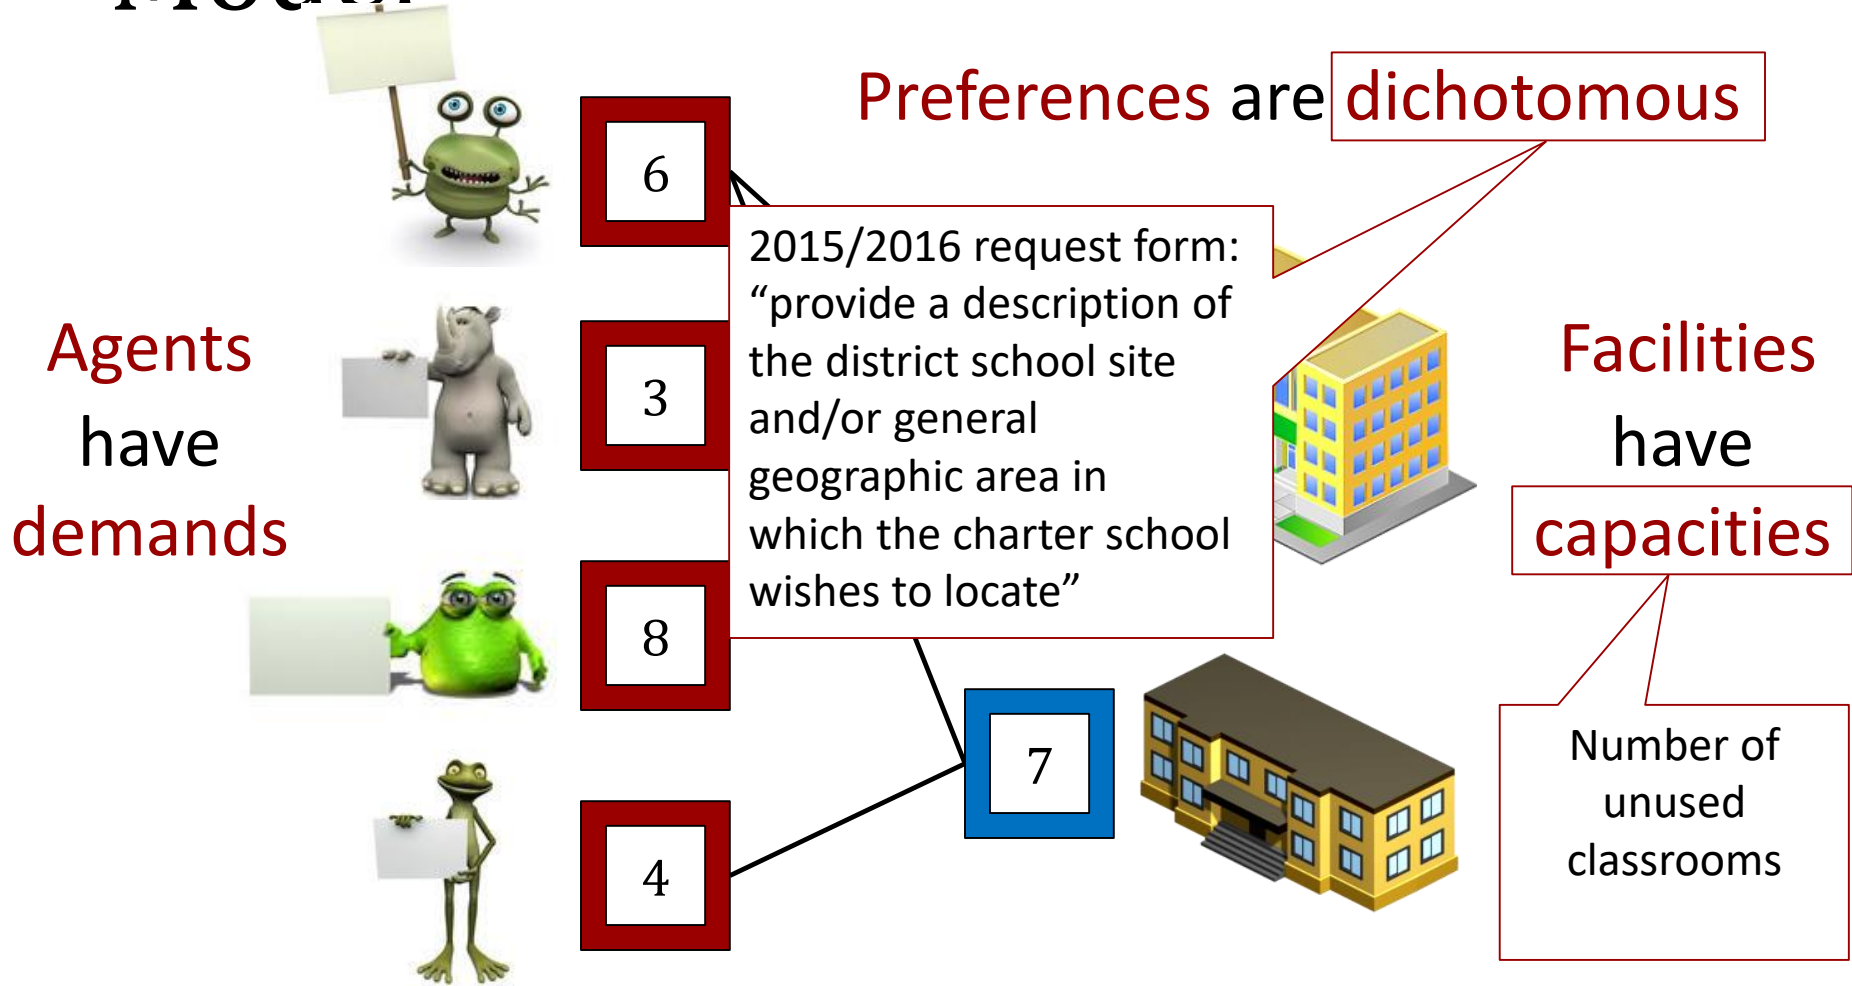
- **Randomized leximin mechanism:**
 - Compute the leximin fractional matching, and implement it as a distribution over integral matchings.
 - Both steps are doable in polynomial time!
- **Theorem [Bogomolnaia, Moulin '04]:**
 - The randomized leximin mechanism satisfies proportionality, envy-freeness, Pareto optimality, and group-strategyproofness (for both sides).
- In contrast: For ranked preferences, no algorithm can be strategyproof for both sides.

Matching with Capacities

- **Proposition 39** in California
 - “Unused resources in public schools should be *fairly* allocated to local charter schools that desire them.”
- Each charter school (**agent**) i wants d_i unused classrooms at one of the acceptable public schools (**facilities**) F_i .
 - If the demand is met, the charter school can relocate to the public school facility.
- Each facility j has c_j unused classrooms.
 - We assume facilities don't have preferences over agents.

Leximin (Classroom Allocation)

Model



Leximin Strikes Again

- Utility of agent i under a randomized allocation = probability of being allocated d_i classrooms at one of the facilities in F_i .
- **Theorem [Kurokawa, Procaccia, S '15]:**
 - The randomized leximin mechanism satisfies proportionality, envy-freeness, Pareto optimality, and group strategyproofness.
- Computing this allocation is NP-hard.
 - Unlike DRF and matching under dichotomous preferences.

Leximin Strikes Again

- The result holds in a generic domain which satisfies:
 - **Convexity:** If two utility vectors are feasible, then so should be their convex combinations.
 - Holds if fractional or randomized allocations are allowed.
 - **Equality:** The maximum utility of each agent should be the same.
 - Normalize utilities.
 - **Shifting Allocations:** Swapping allocations of two agents should be allowed.
 - **Maximal Utilization:** No agent should have a higher utility for agent i 's allocation than agent i has.
 - This should hold after the normalization. This is the most restrictive assumption.
- Captures DRF, matching with dichotomous preferences, classroom allocation, and many other settings from the literature.

Rent Division

Rent Division

- An apartment with n roommates & n rooms
- Roommates have preferences over the rooms
- Total rent is R
- **Goal:** Find an allocation of rooms to roommates & a division of the total rent that is envy-free.

Model

- Allocation A

- A_i = room given to i

- Rent division p

- p_r = rent charged for room r , $\sum_r p_r = R$

- Utilities


- $u_i(r, x)$ = utility to agent i for getting room r at rent x

- Envy-freeness:

$$\forall i, j : u_i(A_i, p_{A_i}) \geq u_i(A_j, p_{A_j})$$


Sperner's Lemma

- Triangle T partitioned into **elementary** triangles
- **Sperner Labeling:**
 - Label vertices $\{1,2,3\}$
 - Main vertices are different
 - Vertices between main vertices i and j are each labeled i or j
- **Lemma:**
 - Any Sperner labeling contains at least one “fully labeled” (1-2-3) elementary triangle.




Sperner's Lemma

- **Doors:** 1-2 edges
- **Rooms:** elementary triangles
- **Claim:** #doors on the boundary of T is odd
- **Claim:** A fully labeled (123) room has 1 door. Every other room has 0 or 2 doors.




Sperner's Lemma

- Start at a door on boundary, and walk through it
- Either found a fully labeled room, or it has another door
- No room visited twice
- Eventually, find a fully labeled room or back out through another door on boundary
- But #doors on boundary is odd. ■




Fair Rent Division

- Three housemates A, B, C
- **Goal:** Divide total rent between three rooms so that at those rents, each person wants a different room.
- Without loss of generality, say the total rent is 1.
 - Represent possible partitions of rent as a triangle.



Fair Rent Division

- “Triangulate” and assign “ownership” of each vertex to A, B, or C so that each elementary triangle is an ABC triangle





Fair Rent Division

- Ask the owner of each vertex v :
 - Which room do you prefer if the rent division is given by the coordinates of v ?
- Gives us a 1-2-3 labeling of the triangulation.
- Assumption: Each roommate prefers any free room over any paid room.
 - “Miserly roommates” assumption



Fair Rent Division

- This dictates the choice of rooms on the edges of T



Fair Rent Division

- Sperner's Lemma: There must be a 1-2-3 triangle.



Fair Rent Division

- The three roommates prefer different rooms...
 - But at slightly different rent divisions.
 - Approximately envy-free.
- By making the triangulations finer, we can increase accuracy.
 - In the limit, we obtain an envy-free allocation.
- This technique generalizes to more roommates [Su 1999].

Quasi-Linear Utilities

- A special case of the general model, but different than assuming “miserly roommates”
 - Value of roommate i for room $r = v_{i,r}$
 - Rent for room $r = p_r$
 - Utility to agent i for getting room $r = v_{i,r} - p_r$
- We need to find an assignment A of rooms to roommates and a price vector p such that
 - Total rent: $R = \sum_r p_r$
 - Envy-freeness: $v_{i,A_i} - p_{A_i} \geq v_{i,A_j} - p_{A_j}$

Quasi-Linear Utilities

- **Theorem:** An envy-free (A, p) always exists!
 - We'll skip this proof.
- **Theorem:** If (A, p) is envy-free, $\sum_i v_{i,A_i}$ is maximized.
 - Implied by “1st fundamental theorem of welfare economics”
 - As a consequence, (A, p) is Pareto optimal.
 - Easy proof!
- **Theorem:** If (A, p) is envy-free and A' maximizes $\sum_i v_{i,A'_i}$ then (A', p) is envy-free.
 - Further, $v_{i,A_i} - p_{A_i} = v_{i,A'_i} - p_{A'_i}$ for every agent i
 - Implied by “2nd fundamental theorem of welfare economics”
 - Easy proof!

PROVABLY FAIR SOLUTIONS.

Spliddit offers quick, free solutions to everyday fair division problems, using methods that provide indisputable fairness guarantees and build on decades of research in economics, mathematics, and computer science.



Share Rent



Split Fare



Assign Credit



Divide Goods



Distribute Tasks



Suggest an App

Which Model Is Better?

- Advantage of quasi-linear utilities:
 - One-shot preference elicitation
 - Players directly report their values for the different rooms
 - Easy to explain the fairness guarantee

Fairness Properties

Why is my assignment envy free? You were assigned the room called 'Smaller Bedroom' for \$314.33. Since you valued the room at \$427.00, you gained \$112.67. You valued the room called 'Master Bedroom' at \$247.00. Since this room costs \$331.33, you would have lost \$84.33. You valued the room called 'Attic' at \$326.00. Since this room costs \$354.33, you would have lost \$28.33.

Spliddit

Which Model Is Better?

- Advantage of miserly roommates model:
 - **Allows arbitrary preferences** subject to a simple assumption
 - **Easy queries**: “Which room do you prefer at these prices?”

What's your total rent? \$ How many of you are there?

If the rooms have the following prices, which room would you choose?

Choices will not necessarily be in order and the same roommate may be asked to choose multiple times in a row. Each roommate keeps choosing until a fair division is found.

Roommate A	No latest choice		
Roommate B	<input type="checkbox"/> \$500 Room 1	<input type="checkbox"/> \$500 Room 2	

The New York Times