CSC304 Lecture 20

Fair Division 3:

Leximin Allocation
(computational resources, matching with dichotomous prefs, classroom allocation)

Utilitarian Allocation
(rent division)

Computational Resources

- Setting: We have a cluster with a number of different resources (CPU, RAM, network bandwidth, etc.)
- A set of players collectively own the cluster.
- Assumption: 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)
  - That is, “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
- Allocation: $A_i = (A_{i1}, \dots, A_{im})$ where A_{ir} is the fraction of available resource r allocated to i
  - Thus, the utility to player i is $u_i(A_i) = \min_{r \in R} A_{ir}/d_{ir}$.
- We'll assume a non-wasteful allocation:
  - Allocates resources proportionally to the demand.

Dominant Resource Fairness

- Dominant resource of i = r such that $d_{ir} = 1$

- Dominant share of i = A_{ir} for dominant resource r

- Dominant Resource Fairness (DRF) Mechanism
  - Allocate maximal resources while maintaining equal dominant shares.

DRF animated

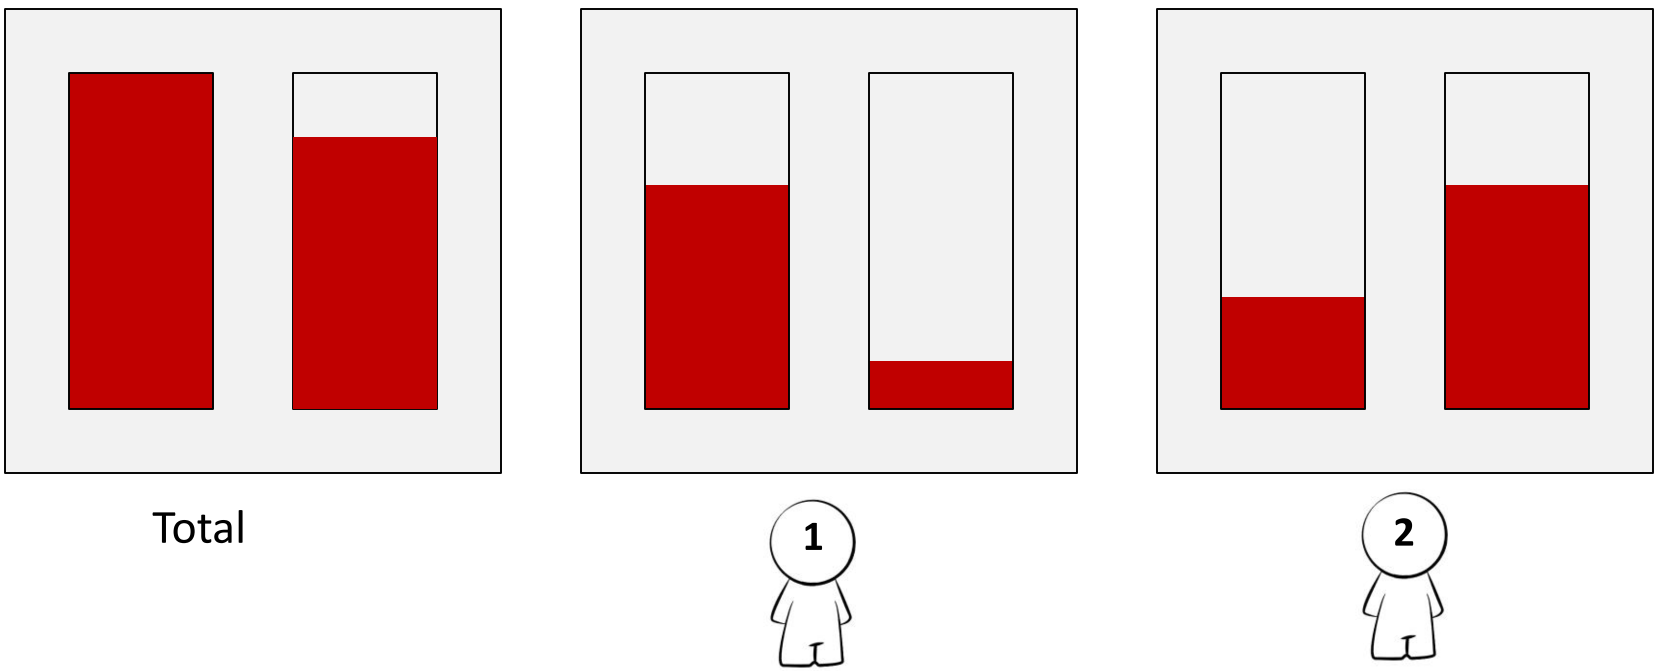

Properties of DRF

- Proportionality: $u_i(A_i) \geq 1/n$ for every player i
  - Why?
- Envy-free: $u_i(A_i) \geq u_i(A_j)$ for all players i, j
  - Why?
  - Note that we no longer have additive values across resources, so EF does not imply Proportionality (WHY?)
- Pareto optimality (Why?)
- Group strategyproofness:
  - If a group of players manipulate, it can't be that none of them lose, and some of them strictly gain
  - OK, this one is complicated.

The Leximin Mechanism

- Generalizes the DRF Mechanism

- Mechanism:
  - Choose an allocation A that maximizes the minimum of all utilities $\{u_i(A_i)\}_{i \in N}$
    - Sum = utilitarian welfare, product = Nash welfare, minimum = egalitarian welfare
  - If there are ties…
    - Break in favor of allocations that has a higher second minimum
    - Then break in favor of a higher third minimum
    - And so on…

The Leximin Mechanism

- DRF is the leximin mechanism applied to allocation of computational resources
  - ➢ It does not need to use 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.

- 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).

Dynamic Environments

- 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 variant of the leximin mechanism satisfies proportionality, Pareto optimality, and strategyproofness along with a relaxed version of envy-freeness when agents arrive over time.

Dynamic Environments

- Fair and game-theoretic allocation of resources in dynamic environments is a relatively new research area, and we do not know much.

- E.g., we do not have good algorithms that can handle departing agents, demands changing over time, or agents submitting/withdrawing multiple jobs over time.
  - Lots of open questions!

Matching + Dichotomous Prefs

- Let's revisit the problem of matching n men to n women.
- Recall that the Gale-Shapley algorithm used ranked preferences from both sides to find a stable matching.
- Consider a different case in which every man (resp. woman) has a subset of women (resp. men) that are acceptable (utility 1) and the rest are unacceptable.

Matching + Dichotomous Prefs

- Formally, for each man m, there is a subset of "acceptable" women P_m such that the man has utility 1 for being matched to any woman in P_m, and utility 0 otherwise.

- If there exists a perfect matching, that's awesome.
  - But what if there isn't?
- Any solution that wants to achieve fairness (proportionality or envy-freeness) must randomize!
  - Utility to agent = probability of being matched to an acceptable partner

Matching + Dichotomous Prefs

- Randomized mechanisms:
  - We can think of all men and women as “divisible” (oops!)
  - When we say that a woman w is “allocated” 0.3 fraction of a man m, it means the probability that w will be matched to m is 0.3.
  - You can just compute the fractional allocation that maximizes the minimum utility (then the second minimum etc).
    - Birkoff von-Neumann Theorem: Every fractional assignment can be written as a probability distribution over integral assignments.

Matching + Dichotomous Prefs

- Theorem [Bogomolnaia, Moulin '04]:
  - The randomized leximin mechanism satisfies proportionality, envy-freeness, Pareto optimality, and group-strategyproofness (for both sides simultaneously!).

- Compare this to the case of ranked preferences in which an algorithm can only be strategyproof for one side of the market, but not both.

Matching with Capacities

- Proposition 39 in California mandates that unused classrooms in public schools be *fairly* assigned to charter schools that want it.
  - If the charter school receives a sufficient number of classrooms to fit all its students, it can physically relocate to the public school facility (e.g., and save on rent).
- Each charter school (agent) i has a set of acceptable public schools (facilities) F_i, but also has a demand d_i for the number of classrooms.
- Each facility j has a capacity c_j (#classrooms available)

Model

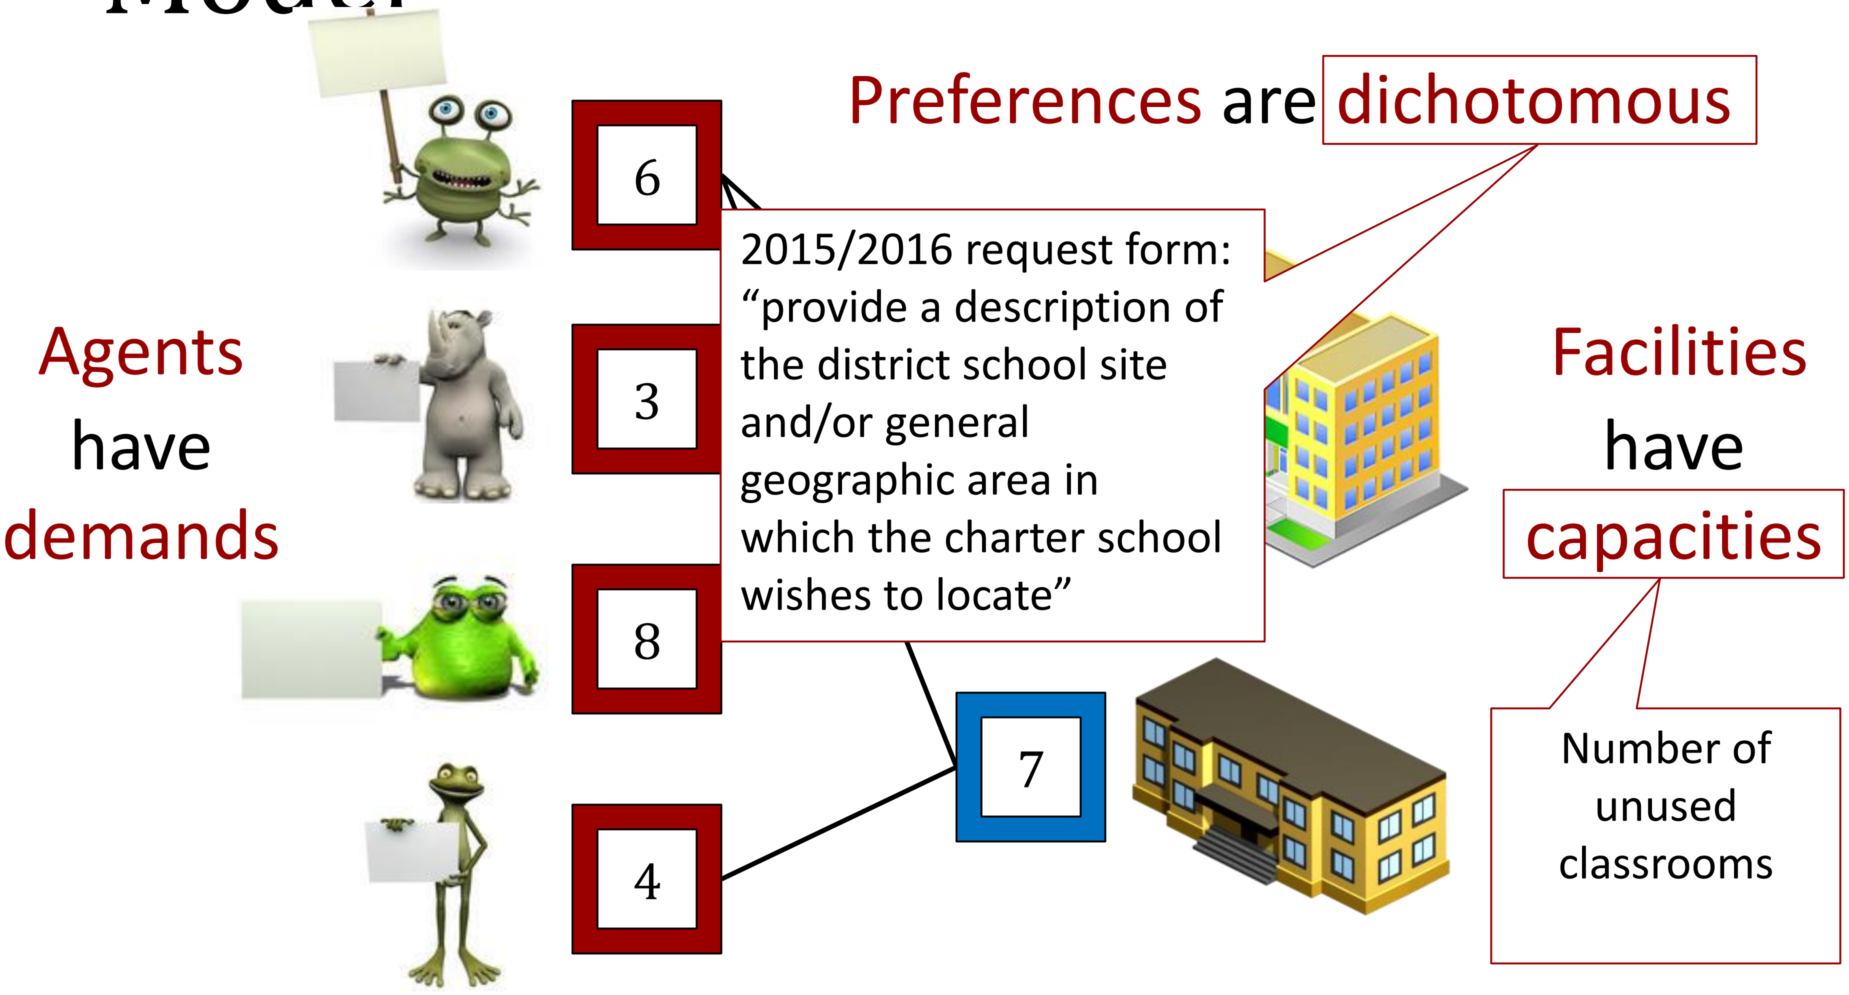
Preferences are dichotomous
Agents have demands
Facilities have capacities
6
3
8
4
7
2015/2016 request form: “provide a description of the district school site and/or general geographic area in which the charter school wishes to locate”
Number of unused classrooms

Leximin Strikes Again

- Theorem [Kurokawa, Procaccia, S ‘15]:
  - The randomized leximin mechanism satisfies proportionality, envy-freeness, Pareto optimality, and group strategyproofness for classroom allocation.

- In fact, the result holds under a wider domain satisfying a “maximal utilization” property.
  - Generalizes DRF, matching with dichotomous preferences, and 8-10 other settings
- For allocating computational resources or matching under dichotomous preferences, the leximin mechanism can be computed in polynomial time.
  - In contrast, it is NP-hard to compute for classroom allocation.

Rent Division

- n roommates rent an apartment with n rooms.
- Roommate i has value $v_{i,r}$ for room r.
- The total rent is R.
  - Assume that $\sum_r v_{i,r} \geq R$ for every roommate i.

- We need to find an allocation A of rooms to roommates and a price vector p such that
  - Total rent: $R = \sum_r p_r$
  - Envy-freeness: $v_{i,A_i} - p_{A_i} \geq v_{i,A_j} - p_{A_j}$

Rent Division: Fascinating Facts

- Existence: An envy-free allocation (A, p) always exists! (hard proof ☹)
- 1st Fundamental Theorem of Welfare Economics:
  - If (A, p) is an envy-free allocation, then A must maximize the sum of values (utilitarian welfare)!
  - Easy proof!
- 2nd Fundamental Theorem of Welfare Economics:
  - If (A, p) is an envy-free allocation, and A' is any allocation maximizing utilitarian welfare, 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.
  - Easy proof!

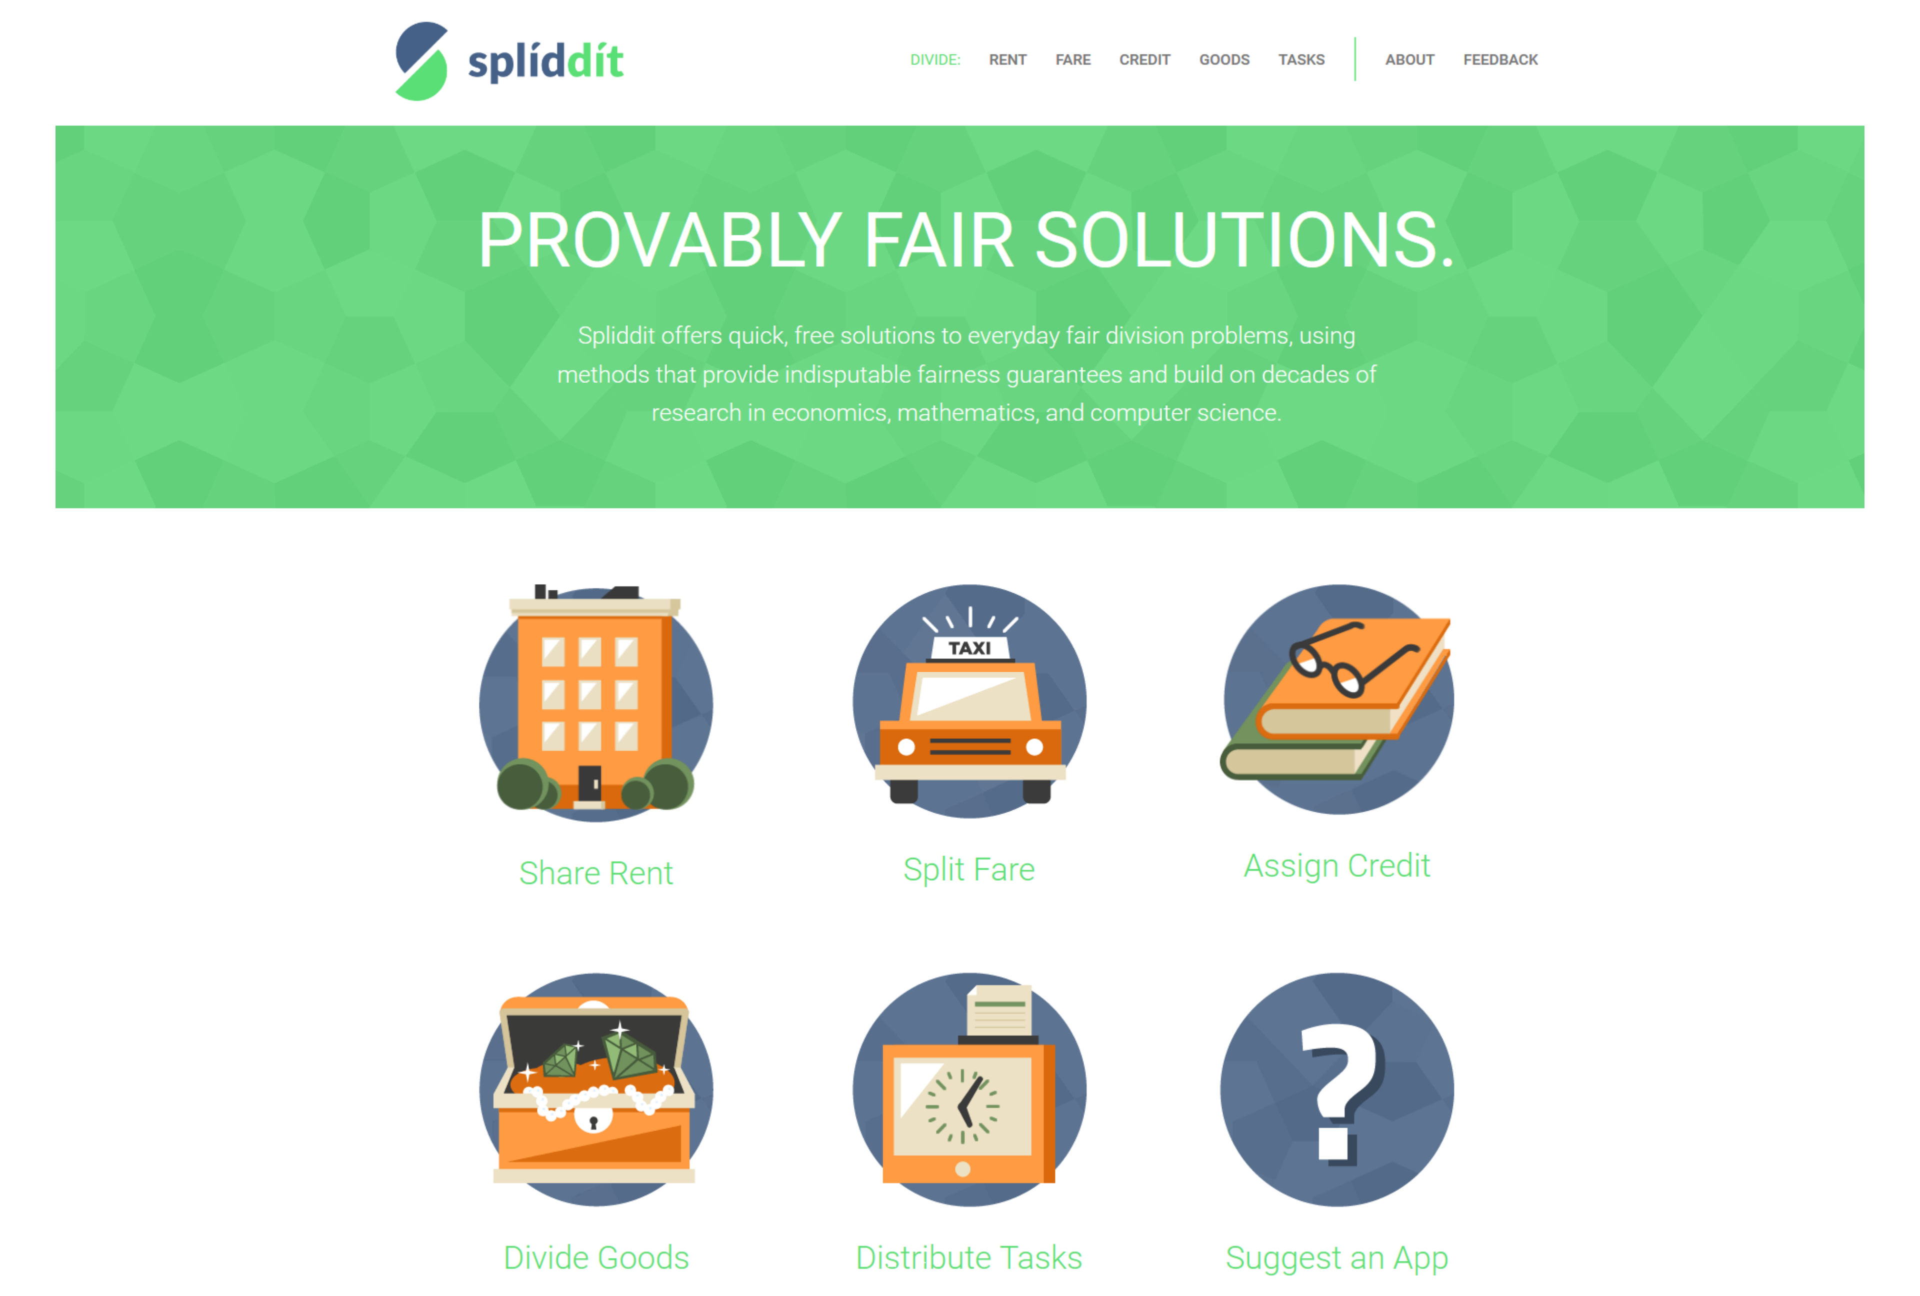
spliddit
DIVIDE: RENT FARE CREDIT GOODS TASKS
ABOUT FEEDBACK
PROVABLY FAIR SOLUTIONS.
Spliddit offers quick, free solutions to everyday fair division problems, using methods that provide indisputable fairness guarantees and build on decades of research in economics, mathematics, and computer science.
TAXI
Share Rent
Split Fare
Assign Credit
Divide Goods
Distribute Tasks
Suggest an App