CSC373 Fall 2021

Pre-module Video

Embedded Ethics Module

Overview of Algorithmic Fairness

Algorithms Making Decisions

Loans

Bails

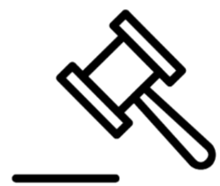

Self-Driving Cars

Ads

Hiring

Organ Exchange

The Setting: Resource Allocation

- m *resources*, to be distributed amongst n *agents*
- For concreteness, may think of the inheritance problem from the philosophy video
- Each agent may value a given resource differently
- Question: What is a “fair” and “effective” way to distribute the resources?

Formalizing Valuations: Utility Functions

- How to model how much value an agent assigns to a particular resource?
- Utility Functions! Agent i is assumed to have an associated function $v_i(\cdot)$
- $v_i(r)$ is a numeric value for each resource r, typically a non-negative real number
- Often, we assume that they are *additive:* for disjoint subsets X,Y of resources, the valuation $v_i(X \cup Y)$ of agent i for X *and* Y is the sum $v_i(X) + v_i(Y)$

Formalizing Fairness

Allocations

- Given the utility functions, makes sense to talk about fairness *axioms* we want our final allocation to satisfy
- Let $A = (A_1, \dots, A_n)$ be an allocation i.e., a *partition* of the set of m resources into n subsets
- Here, A_i indicates the set of resources allocated to agent i

Fairness Axioms

- What kind of fairness might we want from an allocation?
- Proportionality:

For all agents i, $v_i(A_i) \geq \frac{1}{n} \cdot (\sum_{j=1}^{m} v_i(r_j))$

(Each agent receives at least a $1/n$ – fraction of her total valuation of all the resources)

- Envy-Freeness:

For all agents i, j, $v_i(A_i) \geq v_i(A_j)$

(No agent envies another agent)

- Equitability:

For all agents i, j, $v_i(A_i) = v_j(A_j)$

(All agents get equal utility)

Example

- Resources are intervals

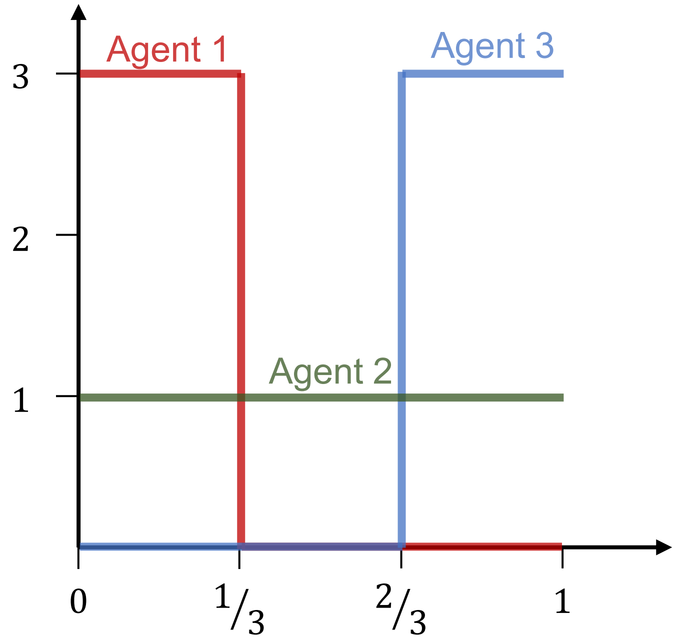

- Agent 1 wants $[0, {}^1/_3]$ uniformly and does not want anything else

- Agent 2 wants the entire interval uniformly

- Agent 3 wants $[{}^2/_3, 1]$ uniformly and does not want anything else

Example

- Resources are intervals

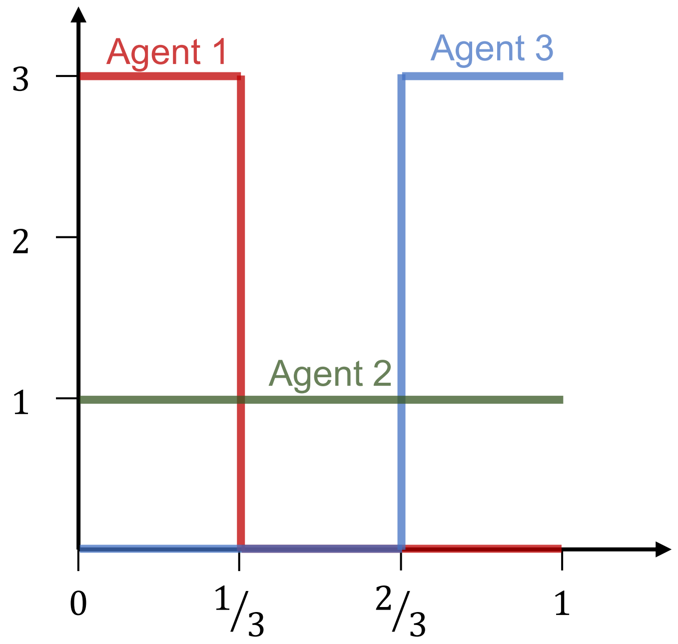

- Consider the following allocation

- $A_1 = [0, {}^1/_9] \Rightarrow v_1(A_1) = {}^1/_3$
- $A_2 = [{}^1/_9, {}^8/_9] \Rightarrow v_2(A_2) = {}^7/_9$
- $A_3 = [{}^8/_9, 1] \Rightarrow v_3(A_3) = {}^1/_3$

- The allocation is proportional, but not envy-free or equitable

Example

- Resources are intervals

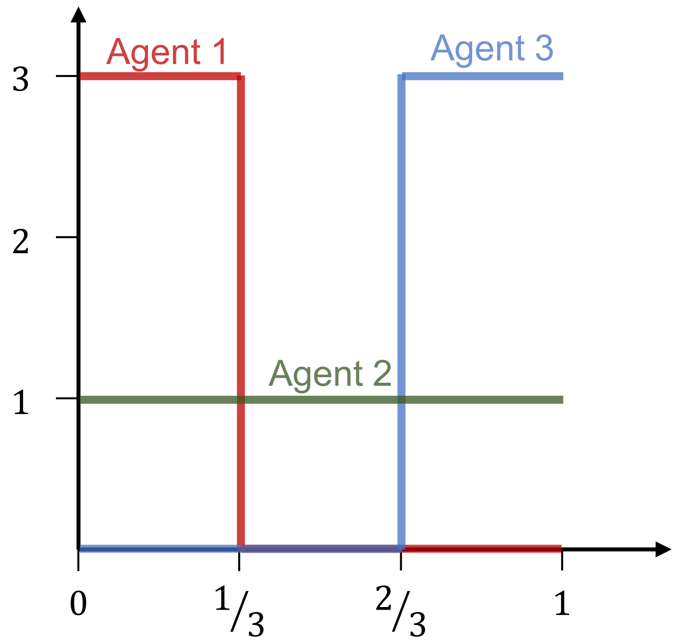

- Consider the following allocation

- $A_1 = [0, {}^1/_6] \Rightarrow v_1(A_1) = {}^1/_2$
- $A_2 = [{}^1/_6, {}^5/_6] \Rightarrow v_2(A_2) = {}^2/_3$
- $A_3 = [{}^5/_6, 1] \Rightarrow v_3(A_3) = {}^1/_2$

- The allocation is proportional and envy-free, but not equitable

Example

- Resources are intervals

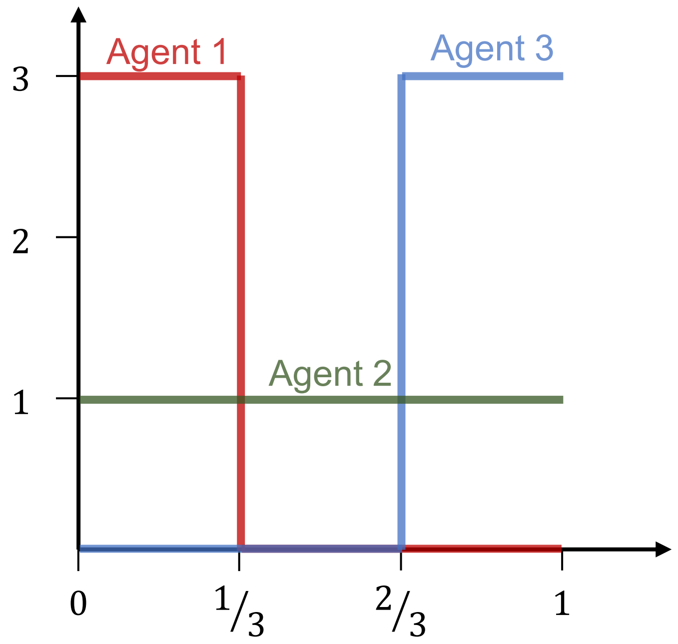

- Consider the following allocation
- $A_1 = [0, {}^1/_5] \Rightarrow v_1(A_1) = {}^3/_5$
- $A_2 = [{}^1/_5, {}^4/_5] \Rightarrow v_2(A_2) = {}^3/_5$
- $A_3 = [{}^4/_5, 1] \Rightarrow v_3(A_3) = {}^3/_5$
- The allocation is proportional, envy-free, and equitable

Relations Between Fairness Desiderata

- Envy-freeness implies proportionality
  - Summing $v_i(A_i) \geq v_i(A_j)$ over all j gives proportionality
- For 2 agents, proportionality also implies envy-freeness
  - Hence, they are equivalent.
- Equitability is incomparable to proportionality and envy-freeness
  - E.g. if each agent has value 0 for her own allocation and 1 for the other agent's allocation, it is equitable but not proportional or envy-free.

Welfare Functions

Welfare Functions

- Functions that you eventually want to optimize, subject to select fairness axioms
- Ways to combine the utility functions of different stakeholders into a single objective to be optimized
- Important to keep their complexity in mind

Popular Choices

- Social welfare:
  - Maximize the sum of utilities of an allocation, $\sum_{i=1}^{n} v_i(A_i)$
  - Optimizes for the "efficiency" building block (from phil video)
- Egalitarian welfare:
  - Maximize the minimum among the various utilities of an allocation, $\min_{i=1}^{n} v_i(A_i)$
  - Formalizes the "egalitarianism" building block
- Nash welfare:
  - Maximize the product of utilities of an allocation, $\prod_{i=1}^{n} v_i(A_i)$
  - Offers a potential trade-off between the two above

Dangers of Interpersonal Comparisons

- Important to keep in mind the risks of making interpersonal comparisons across diff. utility functions
- There's no reason for different agents to stick to the same "scale"
- For example, $v_1(apple) = 1, v_1(orange) = 2$ and $v_2(apple) = 2000, v_2(orange) = 1000$
- Note however that Nash Welfare is scale-free!
- (Doesn't mean that it's the best to use in every situation)