

CPS 296.1

Mechanism design

Vincent Conitzer
conitzer@cs.duke.edu

Mechanism design: setting

- The **center** has a set of outcomes O that she can choose from
 - Allocations of tasks/resources, joint plans, ...
- Each agent i draws a **type** θ_i from Θ_i
 - usually, but not necessarily, according to some probability distribution
- Each agent has a (commonly known) **valuation function** $v_i: \Theta_i \times O \rightarrow \mathbb{R}$
 - Note: depends on θ_i , which is **not** commonly known
- The center has some **objective function** $g: \Theta \times O \rightarrow \mathbb{R}$
 - $\Theta = \Theta_1 \times \dots \times \Theta_n$
 - E.g., efficiency ($\sum_i v_i(\theta_i, o)$)
 - May also depend on payments (more on those later)
 - The center does **not** know the types

What should the center do?

- She would like to know the agents' types to make the best decision
- Why not just ask them for their types?
- Problem: agents might **lie**
- E.g., an agent that slightly prefers outcome 1 may say that outcome 1 will give him a value of 1,000,000 and everything else will give him a value of 0, to force the decision in his favor
- But maybe, if the center is clever about choosing outcomes and/or requires the agents to make some **payments** depending on the types they report, the incentive to lie disappears...