Truthful Verifiable MechanismS

Simina Brânzei
(Aarhus University)

Joint with Ariel Procaccia (CMU)

Social Choice

- Set A of alternatives
- Set N = {1, …, n} of strategic agents
- Each agent has an order over the alternatives in A
- Social Choice Function (i.e. mechanism):
  - Map $F : L^n \rightarrow A$ from the preferences of the agents to an alternative.

Social Choice

- **Desirable properties for social choice functions:**

- ***Onto:*** Every outcome can be selected.
- ***Non-dictatorial:*** For each agent i there is a profile $>_1, \ldots, >_n$ such that $F(>_1, \ldots, >_n) \neq$ agent i's top choice
- ***Truthful (strategyproof):*** An agent can never benefit from lying regardless of the others' strategies

Social Choice

Gibbard-Satterthwaite Theorem:

- ***Every strategyproof and onto social choice function over more than two alternatives is dictatorial.*** 

Better news with:

- Money
- Randomization
- Restricted preference domain (e.g. single peaked)

Social Choice

Single Peaked Preferences

- Set of alternatives is one-dimensional
- Each agent i has a single peak x_i^* (bliss point).

That is $\forall a, b \in A$:

- If $b < a < x_i^*$ then $a >_i b$
- If $x_i^* > a > b$ then $a >_i b$

Facility Location

- Set of agents $N = \{1, \ldots, n\}$
- Set of outcomes: real line
- A mechanism takes a set of preferences and outputs a facility location.
- When is a mechanism ***good***?
  - ***Truthful:*** Each agent truthfully tells its ideal location
  - ***Approximates well*** an objective function (social welfare)

Facility Location

- The cost of an agent is the distance between its ideal location and the facility

- Objective functions:
  - *Social Cost* = sum of distances to the facility
  - *Maximum Cost* = maximum distance (of any agent)

This work

- Mechanism *M* has an approximation of α if for every instance x, $Cost(M(x))/OPT \leq \alpha$.
  - Mechanisms with good approximations ratios are ***good*** : proxy for happy agents on average.
- But why truthfulness?
  - If the mechanism is truthful, the agents have ***simple*** strategies
  - Bounded rational agents might not come up with/execute complicated strategies even if they wanted

This work

- ***Computational Model:***
  - The mechanism is a Turing machine: takes the inputs from agents and outputs an outcome (facility)
  - The agents are programs, too

This work

- ***Our research question is:***
  - ***Can the agents (or the center) convince themselves before play that the mechanism is truthful?***
  - ***Can the "proof" be "short"? Say less than 3 pages?***

  - If not, agents might try to tweak the system anyway......

    (e.g. under partial information they might come up with strange manipulations that seem improving...)

This work

- Undecidability is a barrier:
  - The agents cannot generate and verify a proof of truthfulness for every mechanism.
- ***Restricted Computational Model:***
  - The mechanism is a decision tree
  - Agents and proofs are programs
- Case study of this model in facility location.

Facility Location

Average Mechanism:

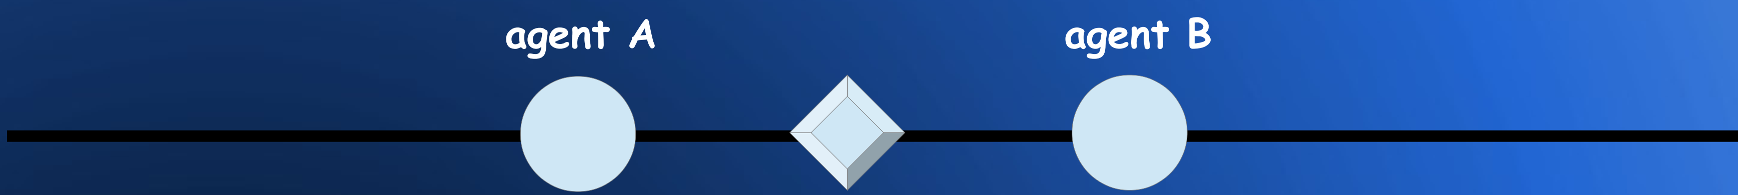

Manipulation....

Facility Location

Rightmost Mechanism

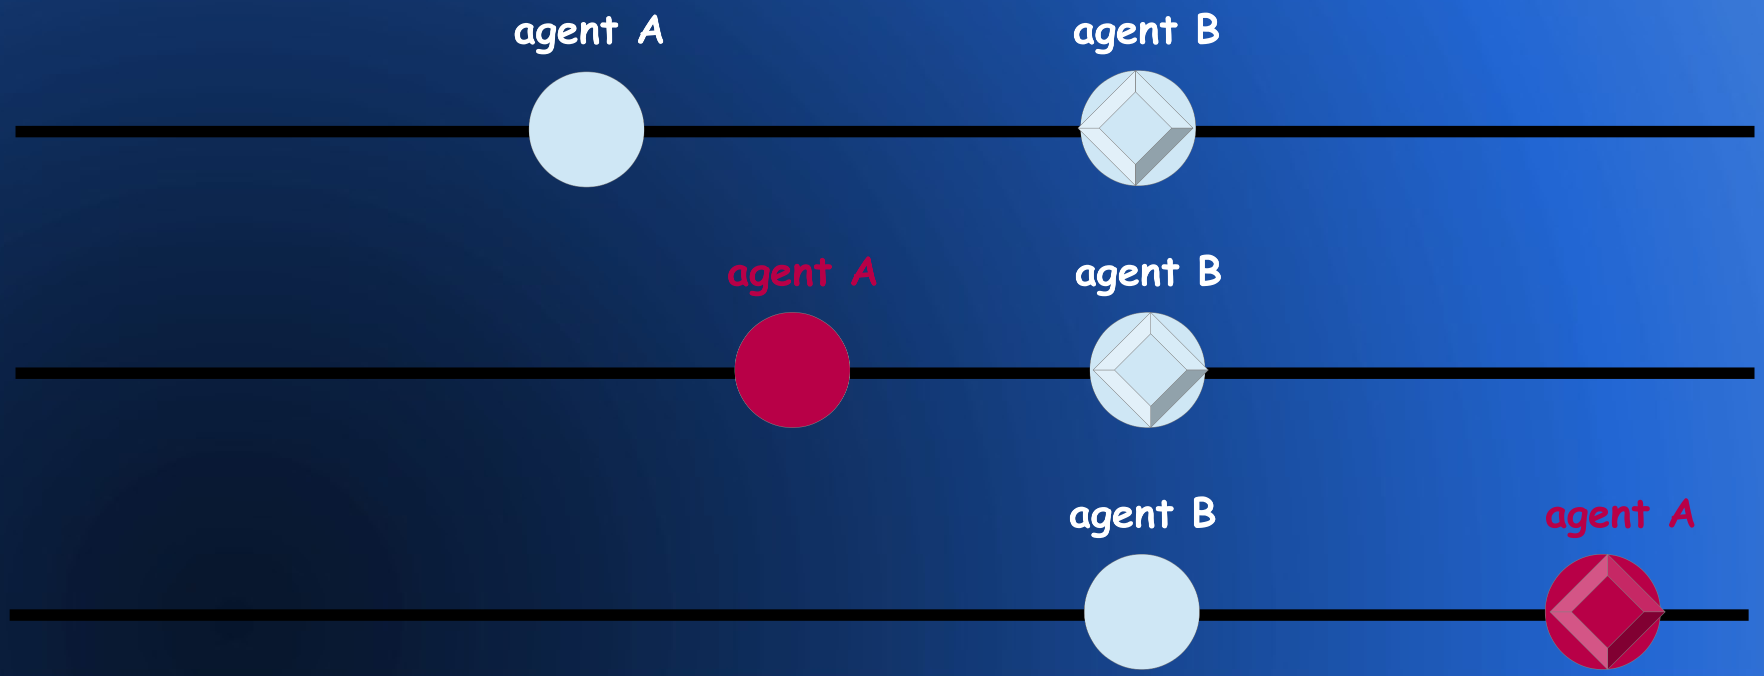

Facility Location

Median Mechanism:

Deterministic Mechanisms

Decision Tree Mechanisms:

- Internal nodes contain binary comparisons of the inputs: $(x_i \geq x_j)$, $(x_i > x_j)$, $(x_i \leq x_j)$, $(x_i < x_j)$, where $i, j \in N$.
- Each leaf contains a location, as convex combinations of the inputs: $\lambda_1 x_1 + \ldots \lambda_n x_n$. Each $\lambda_i \geq 0$ and $\sum \lambda_i = 1$ (the λs might be different for every leaf).

Deterministic Mechanisms

Median for 3 agents:

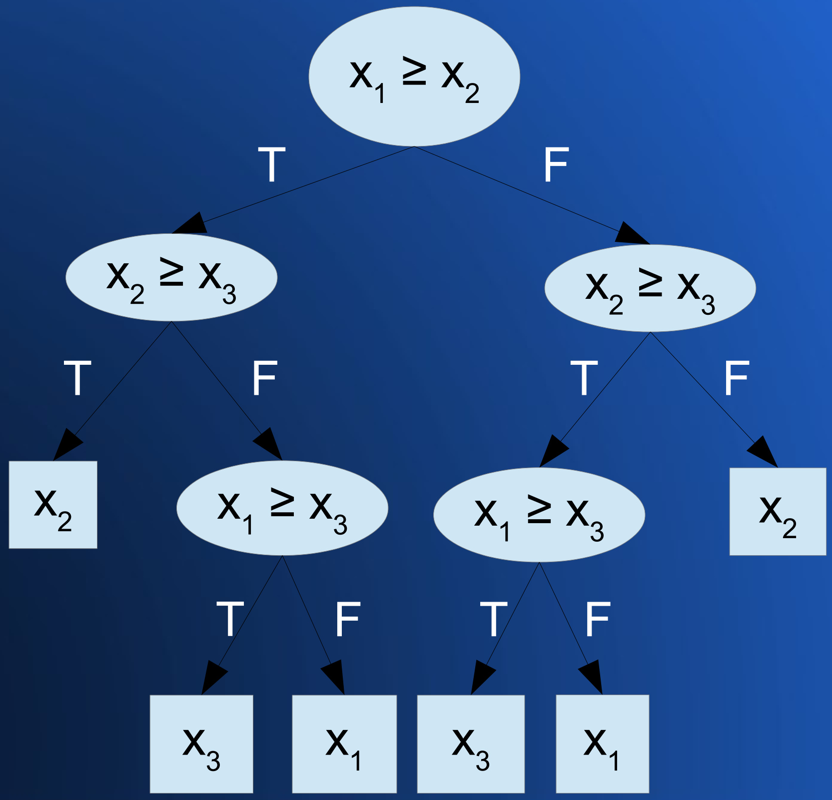

Deterministic Mechanisms

Dictatorship

x_1

Average Mechanism

$$\frac{x_1 + x_2 + \ldots + x_n}{n}$$

Deterministic Mechanisms

When is a tree mechanism M ***not*** truthful?

Example: *Agent 3* can influence the outcome by reporting $x_3' < x_1$

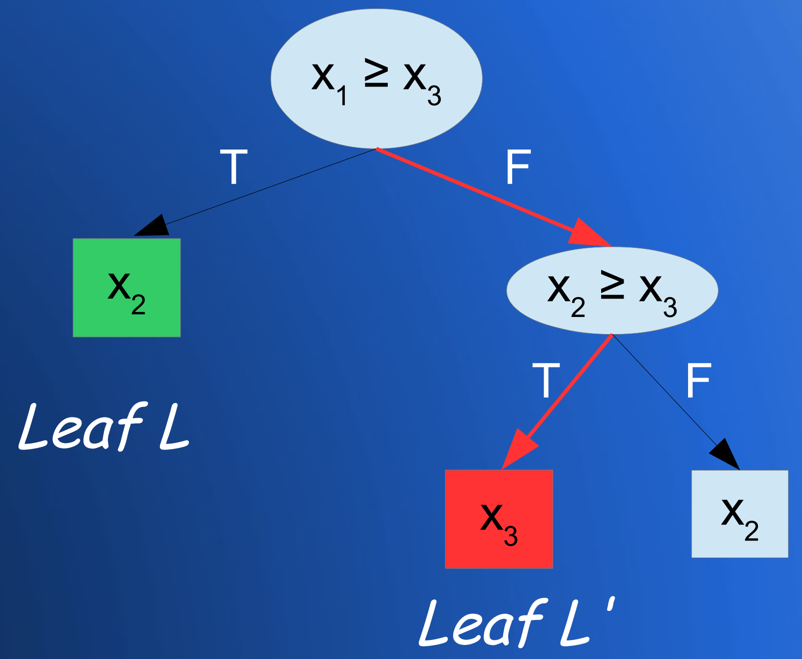

Verifier for Deterministic Mechanisms

Input: tree mechanism T

For every agent *i*:

For every two leaves *L*, *L'*:

Solve an $LP(x_1, \ldots, x_n, x_i')$, where x verifies the constraints to L, $x' = (x_i', x_{-i})$ to L', and $dist(x_i, L'(x')) < dist(x_i, L(x))$

If solution exists, return False

Return True

Deterministic Mechanisms

Theorem: There is a verifier that establishes truthfulness for every decision tree T and runs in $O(poly(n, |T|))$

In worst case, ***any*** verifier must check all the leaves, so T must have poly-size for the total runtime to be $O(poly(n))$

Social Cost: What can be done with poly-size decision trees?

Deterministic Mechanisms

The median computes the optimal social cost and is strategyproof

Some bad news: The median decision tree has height $> n$ and $< 6n$

What approximation ratios do trees of logarithmic height have for the social cost?

Deterministic Mechanisms

Theorem: Deterministic decision trees of polynomial size approximate the ***social cost*** within a factor of $\Theta(n/\log(n))$.

Output the median of the first k agents, where $k = \log(n)$

Theorem: Deterministic decision trees approximate the ***maximum cost*** within a factor of 2 and this is tight (just as bad as deterministic Turing machines).

Just pick any dictator.

Randomized Mechanisms

Randomized decision tree:

The root node is randomized (selects a subtree with some probability)

Each subtree with the root as parent is deterministic

Universal Truthfulness:

The agents have no incentives to lie even after knowing the outcomes of the random coin tosses of the mechanism

Randomized Mechanisms

Randomized mechanism example:

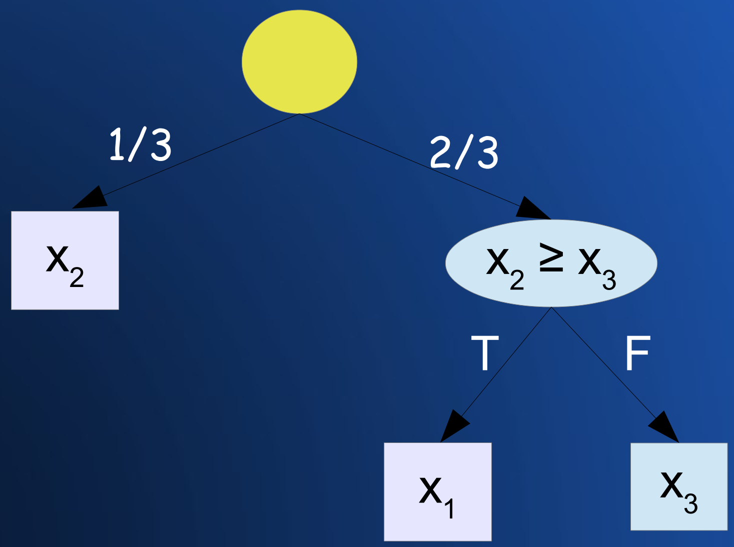

Randomized Mechanisms

Social Cost: There exists a randomized decision tree of polynomial size that is universally truthful and approximates the social cost within a factor of 2 – 2/n.

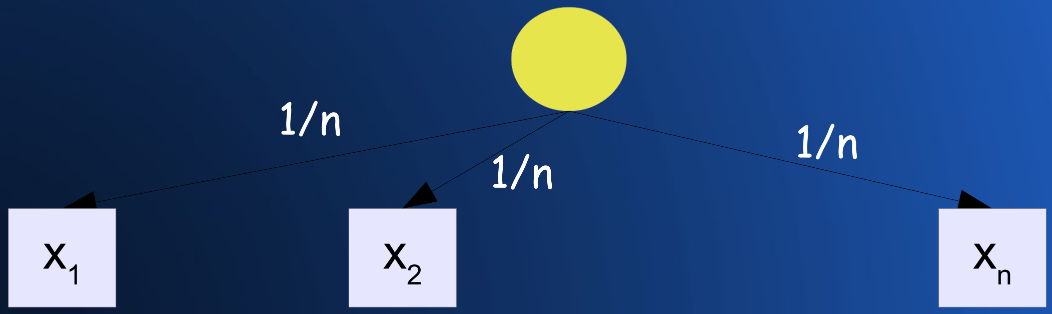

Randomized Mechanisms

Parameterized Mechanism:

Select k inputs from $X = \{x_1, \ldots, x_n\}$ at random with replacement

Output the median of the sampled set

The approximation ratio improves quite quickly, so for small k we get a mechanism of poly-size with better approximation ratio.

Randomized Mechanisms

For maximum cost, randomization doesn't add anything for truthfulness in expectation.

Theorem: there is no randomized decision tree mechanism that is universally truthful and has an approximation better than *2* for maximum cost (for any tree size).

Randomized Mechanisms

Truthfulness in Expectation:

The agents have no incentives to lie before seeing the coin tosses (but may regret their choice afterwards).

Verifier that checks truthfulness in expectation runs in

$O(poly(n, 2^{|T|}))$.

Randomized Mechanisms

Left-Right-Middle Mechanism (Procaccia & Tennenholtz, EC '09)

- With probability $\frac{1}{4}$ output the leftmost location
- With probability $\frac{1}{4}$ output the rightmost location
- With probability $\frac{1}{2}$ output the middle point

Approximation ratio of 3/2 for max cost; optimal over all truthful in expectation mechanisms.

Randomized Mechanisms

Good news: LRM can be implemented as a randomized decision tree (thus the best approximation ratio can be achieved)

Bad news: the decision tree has exponential size in *n*

Parameterized LRM:

- Select a random subset of size k from {1, …, n}
- Run LRM on the selected subset

The tree has poly size, but our verifier is still inefficient…

Discussion

Deterministic mechanisms:

– Verifier runs in polynomial time in *n* and the tree size

– Social Cost: approx ratio = 1 (exponential size); $\Theta(n/\log(n))$ (poly size)

– Maximum Cost: approx ratio = 2 (any cost)

Discussion

Universally Truthful Mechanisms:

– Verifier still runs in poly-time in n and the tree size

– Social Cost: approx ratio « 2 – 2/n

– Maximum Cost: approx ratio = 2 (any size)

Truthful in Expectation Mechanisms:

– Verifier runs in exponential time in n and the tree size

– Maximum Cost: approx ratio = 3/2 (exp size tree)