CMU 15-896

Social choice 3: Advanced manipulation

Teacher:
Ariel Procaccia

Recap

- A Complexity-theoretic barrier to manipulation
- Polynomial-time greedy alg can decide instances of R-Manipulation for R = scoring rules, Copeland, Maximin,... $\Rightarrow$ these rules are easy to manipulate in practice
- Some rules are NP-hard to manipulate: STV, ranked pairs,...

CRITICISMS

- What is the complexity of the Dictatorship-MANIPULATION problem?
- NP-hardness is worst-case, but perhaps a manipulator can usually succeed
- Approaches:
  - Algorithmic: for specific voting rules but works for every reasonable distribution
  - Quantitative G-S: for a specific distribution but works for every reasonable voting rule

Quantitative G-S

- We'll do this roughly, to capture intuitions rather than aiming for accuracy
- The distance between two voting rules is the fraction of inputs on which they differ

$$D(f,g) = \Pr[f(\vec{\precsim}) \neq g(\vec{\precsim})]$$

  where the Pr is over uniformly random preference profiles $\vec{\precsim}$
- For a set F, $D(f,F) = \min_{g \in F} D(f,g)$
- F_{dic} = set of dictatorships, $|F_{dic}| = n$

Quantitative G-S

- $(\vec{\prec}, \prec_i')$ is a manipulation pair for f if
$$f(\prec_i', \vec{\prec}_{-i}) \succ_i f(\vec{\prec})$$
- Theorem [Mossel and Racz 2012]: $m \geq 3$, f is onto, $D(f, F_{dic}) \geq \epsilon$. Then
$$\Pr\left[(\vec{\prec}, \prec_i') \text{ manip. pair for } f\right] \geq p\left(\epsilon, \frac{1}{n}, \frac{1}{m}\right)$$
for a polynomial p, where $\vec{\prec}$ and $\prec_i'$ are chosen uniformly at random

Randomized voting rules

- Randomized voting rule: outputs a distribution over alternatives
- To think about successful manipulations we need utilities (assume strict preferences)
- $\prec_i$ is consistent with u_i if
$$x \succ_i y \Leftrightarrow u_i(x) > u_i(y)$$
- Strategyproofness: $\forall i \in N, \forall u_i, \forall \vec{\precsim}_{-i}, \forall \prec_i',$
$$\mathbb{E}\left[u_i\left(f(\vec{\precsim})\right)\right] \geq \mathbb{E}\left[u_i\left(f(\prec_i', \vec{\precsim}_{-i})\right)\right]$$
where $\prec_i$ is consistent with u_i

Randomized voting rules

- A (deterministic) voting rule is
  - unilateral if it only depends on one voter
  - duple if its range is of size at most 2
- A randomized rule is a probability mixture over rules $f_1, \ldots, f_k$ if there exist $\alpha_1, \ldots, \alpha_k$ such that for all $\vec{\precsim}$, $\Pr[f(\vec{\precsim}) = f_j(\vec{\precsim})] = \alpha_j$

Randomized voting rules

- Theorem [Gibbard 1977]: A randomized voting rule is strategyproof only if it is a probability mixture over unilaterals and duples

Mixture over unilaterals and duples that is not SP?

Randomization+approximation

- Idea: can strategyproof randomized rules approximate popular rules?
- Fix a rule with a clear notion of score (e.g., Borda) denoted $\mathrm{sc}(\vec{\prec}, x)$
- Randomized rule f is a c-approximation if for every preference profile $\vec{\prec}$,

$$\frac{\mathbb{E}\left[\mathrm{sc}\left(\vec{\prec}, f(\vec{\prec})\right)\right]}{\max_{x \in A} \mathrm{sc}(\vec{\prec}, x)} \geq c$$

Approximating Borda

Poll 1: What is the approximation ratio to Borda from randomly choosing an alternative?

1. $\Theta(1/n)$
2. $\Theta(1/m)$
3. $\Theta(1/\sqrt{m})$
4. $\Theta(1)$

- Theorem [P 2010]: No strategyproof randomized voting rule can approximate Borda to a factor of $\frac{1}{2} + \omega\left(\frac{1}{\sqrt{m}}\right)$

Interlude: Zero-sum games

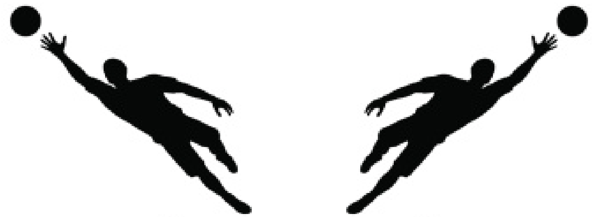

-1	1
1	-1

Interlude: Minimax strategies

- Minimax (randomized) strategy minimizes worst-case expected loss (or maximizes the expected gain)
- In the penalty shot game, minimax strategy for both players is playing $\left(\frac{1}{2},\frac{1}{2}\right)$
- In the game below, if shooter uses $(p, 1-p)$:
  - Jump left: $-\frac{p}{2}+1-p=1-\frac{3}{2}p$
  - Jump right: $p-1+p=2p-1$
  - Maximize $\min\{1-\frac{3}{2}p, 2p-1\}$ over p

$-\frac{1}{2}$	1
1	-1

Interlude: The Minimax Theorem

- Theorem [von Neumann, 1928]: Every 2-player zero-sum game has a unique value v such that:
  - Player 1 can guarantee value at least v
  - Player 2 can guarantee loss at most v

Yao's minimax principle

	$\succsim^1$	…	…	…	…	$\succsim^t$
U_1	$\frac{1}{15}$	…	…	…	…	$\frac{2}{21}$
…	…	…	…	…	…	…
U_k	$\frac{7}{15}$	Approximation ratio				$\frac{5}{21}$
D_1	$\frac{4}{15}$	…	…	…	…	$\frac{8}{21}$
…	…	…	…	…	…	…
D_s	$\frac{13}{15}$	…	…	…	…	$\frac{17}{21}$

Yao's minimax principle

- Maximin Theorem $\Rightarrow$ The expected ratio of the best distribution over unilateral rules and duples against the worst preference profile is equal to the expected ratio of the worst distribution over profiles against the best unilateral rule or duple
- An upper bound on the approximation ratio of the best distribution over unilateral rules and duples is given by some distribution over profiles against the best unilateral rule or duple
- Gibbard's Theorem $\Rightarrow$ this is also an upper bound on the best randomized strategyproof rule

A bad distribution

- $m = n + 1$
- Choose $x^* \in A$ uniformly at random
- Each voter i chooses a random number $k_i \in \{1, \dots, \sqrt{m}\}$ and puts x^* in position k_i
- The other alternatives are ranked cyclically

1	2	3
c	b	d
b	a	b
a	d	c
d	c	a

$x^* = b$
$k_1 = 2$
$k_2 = 1$
$k_3 = 2$

A bad distribution

Poll 2: What is the best feasible lower bound on $\mathrm{sc}(\vec{\prec}, x^*)$?

1. $\sqrt{n}$
2. $\sqrt{m}$
3. $n(m-\sqrt{m})$
4. nm

A bad distribution

- For $x \in A \setminus \{x^*\}, \mathrm{sc}(\vec{\prec}, x) \sim \frac{n(m-1)}{2}$
- Unilateral rule: by looking at one vote there is no way to tell who x^* is; need to "guess" among $\sqrt{m}$ first alternatives
- Duple: by fixing only two alternatives the probability of getting x^* is $\frac{2}{m}$ ■