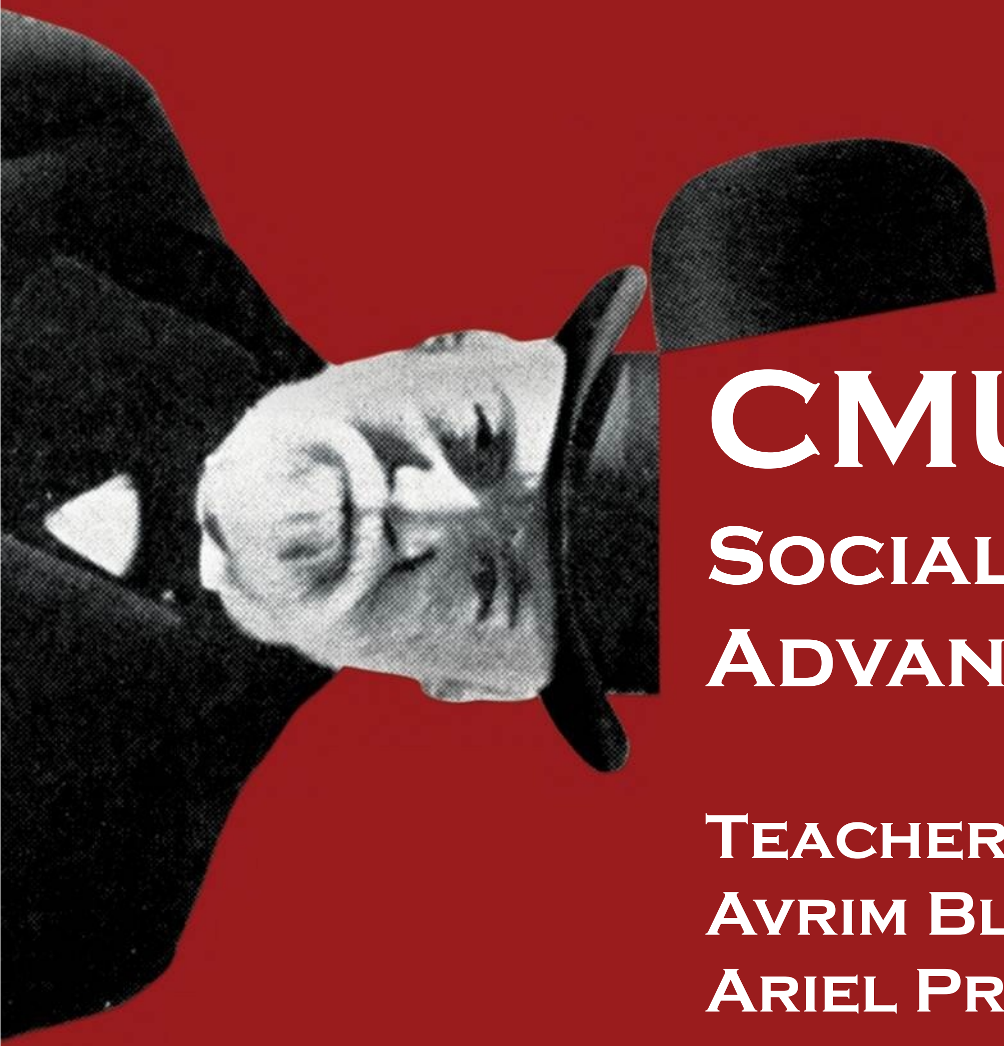

CMU 15-896

Social choice: Advanced manipulation

Teachers:
Avrim Blum
Ariel Procaccia (this time)

Recap

- A Complexity-theoretic barrier to manipulation
- Polynomial-time greedy alg can successfully 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,...
- Today: prove theorem about greedy algorithm, and then: NP-hardness is not enough

A GREEDY ALGORITHM

- Rank p in first place
- While there are unranked alternatives:
  - If there is an alternative that can be placed in next spot without preventing p from winning, place this alternative
  - Otherwise return false

When does the alg work?

- **Theorem [Bartholdi et al., SCW 89]:** Fix $i \in N$ and the votes of other voters. Let R be a rule s.t. $\exists$function $s(\prec_i, x)$ such that:
  - For every $\prec_i$ chooses a candidate that uniquely maximizes $s(\prec_i, x)$
  - $\{y\colon y \prec_i x\} \subseteq \{y\colon y \prec_i' x\} \Rightarrow s(\prec_i, x) \leq s(\prec_i', x)$

  Then the algorithm always decides R-Manipulation correctly
- Does plurality have a function s such that the winner uniquely maximizes score?
- At last: We prove the theorem on the board

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
$$D(f, g) = \Pr[f(\succ) \neq g(\succ)]$$
where the Pr is over uniformly random preference profiles $\succ$
- Distance to a set is defined as usual
- F_{dic} = set of dictatorships, $|F_{dic}| = n$

Quantitative G-S

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

Randomized voting rules

- Randomized voting rule: outputs a distribution over alternatives
- To think about successful manipulations we need utilities (assume strict preferences)
- $\succ_i$ is consistent with u_i if
$$x \succ_i y \Leftrightarrow u_i(x) > u_i(y)$$
- Strategyproofness: $\forall i \in N, \forall u_i, \forall \succ_{-i}, \forall \prec_i',$
$$\mathbb{E}\big[u_i\big(f(\prec)\big)\big] \geq \mathbb{E}\big[u_i\big(f(\prec_i', \prec_{-i})\big)\big]$$
where $\succ_i$ 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 $\succ$, $\Pr[f(\succ) = f_j(\succ)] = \alpha_j$.
- **Theorem [Gibbard 1977]:** randomized rule is strategyproof [vote: if / only if / iff] it is a probability mixture over unilaterals and duples

Randomization+approximation

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

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

Approximating Borda

- Choosing an alternative at random gives a $\frac{1}{2}$ approximation because $\sum_{x \in A} \frac{1}{m} \mathrm{sc}(\succ, x) = \frac{1}{m} \cdot \frac{nm(m-1)}{2} = \frac{n(m-1)}{2}$
- **Theorem [P 2010]:** No strategyproof randomized voting rule can approximate Borda to a factor of $\frac{1}{2} + \omega\left(\frac{1}{\sqrt{m}}\right)$

Yao's minimax principle

	$\succ^1$	…	…	…	…	$\succ^t$
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

- Choose $x^* \in A$ uniformly at random
- Each voter i chooses a random number $k_i \in \{1, \ldots, \sqrt{m}\}$ 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

- $\mathrm{sc}(\succ, x^*) \geq n(m - \sqrt{m})$
- For $x \in A \setminus \{x^*\}, sc(\succ, x) \sim \frac{n(m-1)}{2}$
- 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 $\frac{2}{m}$