CSC2556

Lecture 2

Manipulation in Voting

Credit for many visuals: Ariel D. Procaccia

Recap

- Voting
  - n voters, m alternatives
  - Each voter i expresses a ranked preference $\succ_i$
  - Voting rule f
    - Takes as input the collection of preferences $\vec{\succ}$
    - Returns a single alternative

- A plethora of voting rule
  - Plurality, Borda count, STV, Kemeny, Copeland, maximin, …

Incentives

- Can a voting rule incentivize voters to truthfully report their preferences?

- Strategyproofness
  - A voting rule is strategyproof if a voter cannot submit a false preference and get a more preferred alternative (under her true preference) elected, irrespective of the preferences of other voters.
  - Formally, a voting rule f is strategyproof if there is no preference profile $\overrightarrow{\succ}$, voter i, and false preference $\succ_i'$ s.t.

$$f\left(\overrightarrow{\succ}_{-i}, \succ_i'\right) \succ_i f\left(\overrightarrow{\succ}\right)$$

Strategyproofness

- None of the rules we saw are strategyproof!
- Example: Borda Count
  - In the true profile, b wins
  - Voter 3 can make a win by pushing b to the end

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

Winner
b

→

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

Winner
a

Borda's Response to Critics

My scheme is intended only for honest men!

Random 18th century French dude

Strategyproofness

- Are there any strategyproof rules?
  - Sure
- Dictatorial voting rule
  - The winner is always the most preferred alternative of voter i
- Constant voting rule
  - The winner is always the same
- Not satisfactory (for most cases)

Dictatorship

Constant function

Three Properties

- Strategyproof: Already defined. No voter has an incentive to misreport.

- Onto: Every alternative can win under some preference profile.

- Nondictatorial: There is no voter i such that $f(\overrightarrow{\succ})$ is always the alternative most preferred by voter i.

Gibbard-Satterthwaite

- Theorem: For $m \geq 3$, no deterministic social choice function can be strategyproof, onto, and nondictatorial simultaneously ☹

- Proof: We will prove this for $n = 2$ voters.
  - Step 1: Show that SP implies "strong monotonicity" [Assignment]
  - Strong Monotonicity (SM): If $f(\overrightarrow{\succ}) = a$, and $\overrightarrow{\succ}'$ is such that $\forall i \in N, x \in A$: $a \succ_i x \Rightarrow a \succ_i' x$, then $f(\overrightarrow{\succ}') = a$.
    - If a still defeats every alternative it defeated in every vote in $\overrightarrow{\succ}$, it should still win.

Gibbard-Satterthwaite

- Theorem: For $m \geq 3$, no deterministic social choice function can be strategyproof, onto, and nondictatorial simultaneously ☹

- Proof: We will prove this for $n = 2$ voters.
  - Step 2: Show that SP+onto implies "Pareto optimality" [Assignment]
  - Pareto Optimality (PO): If $a \succ_i b$ for all $i \in N$, then $f(\vec{\succ}) \neq b$.
    - If there is a different alternative that *everyone* prefers, your choice is not Pareto optimal (PO).

Gibbard-Satterthwaite

- Proof for n=2: Consider problem instance $I(a, b)$

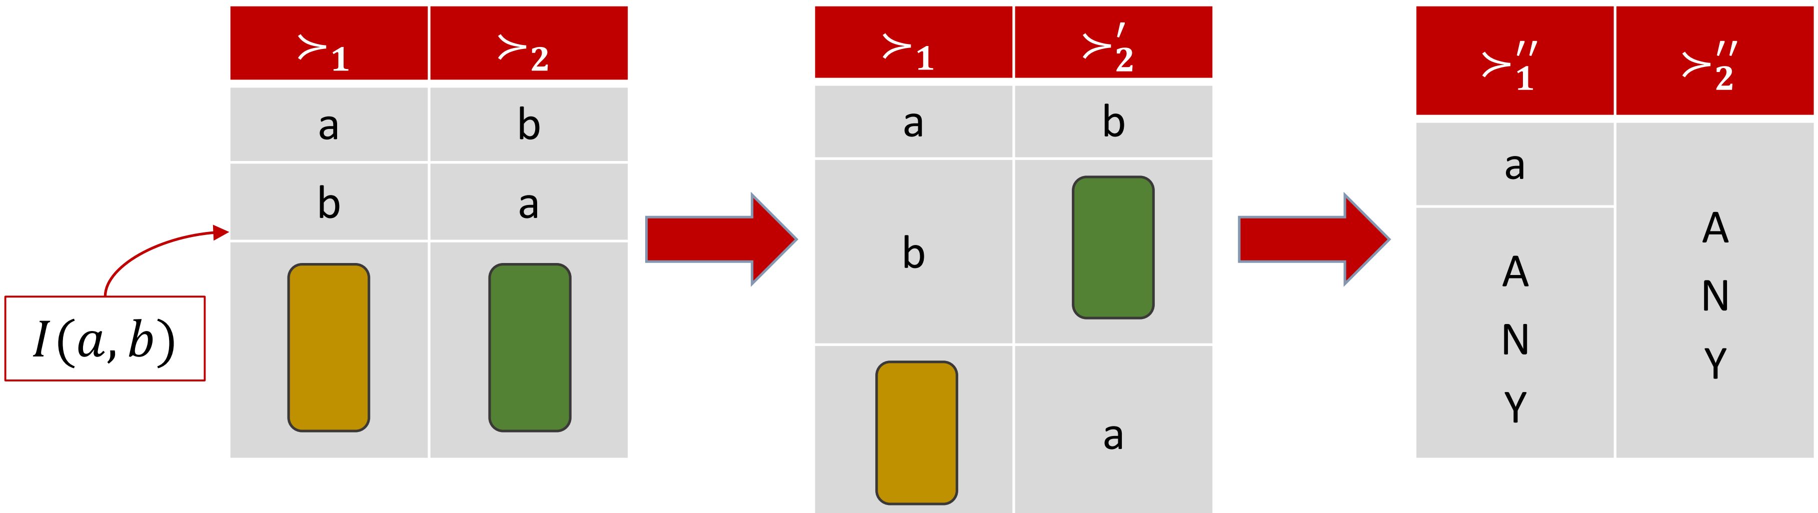

$f(\succ_1, \succ_2) \in \{a, b\}$

➢ PO

Say $f(\succ_1, \succ_2) = a$

$f(\succ_1, \succ_2') = a$

- PO: $f(\succ_1, \succ_2') \in \{\mathrm{a, b}\}$
- SP: $f(\succ_1, \succ_2') \neq b$

$f(\succ'') = a$

- Due to strong monotonicity

Gibbard-Satterthwaite

- Proof for n=2:
  - If f outputs a on instance $I(a, b)$, voter 1 can get a elected whenever she puts a first.
    - In other words, voter 1 becomes dictatorial for a.
    - Denote this by $D(1, a)$.
  - If f outputs b on $I(a, b)$
    - Voter 2 becomes dictatorial for b, i.e., we have $D(2, b)$.

- For every (a, b), we have either $D(1, a)$ or $D(2, b)$.

Gibbard-Satterthwaite

- Proof for n=2:
  - Fix a^* and b^*. Suppose $D(1, a^*)$ holds.
  - Then, we show that voter 1 is a dictator.
    - That is, $D(1, c)$ holds for every $c \neq a^*$ as well.
  - Take $c \neq a^*$. Because $|A| \geq 3$, there exists $d \in A \backslash \{a^*, c\}$.
  - Consider $I(c, d)$. We either have $D(1, c)$ or $D(2, d)$.
  - But $D(2, d)$ is incompatible with $D(1, a^*)$
    - Who would win if voter 1 puts a^* first and voter 2 puts d first?
  - Thus, we have $D(1, c)$, as required.
  - QED!

Circumventing G-S

- Restricted preferences (later in the course)
  - Not allowing all possible preference profiles
  - Example: single-peaked preferences
    - Alternatives are on a line (say 1D political spectrum)
    - Voters are also on the same line
    - Voters prefer alternatives that are closer to them
- Use of money (later in the course)
  - Require payments from voters that depend on the preferences they submit
  - Prevalent in auctions

Circumventing G-S

- Randomization (later in this lecture)
- Equilibrium analysis
  - How will strategic voters act under a voting rule that is not strategyproof?
  - Will they reach an "equilibrium" where each voter is happy with the (possibly false) preference she is submitting?
- Restricting information
  - Can voters successfully manipulate if they don't know the votes of the other voters?

Circumventing G-S

- Computational complexity
  - So we need to use a rule that is the rule is manipulable.
  - Can we make it NP-hard for voters to manipulate? [Bartholdi et al., SC&W 1989]
  - NP-hardness can be a good thing!

- f-Manipulation problem (for a given voting rule f):
  - Input: Manipulator i, alternative p, votes of other voters (non-manipulators)
  - Output: Can the manipulator cast a vote that makes p uniquely win under f?

Example: Borda

- Can voter 3 make a win?

1	2	3
b	b	
a	a	
c	c	
d	d	

→

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

A Greedy Algorithm

- Goal: The manipulator wants to make alternative p win uniquely

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

Example: Borda

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

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

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

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

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

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

Example: Copeland

1	2	3	4	5
a	b	e	e	a
b	a	c	c	
c	d	b	b	
d	e	a	a	
e	c	d	d	

Preference profile

	a	b	c	d	e
a	-	2	3	5	3
b	3	-	2	4	2
c	2	2	-	3	1
d	0	0	1	-	2
e	2	2	3	2	-

Pairwise elections

Example: Copeland

1	2	3	4	5
a	b	e	e	a
b	a	c	c	c
c	d	b	b	
d	e	a	a	
e	c	d	d	

Preference profile

	a	b	c	d	e
a	-	2	3	5	3
b	3	-	2	4	2
c	2	3	-	4	2
d	0	0	1	-	2
e	2	2	3	2	-

Pairwise elections

Example: Copeland

1	2	3	4	5
a	b	e	e	a
b	a	c	c	c
c	d	b	b	d
d	e	a	a	
e	c	d	d	

Preference profile

	a	b	c	d	e
a	-	2	3	5	3
b	3	-	2	4	2
c	2	3	-	4	2
d	0	1	1	-	3
e	2	2	3	2	-

Pairwise elections

Example: Copeland

1	2	3	4	5
a	b	e	e	a
b	a	c	c	c
c	d	b	b	d
d	e	a	a	e
e	c	d	d	

Preference profile

	a	b	c	d	e
a	-	2	3	5	3
b	3	-	2	4	2
c	2	3	-	4	2
d	0	1	1	-	3
e	2	3	3	2	-

Pairwise elections

Example: Copeland

1	2	3	4	5
a	b	e	e	a
b	a	c	c	c
c	d	b	b	d
d	e	a	a	e
e	c	d	d	b

Preference profile

	a	b	c	d	e
a	-	2	3	5	3
b	3	-	2	4	2
c	2	3	-	4	2
d	0	1	1	-	3
e	2	3	3	2	-

Pairwise elections

When does this work?

- Theorem [Bartholdi et al., SCW 89]:

  Fix voter i and votes of other voters. Let f be a rule for which ∃ function $s(\succ_i, x)$ such that:

  1. For every $\succ_i$, f chooses a candidate x that uniquely maximizes $s(\succ_i, x)$.
  2. $\{y : x \succ_i y\} \subseteq \{y : x \succ_i' y\} \Rightarrow s(\succ_i, x) \leq s(\succ_i', x)$

  Then the greedy algorithm solves f-MANIPULATION correctly.

- Question: What is the function s for plurality?

Proof of the Theorem

- Say the algorithm creates a partial ranking $\succ_i$ and then fails, i.e., every next choice prevents p from winning
- Suppose for contradiction that $\succ_i'$ could make p uniquely win
- $U \leftarrow$ alternatives not ranked in $\succ_i$
- $u \leftarrow$ highest ranked alternative in U according to $\succ_i'$
- Complete $\succ_i$ by adding u next, and then other alternatives arbitrarily

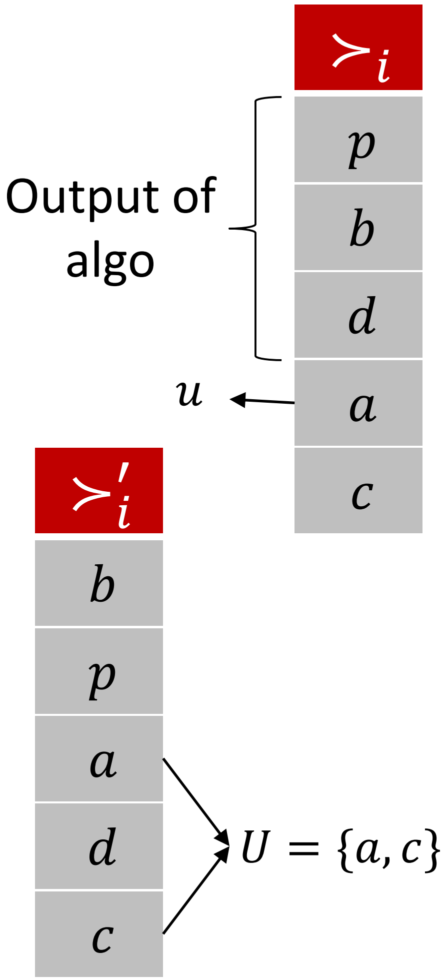

Proof of the Theorem

- $s(\succ_i, p) \geq s(\succ_i', p)$
  - Property 2
- $s(\succ_i', p) > s(\succ_i', u)$
  - Property 1 & p wins under $\succ_i'$
- $s(\succ_i', u) \geq s(\succ_i, u)$
  - Property 2
- Conclusion
  - Putting u in the next position wouldn't have prevented p from winning
  - So the algorithm should have continued

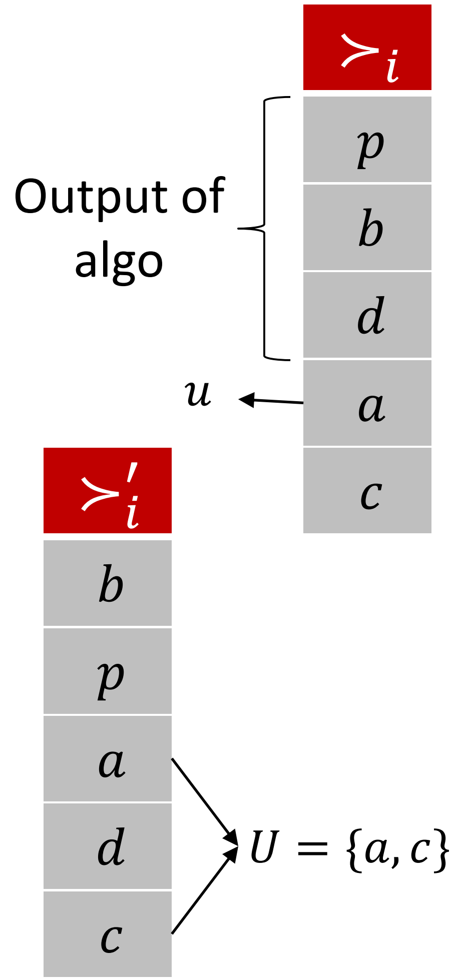

Hard-to-Manipulate Rules

- Natural rules
  - Copeland with second-order tie breaking [Bartholdi et al. SCW 89]
    - In case of a tie, choose the alternative for which the sum of Copeland scores of defeated alternatives is *the largest*
  - STV [Bartholdi & Orlin, SCW 91]
  - Ranked Pairs [Xia et al., IJCAI 09]
    - Iteratively lock in pairwise comparisons by their margin of victory (largest first), ignoring any comparison that would form cycles.
    - Winner is the top ranked candidate in the final order.
- Can also "tweak" easy to manipulate voting rules [Conitzer & Sandholm, IJCAI 03]

Example: Ranked Pairs

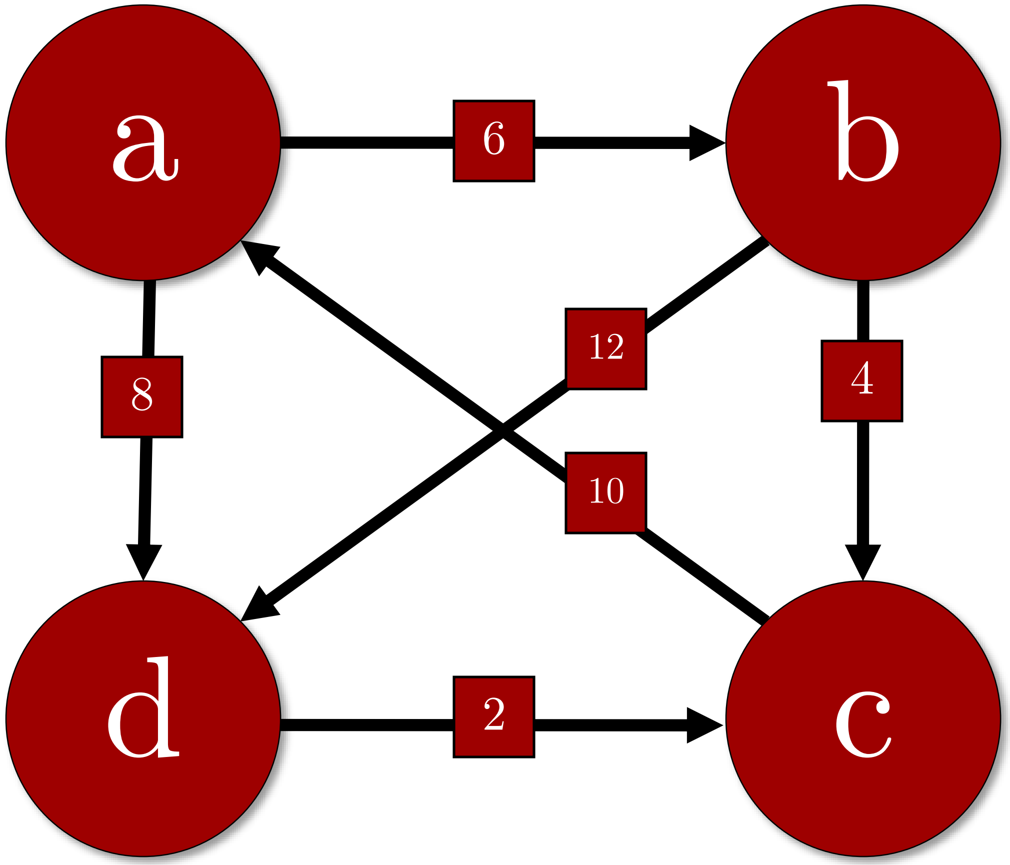

Example: Ranked Pairs

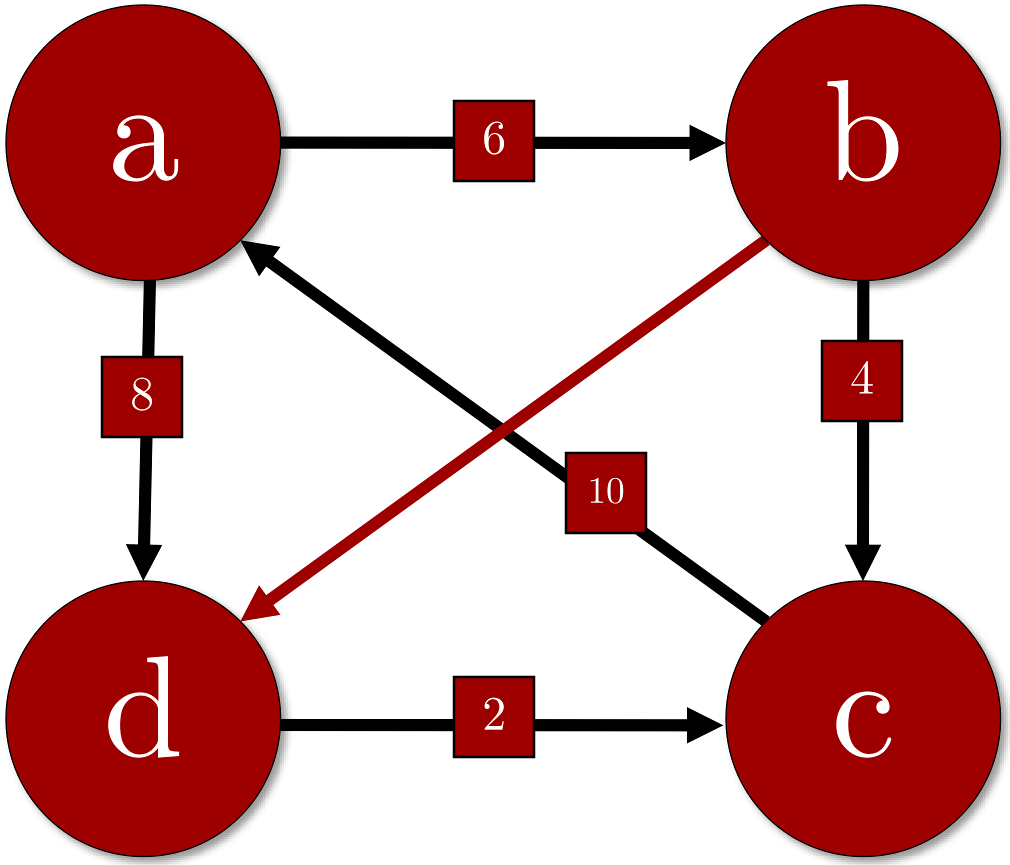

Example: Ranked Pairs

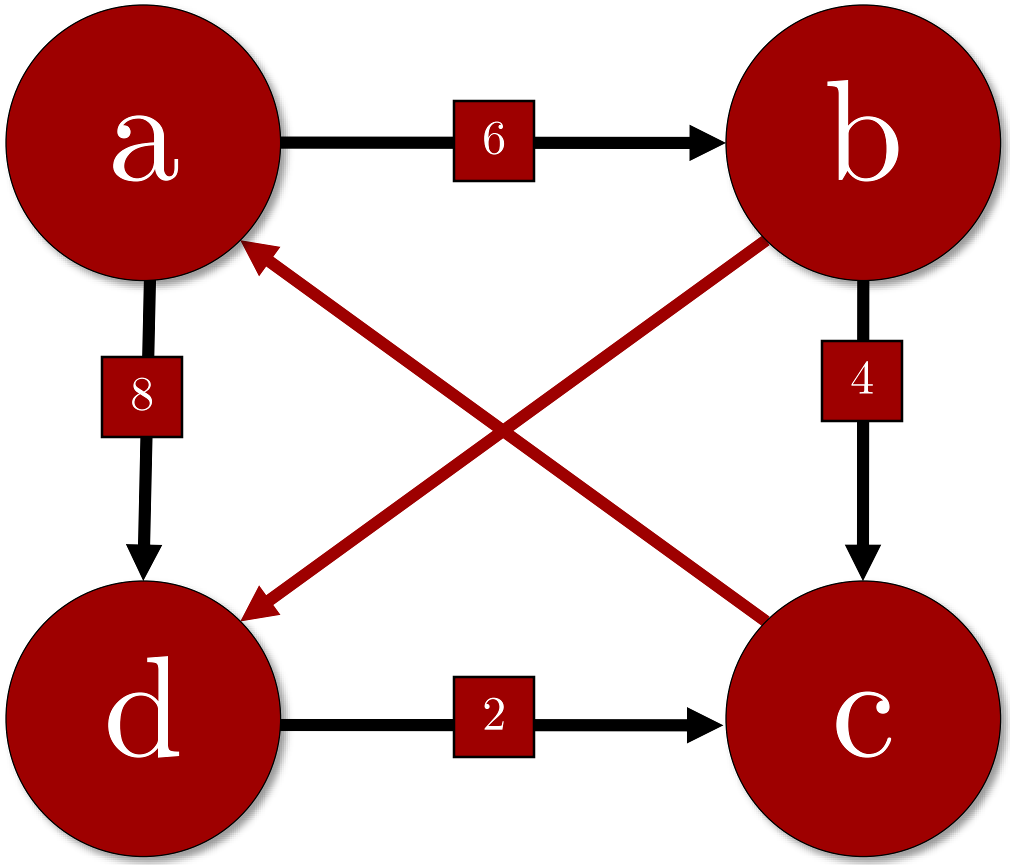

Example: Ranked Pairs

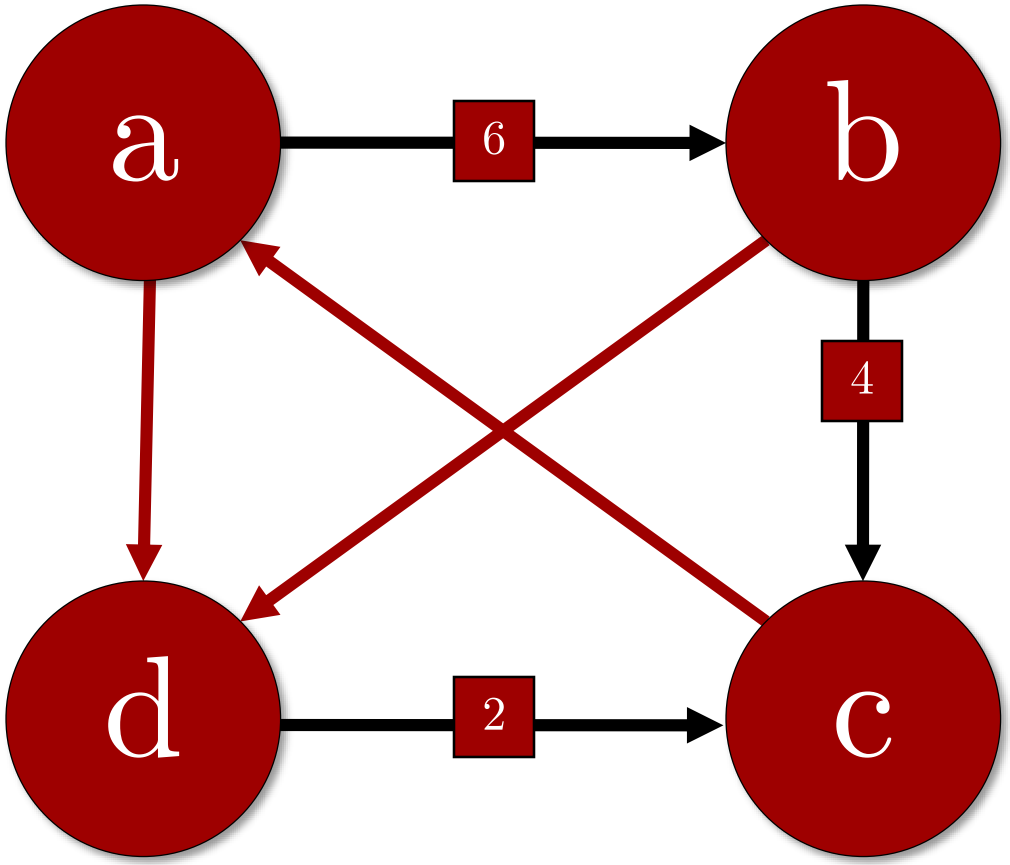

Example: Ranked Pairs

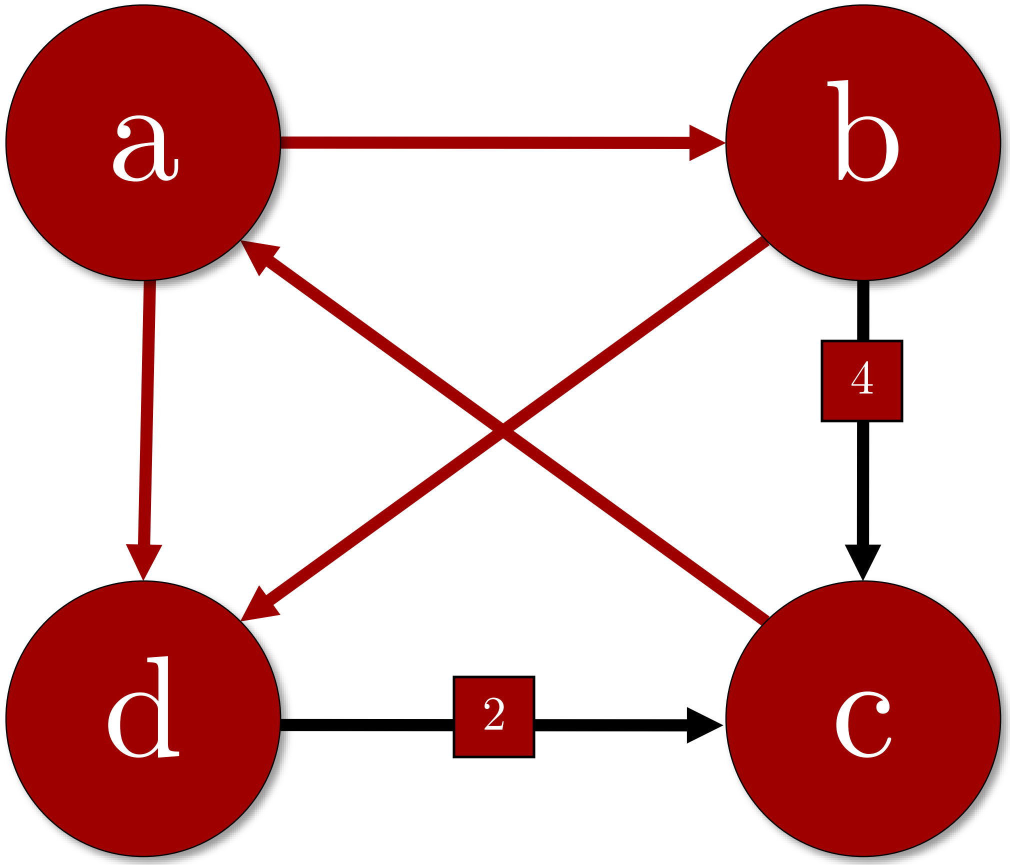

Example: Ranked Pairs

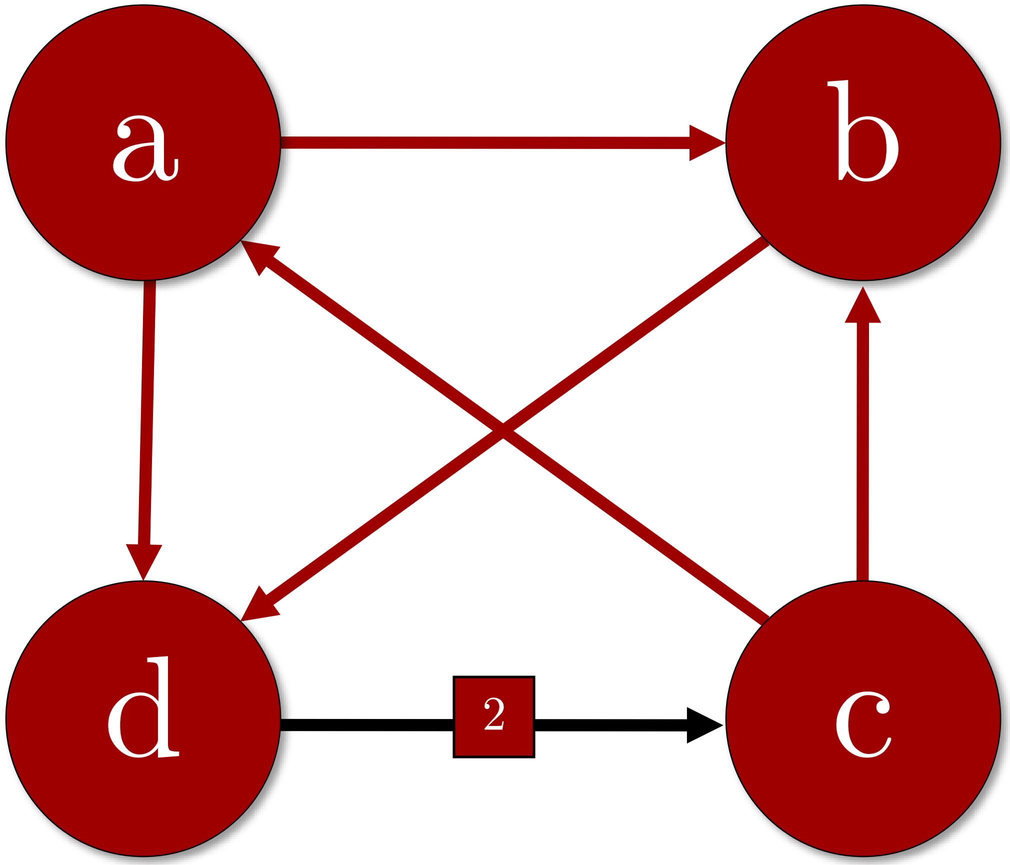

Example: Ranked Pairs

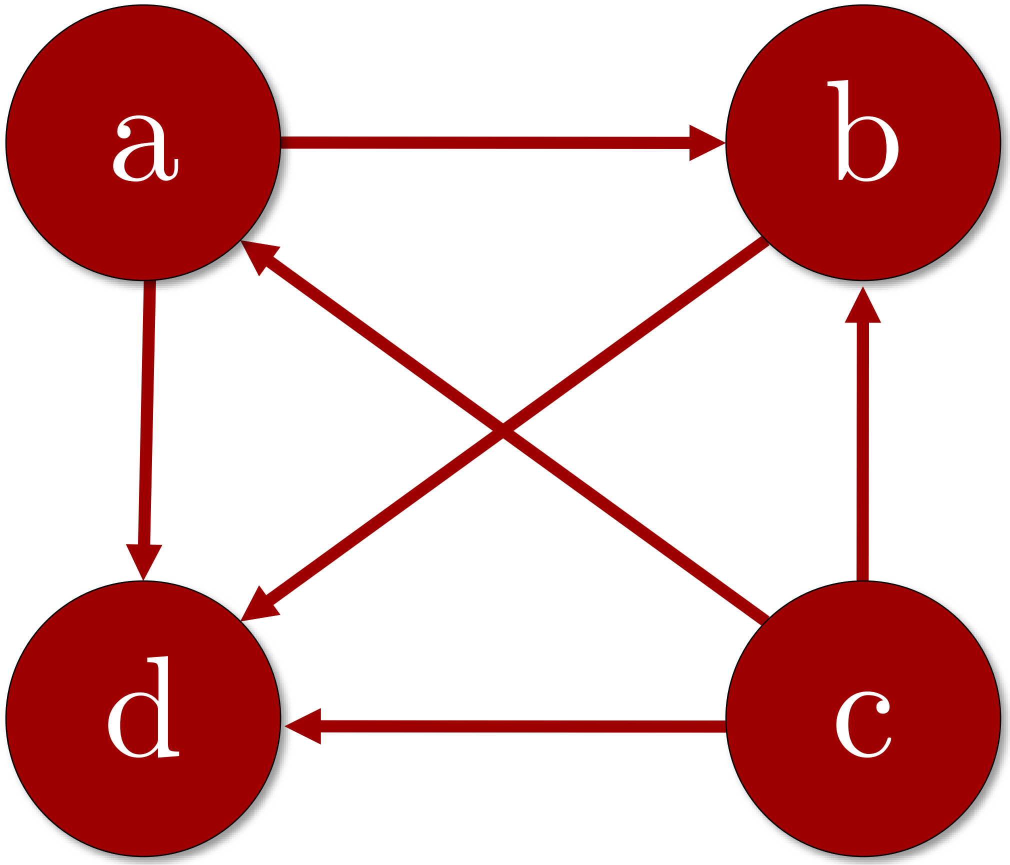

Randomized Voting Rules

- Take as input a preference profile, output a distribution over alternatives
- To think about successful manipulations, we need numerical utilities
- $\succ_i$ is consistent with u_i if
$$a \succ_i b \Leftrightarrow u_i(a) > u_i(b)$$
- Strategyproofness: For all i, u_i, $\vec{\succ}_{-i}$, and $\succ_i'$
$$\mathbb{E}\left[u_i\left(f(\vec{\succ})\right)\right] \ge \mathbb{E}\left[u_i\left(f(\vec{\succ}_{-i}, \succ_i')\right)\right]$$
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 contains at most two alternatives

- A probability mixture f over rules $f_1, \dots, f_k$ is a rule given by some probability distribution $(\alpha_1, \dots, \alpha_k)$ s.t. on every profile $\overrightarrow{\succ}$, f returns $f_j(\overrightarrow{\succ})$ w.p. α_j.

Randomized Voting Rules

- Theorem [Gibbard 77]:
  A randomized voting rule is strategyproof only if it is a probability mixture over unilaterals and duples.

- Example:
  - With probability 0.5, output the top alternative of a randomly chosen voter
  - With the remaining probability 0.5, output the winner of the pairwise election between a^* and b^*

- Question: What is a probability mixture over unilaterals and duples that is *not* strategyproof?

Approximating Voting Rules

- Idea: Can we use strategyproof voting rules to approximate popular voting rules?
- Fix a rule (e.g., Borda) with a clear notion of score denoted $\text{sc}(\vec{\succ}, a)$
- A randomized voting rule f is a c-approximation to sc if for every profile $\vec{\succ}$

$$\frac{\mathbb{E}[\text{sc}\left(\vec{\succ}, f(\vec{\succ})\right)}{\max_a \text{sc}(\vec{\succ}, a)} \geq c$$

Approximating Borda

- Question: How well does choosing a random alternative approximate Borda?
  1. $\Theta(1/n)$
  2. $\Theta(1/m)$
  3. $\Theta(1/\sqrt{m})$
  4. $\Theta(1)$

- Theorem [Procaccia 10]:
  No strategyproof voting rule gives ${}^{1}/_{2} + \omega\left({}^{1}/_{\sqrt{m}}\right)$ approximation to Borda.

Interlude: Zero-Sum Games

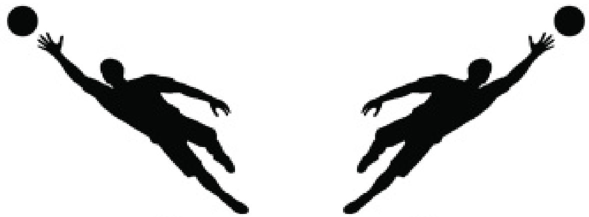

-1	1
1	-1

Interlude: Minimiax Strategies

- A minimax strategy for a player is
  - a (possibly) randomized choice of action by the player
  - that minimizes the expected loss (or maximizes the expected gain)
  - in the worst case over the choice of action of the other player

- In the previous game, the minimax strategy for each player is $(1/2, 1/2)$. Why?

Interlude: Minimiax Strategies

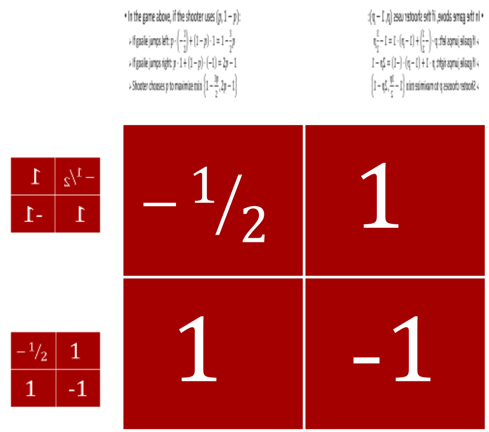

- In the game above, if the shooter uses $(p, 1-p)$:
  - If goalie jumps left: $p \cdot \left(-\frac{1}{2}\right) + (1-p) \cdot 1 = 1 - \frac{3}{2}p$
  - If goalie jumps right: $p \cdot 1 + (1-p) \cdot (-1) = 2p - 1$
  - Shooter chooses p to maximize $\min \left\{1 - \frac{3p}{2}, 2p - 1\right\}$

Interlude: 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

- Rows as inputs
- Columns as deterministic algorithms
- Cell numbers = running times
- Best randomized algorithm
  - Minimax strategy for the column player

$$\min_{rand\ algo} \max_{input} E[time] =$$

$$\max_{dist\ over\ inputs} \min_{det\ algo} E[time]$$

Yao's Minimax Principle

- To show a lower bound T on the best worst-case running time achievable through randomized algorithms:
  - ➢ Show a "bad" distribution over inputs D such that every deterministic algorithm takes time at least T on average, when inputs are drawn according to D

$$\min_{rand\ algo} \max_{input} E[time] =$$

$$\max_{dist\ over\ inputs} \min_{det\ algo} E[time]$$

Randomized Voting Rules

	$\overrightarrow{\precsim}^1$	…	…	…	…	$\overrightarrow{\precsim}^t$
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	$\frac{13}{15}$	…	…	…	…	$\frac{17}{21}$

Randomized Voting Rules

- Rows = unilaterals and duples
- Columns = preference profiles
- Cell numbers = approximation ratios

- The expected ratio of the best strategyproof rule (by Gibbard's theorem, distribution over unilaterals and duples) is at most…
  - The expected ratio of the best unilateral or duple rule when profiles are drawn from a "bad" distribution D

A Bad Distribution

- $m = n + 1$
- Choose a random alternative x^*
- Each voter i chooses a random number $k_i \in \{1, \ldots, \sqrt{m}\}$ and places 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

- Question: What is the best lower bound on $\text{sc}(\vec{\succ}, x^*)$ that holds for every profile $\vec{\succ}$ generated under this distribution?

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

A Bad Distribution

- How bad are other alternatives?
  - For every other alternative x, $\text{sc}(\overrightarrow{\succ}, x) \sim \frac{n(m-1)}{2}$
- How surely can a unilateral/duple rule return x^*?
  - Unilateral: By only looking at a single vote, the rule is essentially guessing x^* among the first $\sqrt{m}$ positions, and captures it with probability at most $1/\sqrt{m}$.
  - Duple: By fixing two alternatives, the rule captures x^* with probability at most $2/m$.
- Putting everything together…

Quantitative GS Theorem

- Regarding the use of NP-hardness to circumvent GS
  - NP-hardness is hardness in the worst case
  - What happens in the average case?

- Theorem [Mossel-Racz '12]:
  For every voting rule that is at least ϵ-far from being a dictatorship or having range of size 2, the probability that a profile chosen uniformly at random admits a manipulation is at least $p(n, m, {}^{1}/_{\epsilon})$ for some polynomial p.

Coalitional Manipulations

- What if multiple voters collude to manipulate?
  - The following result applies to a wide family of voting rules called "generalized scoring rules".

- Theorem [Conitzer-Xia '08]:

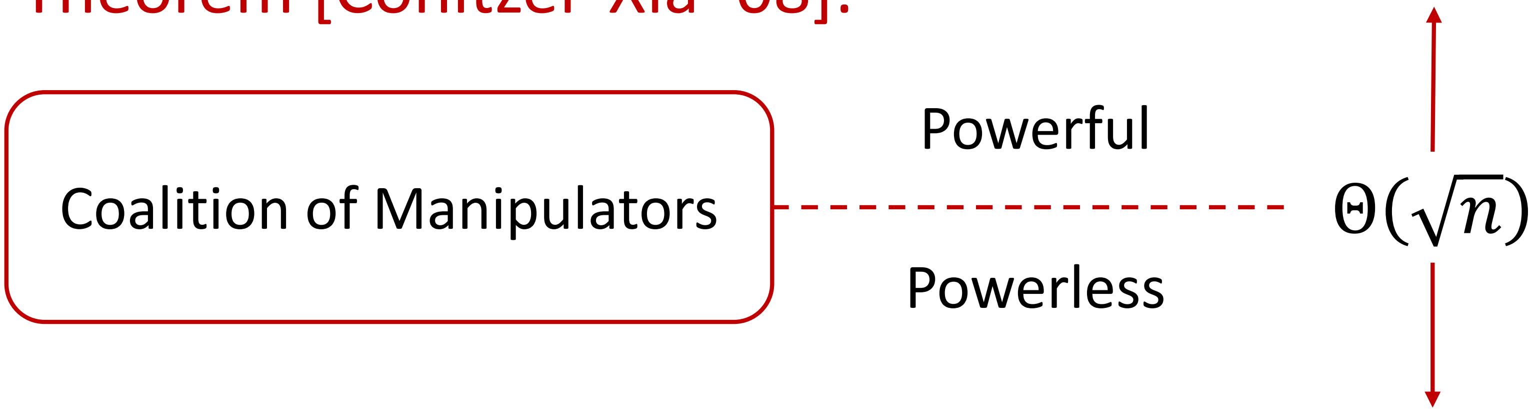

Powerful = can manipulate with high probability

Interesting Tidbit

- Detecting a manipulable profile versus finding a beneficial manipulation

- Theorem [Hemaspaandra, Hemaspaandra, Menton '12]
  If integer factoring is NP-hard, then there exists a generalized scoring rule for which:
  - We can efficiently check if there exists a beneficial manipulation.
  - But finding such a manipulation is NP-hard.

Next Lecture

- Frameworks to compare voting rules
  - Even if we assume that voters will reveal their true preferences, we still don't know if there is one "right" way to choose the winner.
  - There are reasonable profiles where most prominent voting rules return different winners [Assignment]