CSC2556

Lecture 3

Voting III

Credit for many visuals: Ariel D. Procaccia

Recap: Complexity for Good

- Computational complexity
  - 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?

Recap: 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.

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 $\exists$ function $s(\succ_i, x)$ such that:

  1. For every $\succ_i$, f chooses candidates maximizing $s(\succ_i\,,\ \cdot)$
  2. $\{y : x \succ_i y\} \subseteq \{y : x \succ_i' y\} \Rightarrow s(\succ_i, x) \le s(\succ_i', x)$

  Then the greedy algorithm solves f-MANIPULATION correctly.

- Question: What is the function s for the plurality rule?

Proof of the Theorem

- Suppose for contradiction:
  - Algo creates a partial ranking $\succ_i$ and then fails, i.e., every next choice prevents p from winning
  - But $\succ_i'$ could have made 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

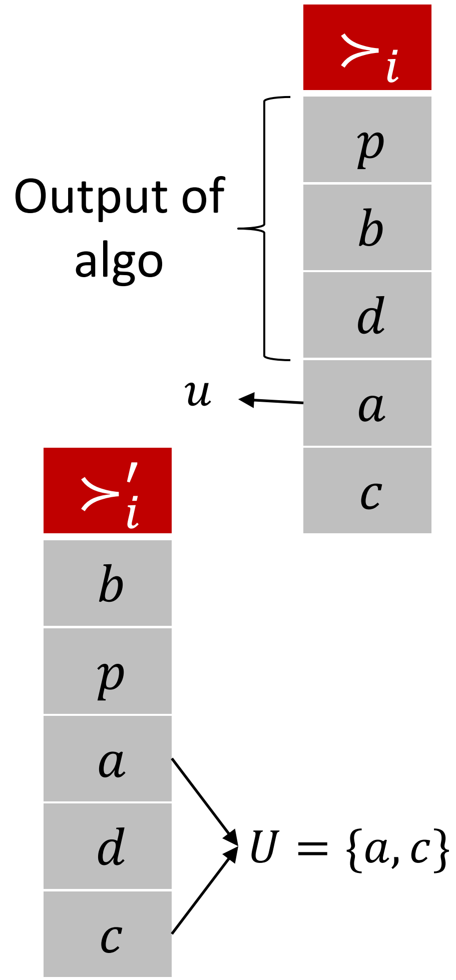

Proof of the Theorem

- $s(\succ_i, p) \geq s(\succ_i', p)$
  - Property 2
- $s(\succ_i', p) > s(\succ_i', u)$
  - Property 1 & p uniquely 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

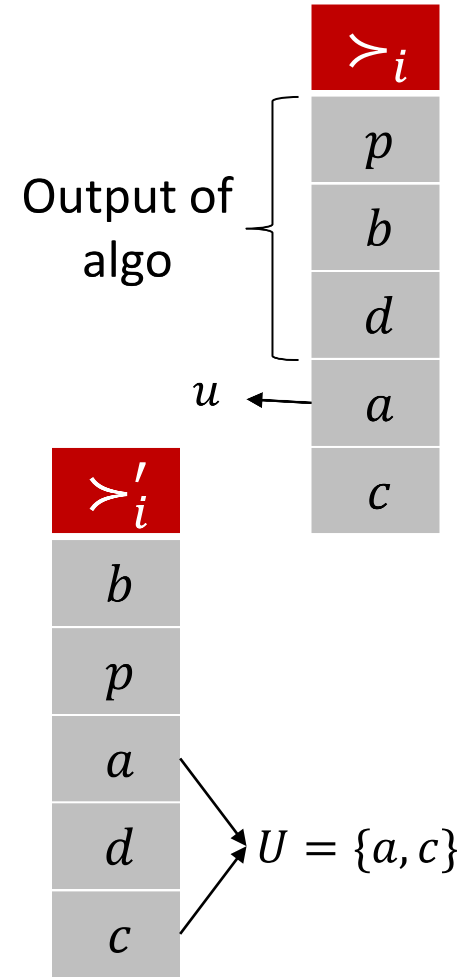

Hard-to-Manipulate Rules

- 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

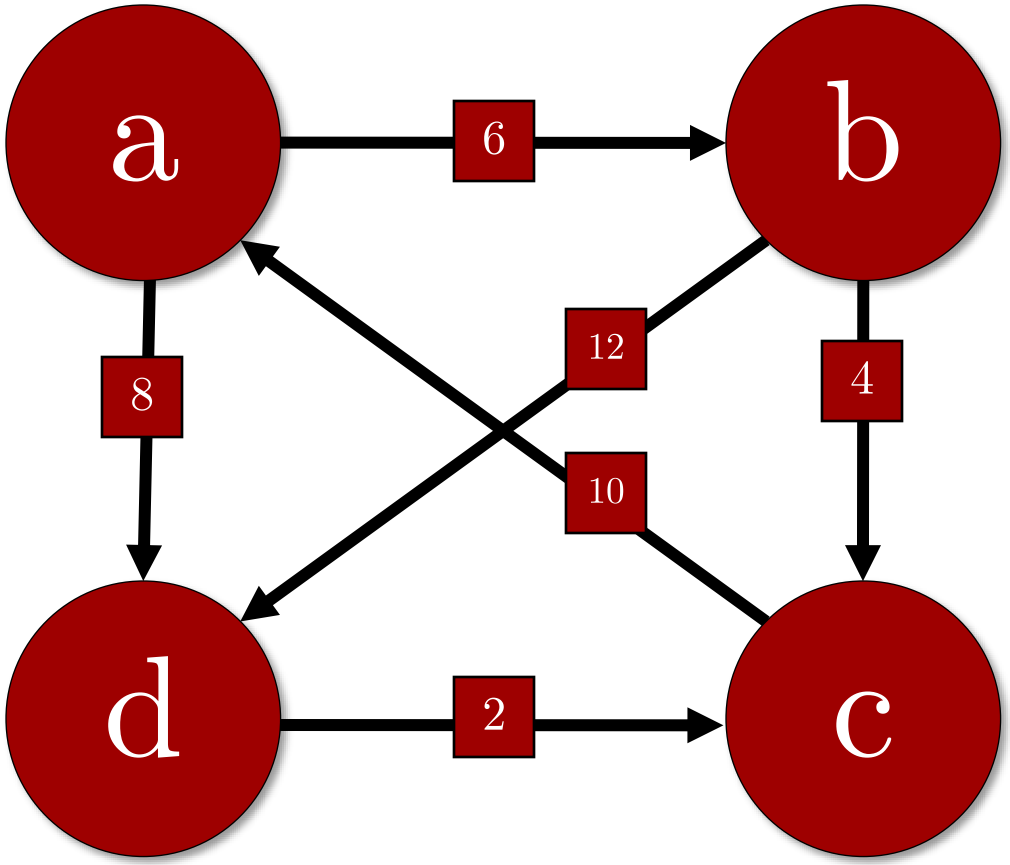

Example: Ranked Pairs

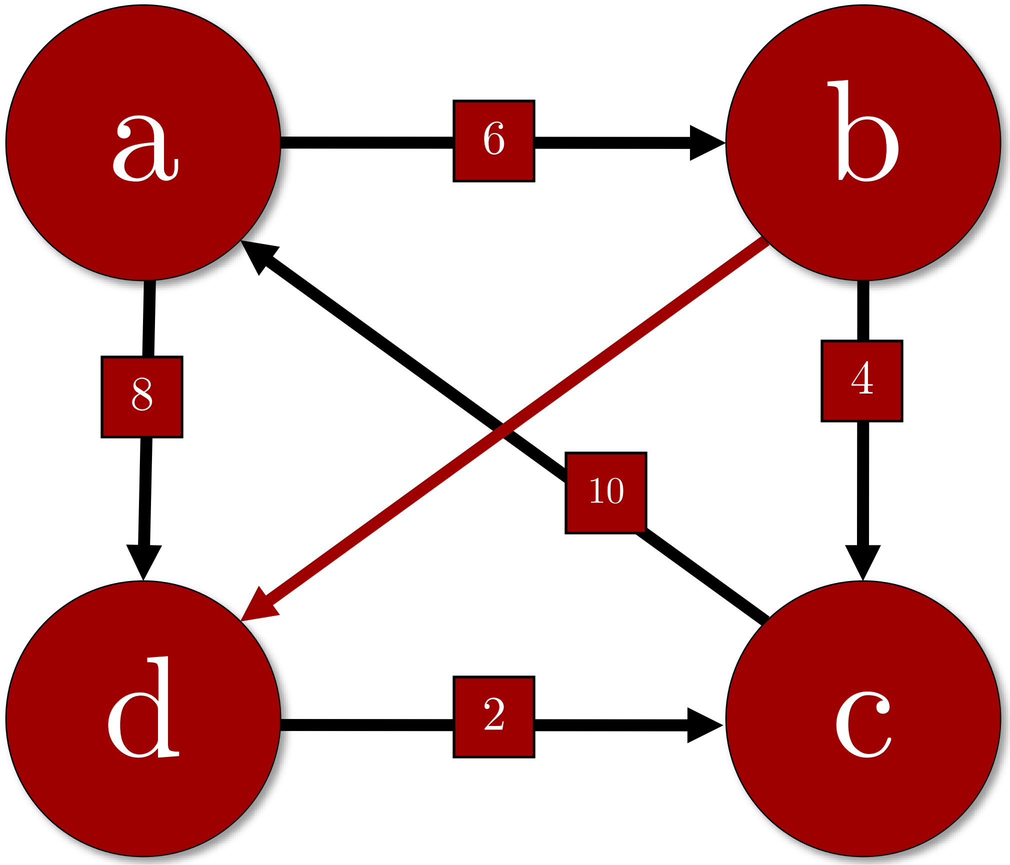

Example: Ranked Pairs

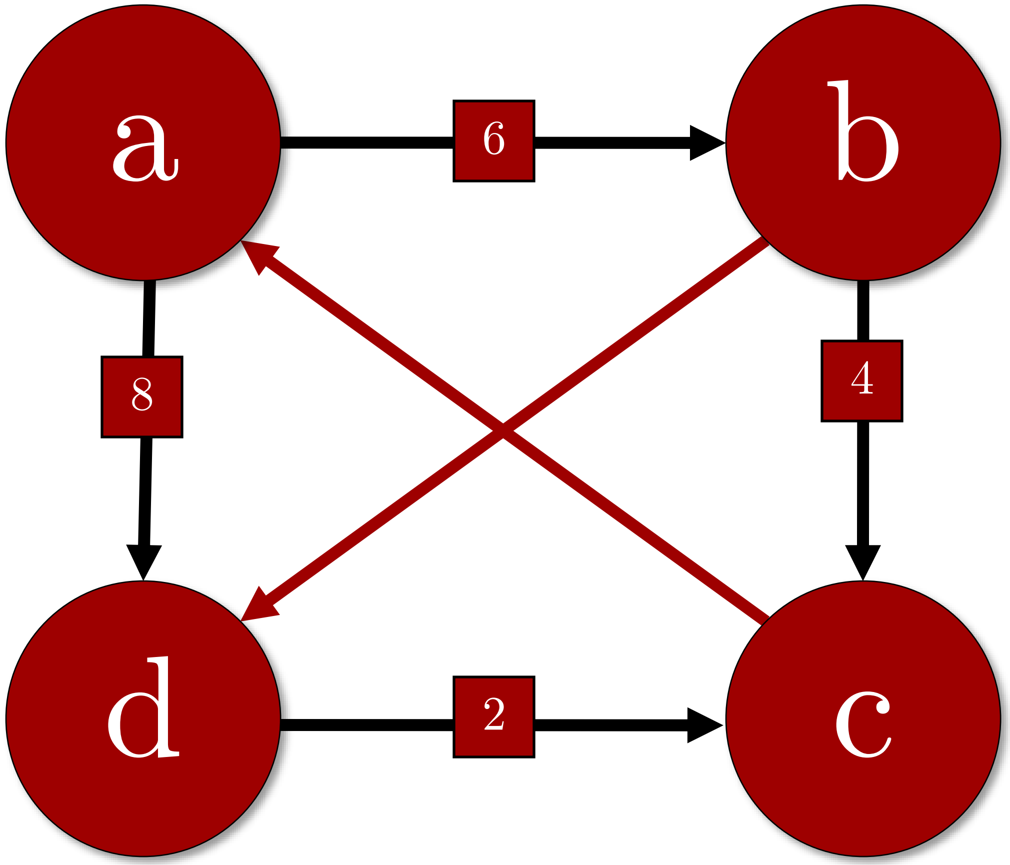

Example: Ranked Pairs

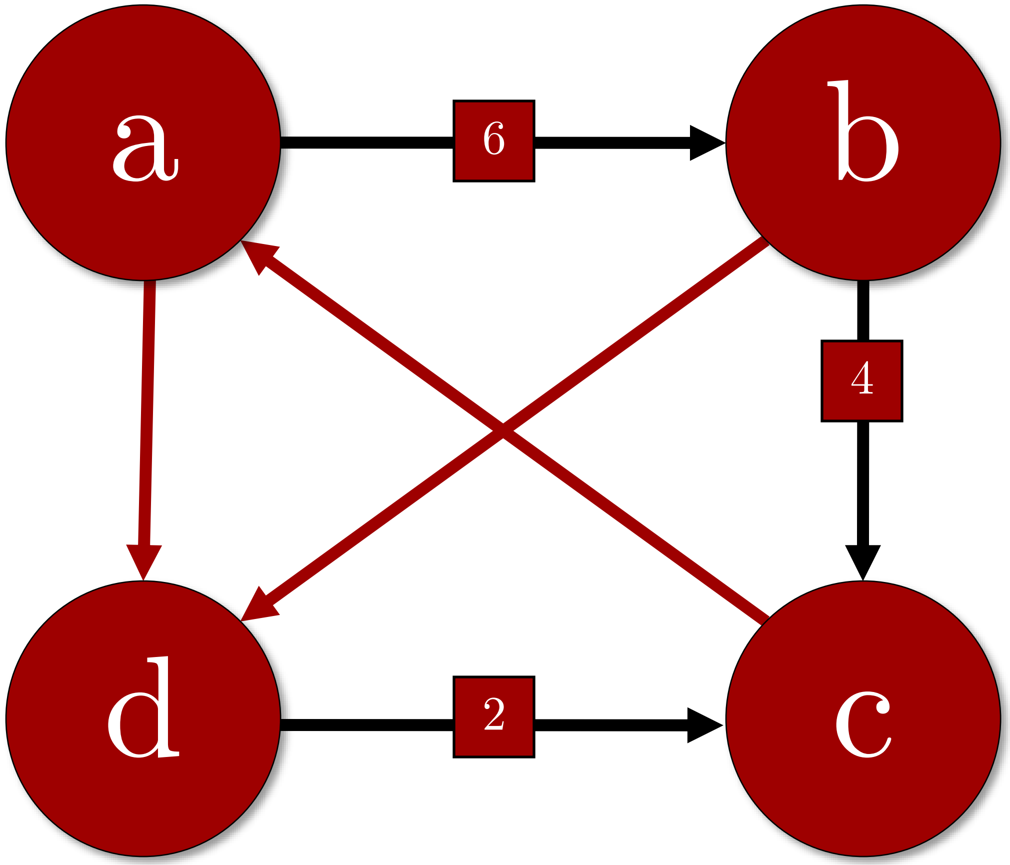

Example: Ranked Pairs

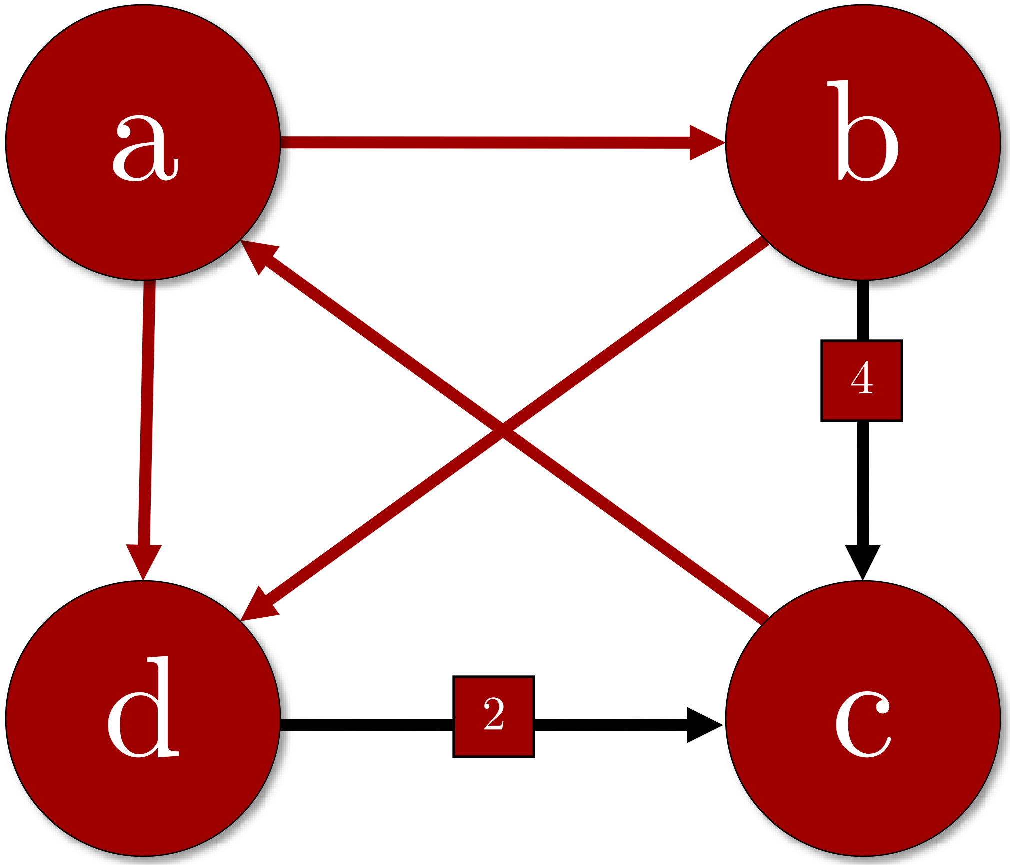

Example: Ranked Pairs

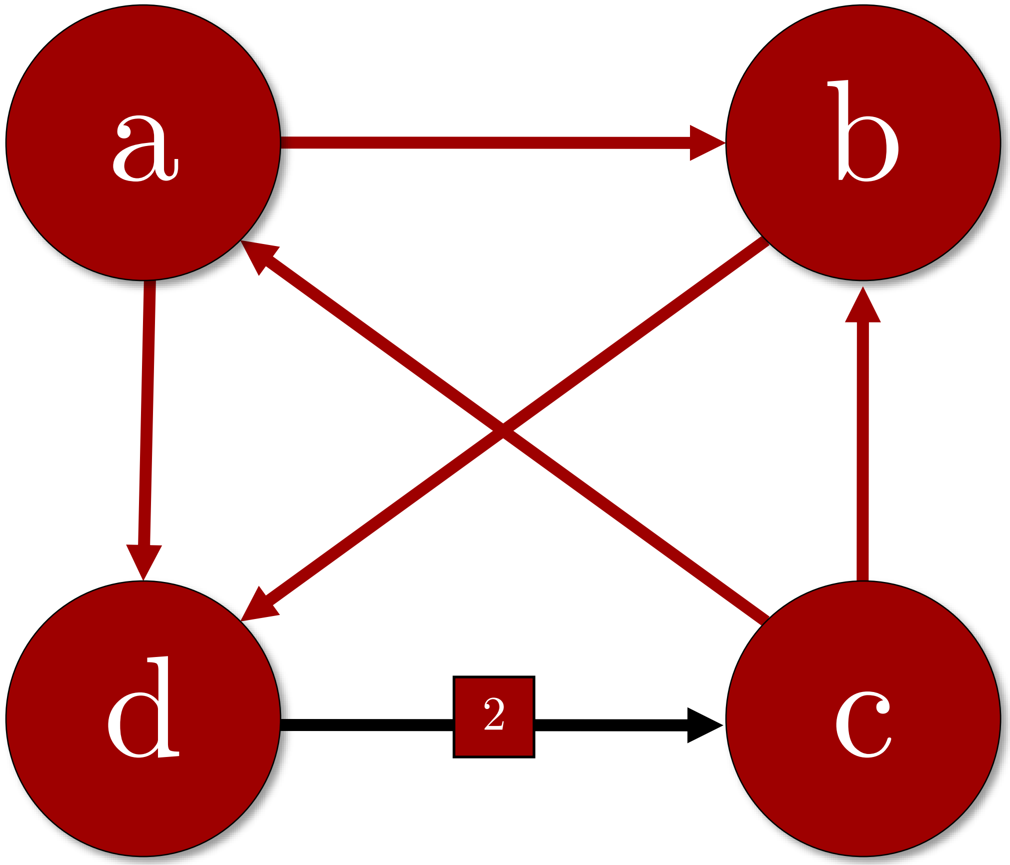

Example: Ranked Pairs

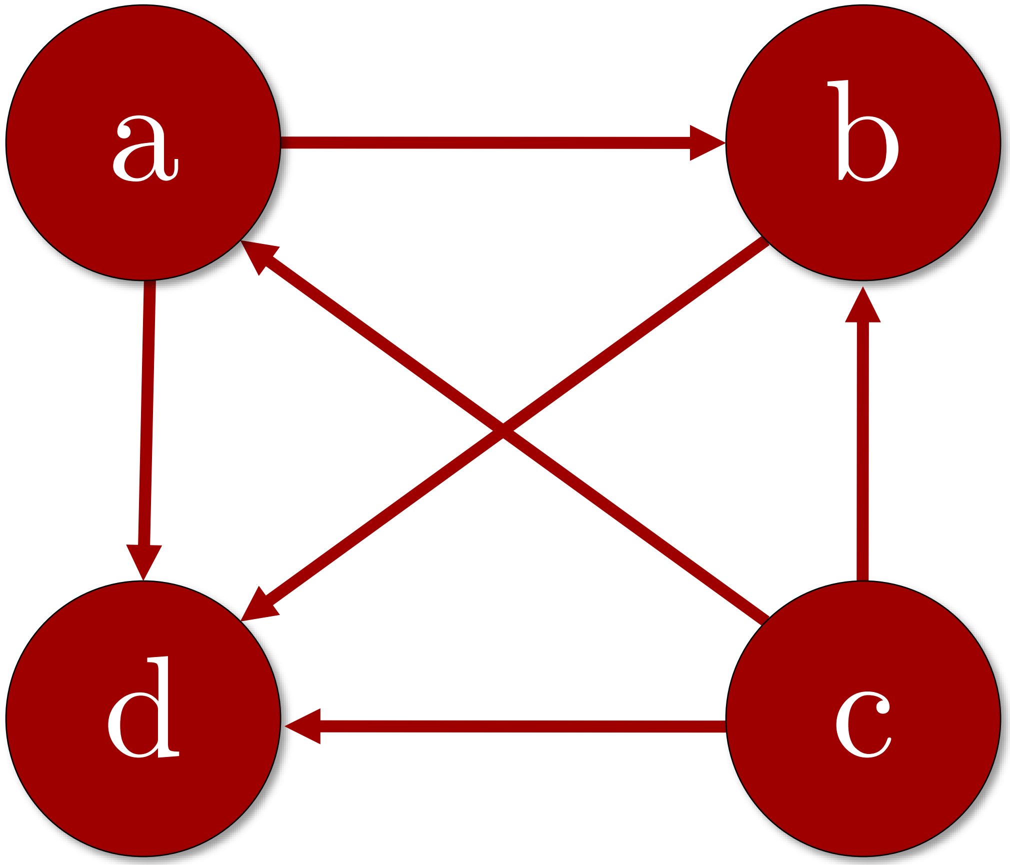

Randomized Voting Rules

- 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 $\succ_i'$

$$\mathbb{E}\left[u_i\left(f(\vec{\succ})\right)\right] \geq \mathbb{E}\left[u_i\left(f(\vec{\succ}_{-i}, \succ_i')\right)\right]$$

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

- Question:
  - What is a unilateral rule that is not strategyproof?
  - What is a duple rule that is not strategyproof?

Randomized Voting Rules

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

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 $\mathrm{sc}(\overrightarrow{\succ}, a)$
- A randomized voting rule f is a c-approximation to sc if for every profile $\overrightarrow{\succ}$

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

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

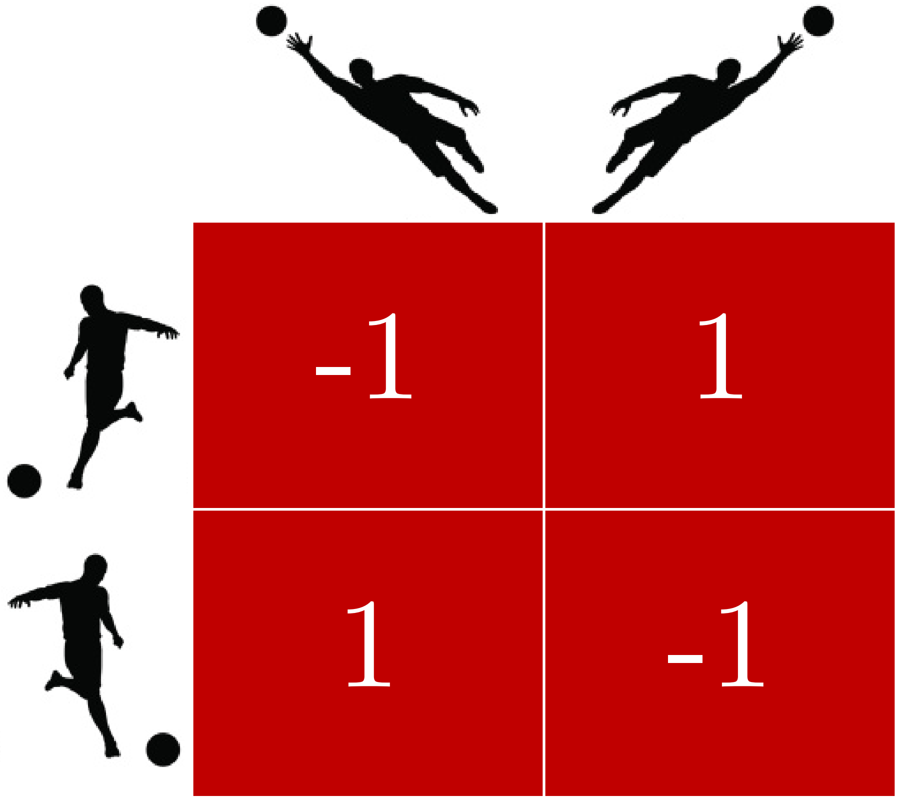

	-1	1
	1	-1

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

- Intuition
  - Suppose I were to act first…
  - …and the other player could observe my strategy and respond to it (thus picking a response that is worst case for me)
  - …then which randomized choice would I make?

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

Interlude: Minimiax Strategies

$-{}^{1}/_{2}$	1
1	-1

- 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

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

Randomized Voting Rules

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

- The expected ratio for the best distribution over unilaterals and duples on the worst profile is equal to the expected ratio of the best unilateral or duple rule when the profiles are drawn from the worst distribution D
  - Best ratio of any best strategyproof rule $\leq$ best ratio of any distribution over unilaterals and duples $\leq$ ratio of the best unilateral/duple under some bad distribution

Back to Borda

- Assume $m = n + 1$
- A bad distribution:
  - Choose a random alternative x^*
  - Each voter i chooses a random number $k_i \in \{1, \dots, \sqrt{m}\}$ 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

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

Back to Borda

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

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

Back to Borda

- How bad are other alternatives?
  - For every other alternative x, $\text{sc}(\vec{\succ}, x) \sim \frac{n(m-1)}{2}$
- 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 uniformly random profile 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]:

Coalition of Manipulators

Powerful

$\Theta(\sqrt{n})$

Powerless

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.