Fair and Efficient Social Decision-Making

CSCI 699

Voting: Axioms II

Evi Micha

Credit for the slides and most of visuals: Nisarg Shah and Ariel D. Procaccia

Which rule to use?

- We just introduced infinitely many rules
  - (Recall positional scoring rules…)
- How do we know which is the “right” rule to use?
  - Various approaches
  - Axiomatic, statistical, utilitarian, …
- How do we ensure good incentives without using money?
  - Bad luck!

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 for every preference profile $\overrightarrow{\succ}$, voter i, and preference $\succ_i'$, we have

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

  - Question: What is the relation between $f(\overrightarrow{\succ})$ and $f(\overrightarrow{\succ}_{-i}, \succ_i')$ according to $\succcurlyeq_i'$?

Strategyproofness

- Question: Is Borda Count strategyproof?

- Example:
  - 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

- None of the rules we saw are strategyproof!

- 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 is strategyproof, onto, and nondictatorial simultaneously ☹

- Proof: We will prove this for $n = 2$ voters.
  - Step 1: Show that SP ⇒ "strong monotonicity" [Assignment]
  - Strong Monotonicity (SM): If $f(\vec{\succ}) = a$, and $\vec{\succ}'$ is such that $\forall i \in N, x \in A$: $a \succ_i x \Rightarrow a \succ_i' x$, then $f(\vec{\succ}') = a$.
    - If, for each i, the set of alternatives defeated by a in $\succ_i'$ is a superset of what it defeats in $\succ_i$, then if it was winning under $\vec{\succ}$, it should also win under $\vec{\succ}'$

Gibbard-Satterthwaite

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

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

Gibbard-Satterthwaite

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

$I(a,b)$

$\succ_1$	$\succ_2$
a	b
b	a

$\succ_1$	$\succ'_2$
a	b
b	
	a

$\succ''_1$	$\succ''_2$
a	A N Y
A N Y	

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

➢ PO

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

$f(\succ_1, \succ'_2) = a$

- PO: $f(\succ_1, \succ'_2) \in \{\mathrm{a}, \mathrm{b}\}$
- SP: $f(\succ_1, \succ'_2) \neq b$

$f(\succ'') = \mathrm{a}$

➢ SM

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 property by the notation $D(1, a)$.
  - 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), f either satisfies the property $D(1, a)$ or the property $D(2, b)$.
  - We're not done! (Why?)

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, $D(1, c)$ also holds for every $c \neq a^*$
  - Take $c \neq a^*$. Because $|A| \geq 3$, there exists $d \in A \backslash \{a^*, c\}$
  - Consider $I(c, d)$; f sastisifies either $D(1, c)$ 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 ■

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

Circumventing G-S

- Randomization (later in the course)
- 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 required for manipulation
  - Can voters successfully manipulate if they don't know the votes of the other voters?

Circumventing G-S

- 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?

Example: Borda

- Can voter 3 make a win?
  - Yes

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 candidates maximizing $s(\succ_i\ , \ \cdot)$
  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 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

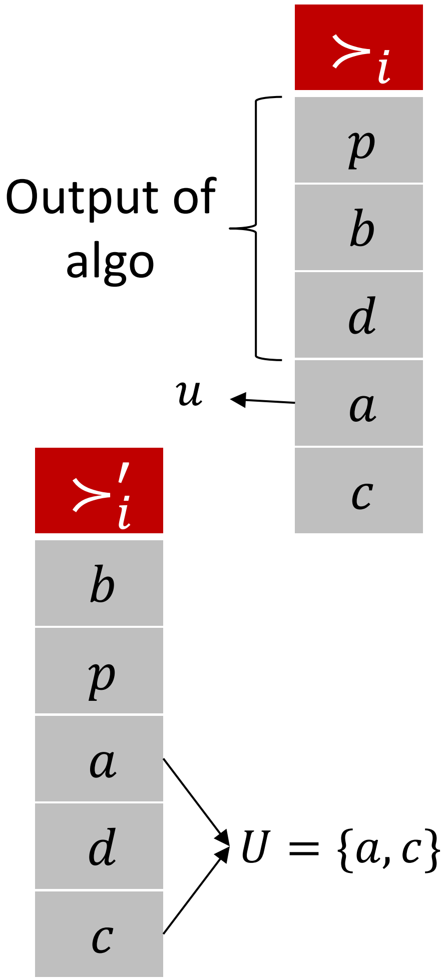

Proof of the Theorem

- $s(\succ_i, p) \geq s(\succ_i', p)$
  - 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

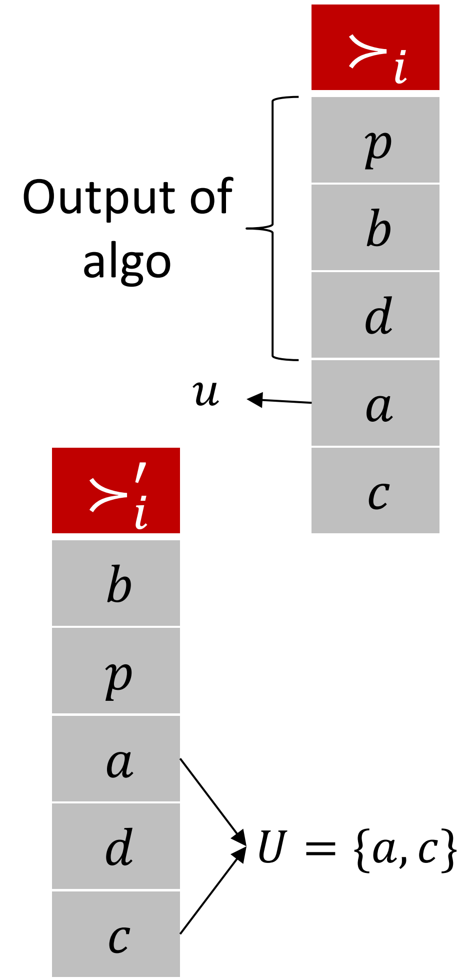

Hard-to-Manipulate Rules

- 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

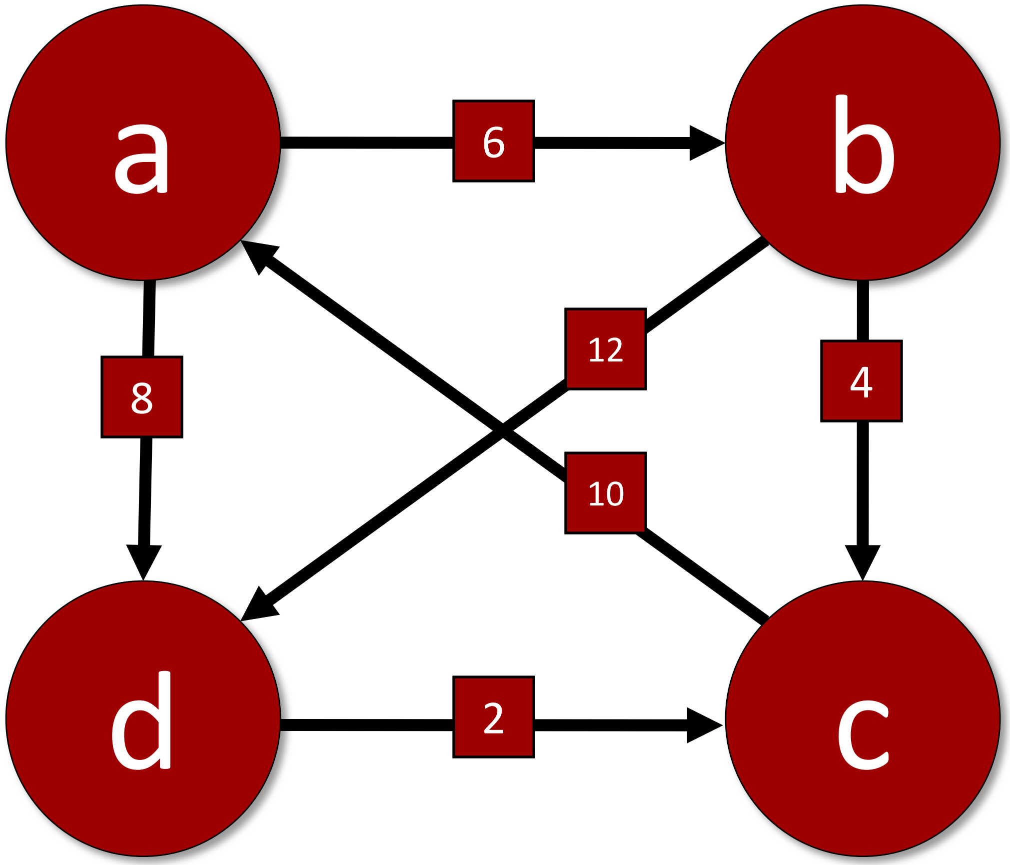

Example: Ranked Pairs

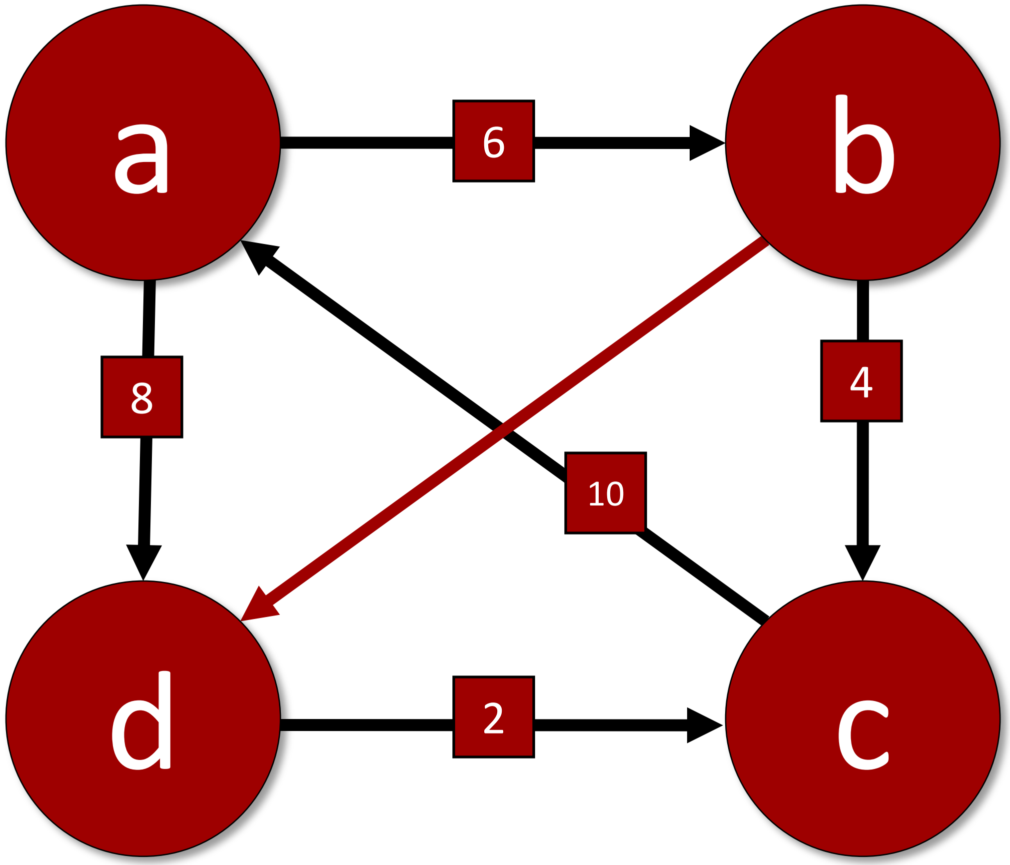

Example: Ranked Pairs

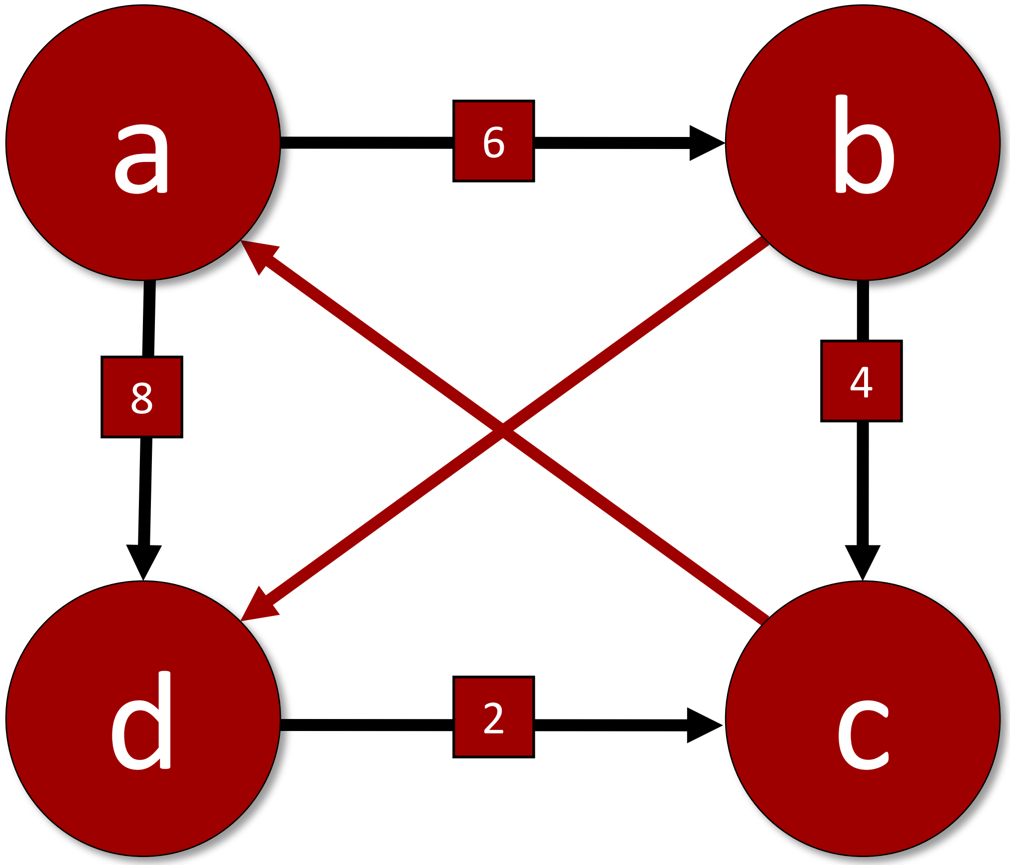

Example: Ranked Pairs

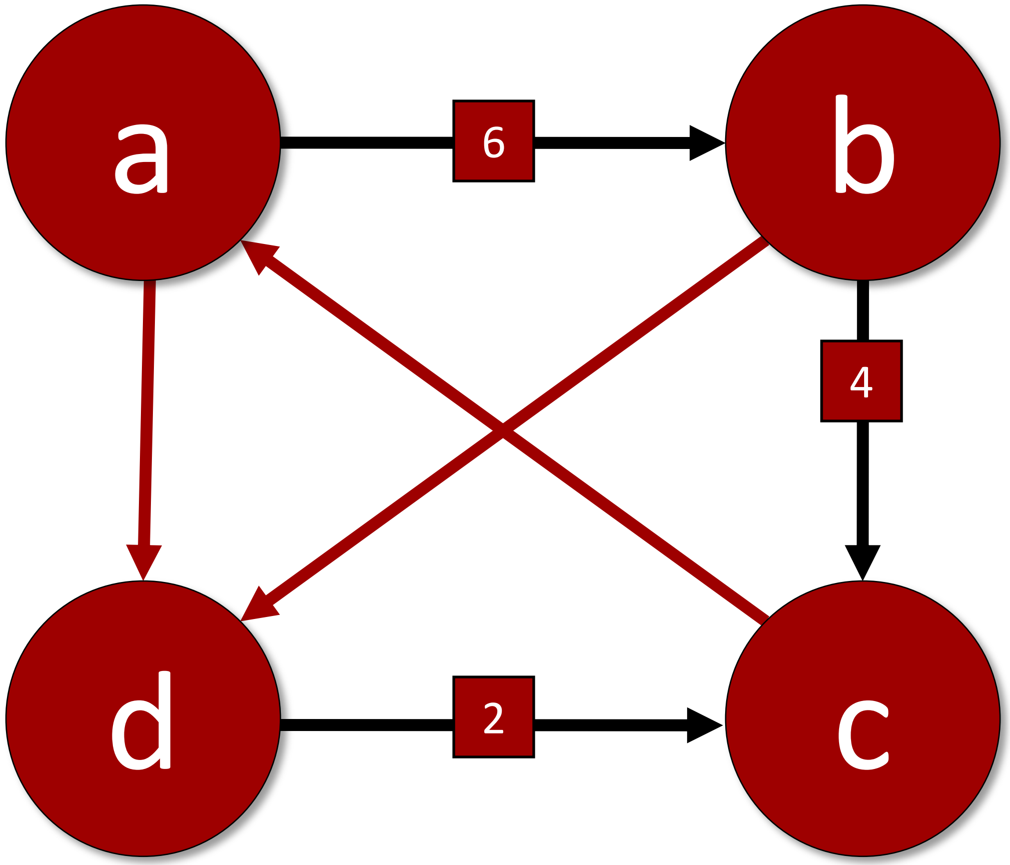

Example: Ranked Pairs

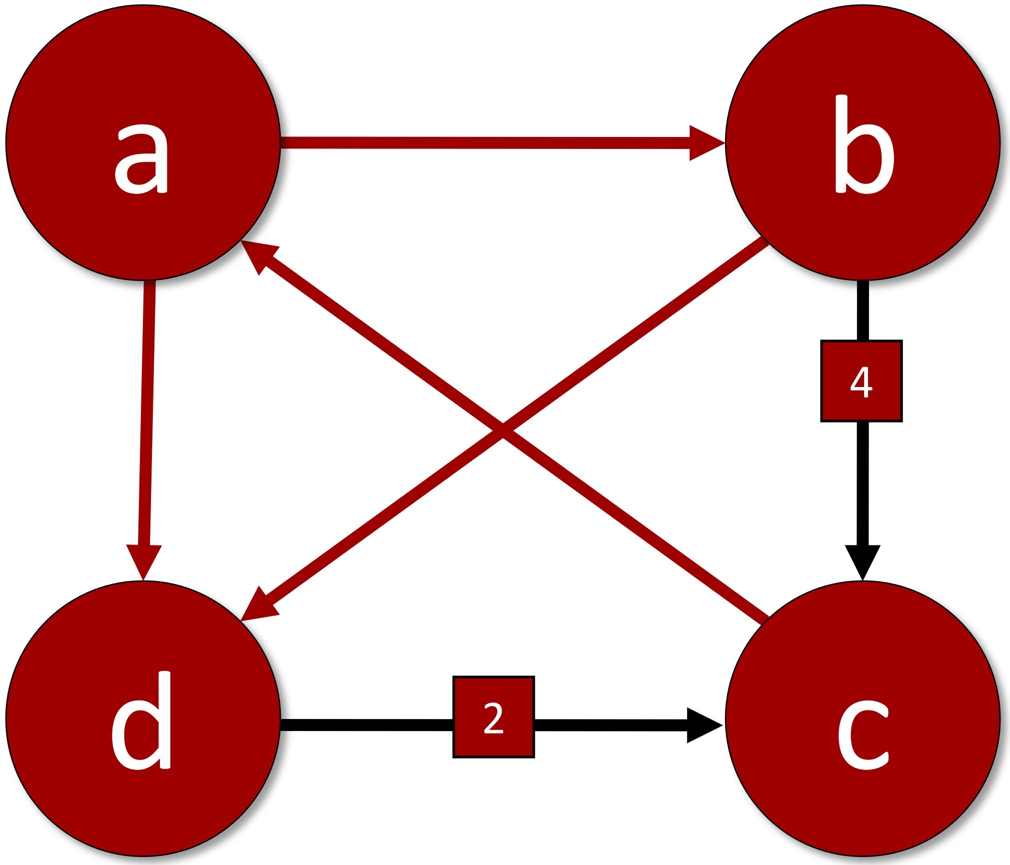

Example: Ranked Pairs

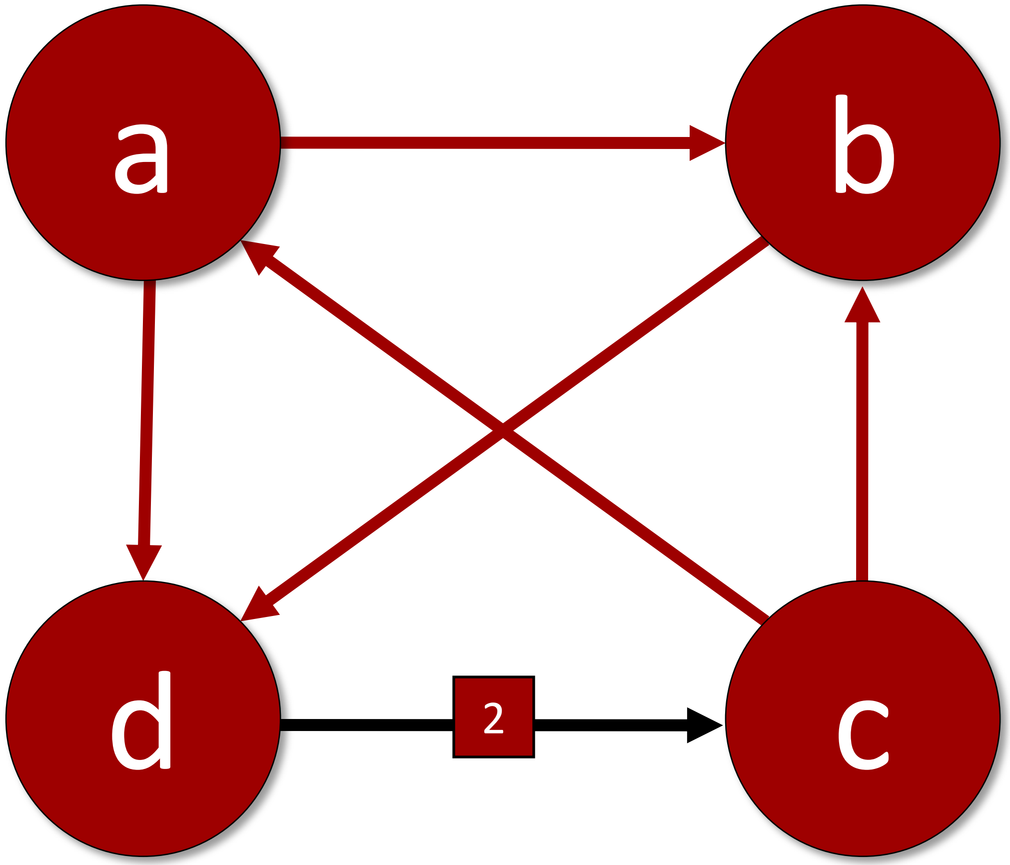

Example: Ranked Pairs

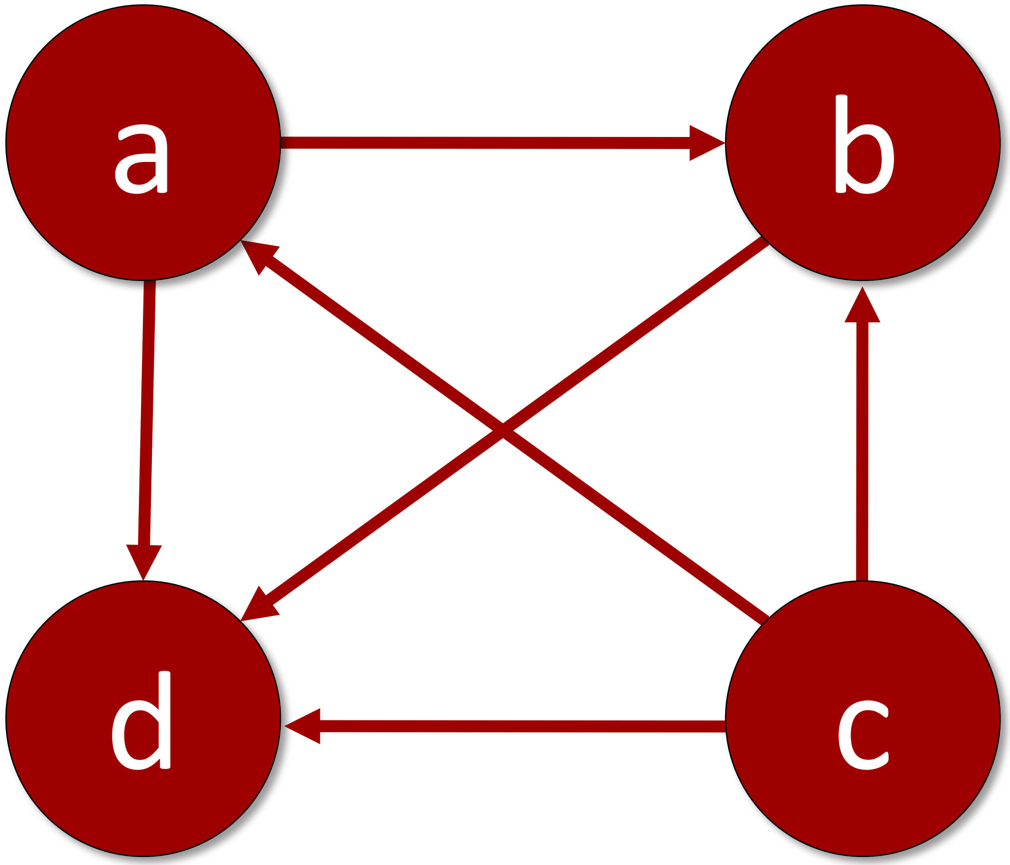

Randomized Voting Rules

- Input: preference profile
- Output: *distribution* over alternatives
  - To think about successful manipulations, we need numerical utilities
- u_i is consistent with $\succ_i$ if

$$a \succ_i b \Rightarrow u_i(a) \geq u_i(b)$$

- Strategyproofness:
  - For all i, $\vec{\succ}_{-i}$, $\succ_i$, $\succ_i'$, and u_i consistent with $\succ_i$

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

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 $f_1, \ldots, f_k$ is a rule given by some probability distribution $(\alpha_1, \ldots, \alpha_k)$ 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 $\mathrm{sc}(\vec{\succ}, a)$

- A randomized voting rule f is a c-approximation to sc if for every profile $\vec{\succ}$

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