Strategic Voting

(tutorial)

Full lecture notes available here

Reshef Meir, Technion, COST COMSOC summer school 2016

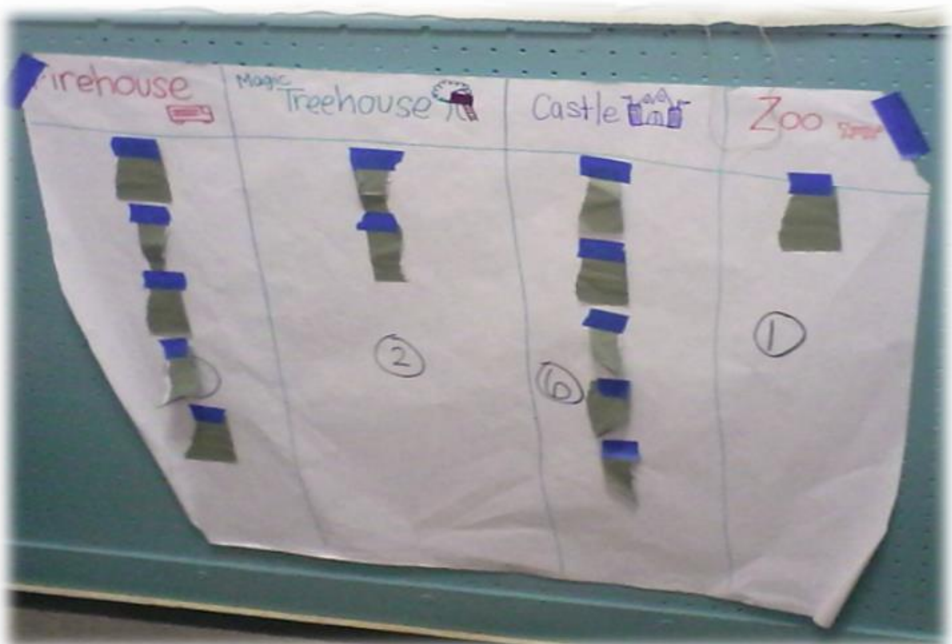

Least preferred

Most preferred

Manipulation

- Consider the following voting profile:

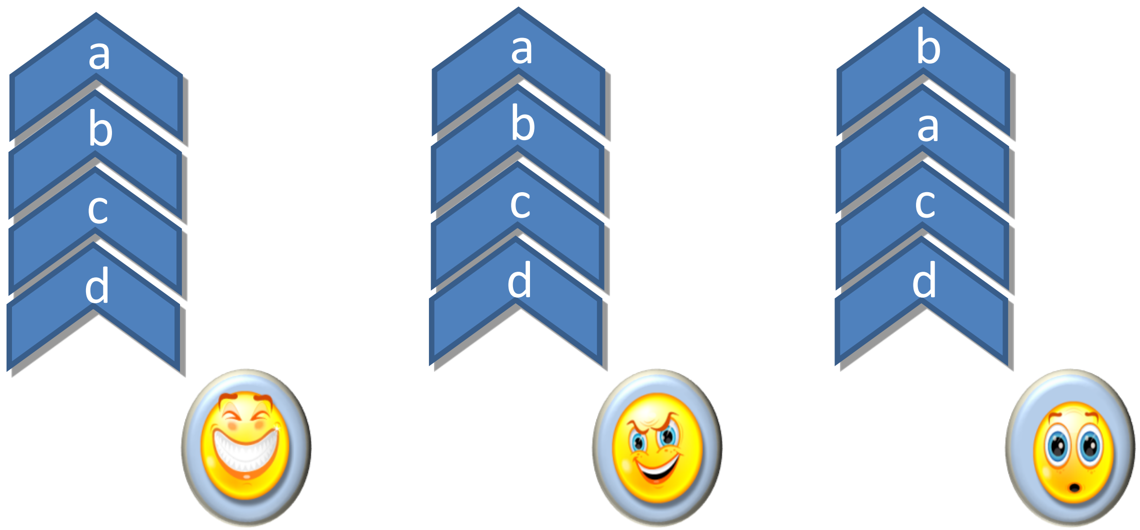

- If the Borda rule is used, then a will win
  - a has 8 points, while b only has 7

Manipulation

- Consider the following voting profile:

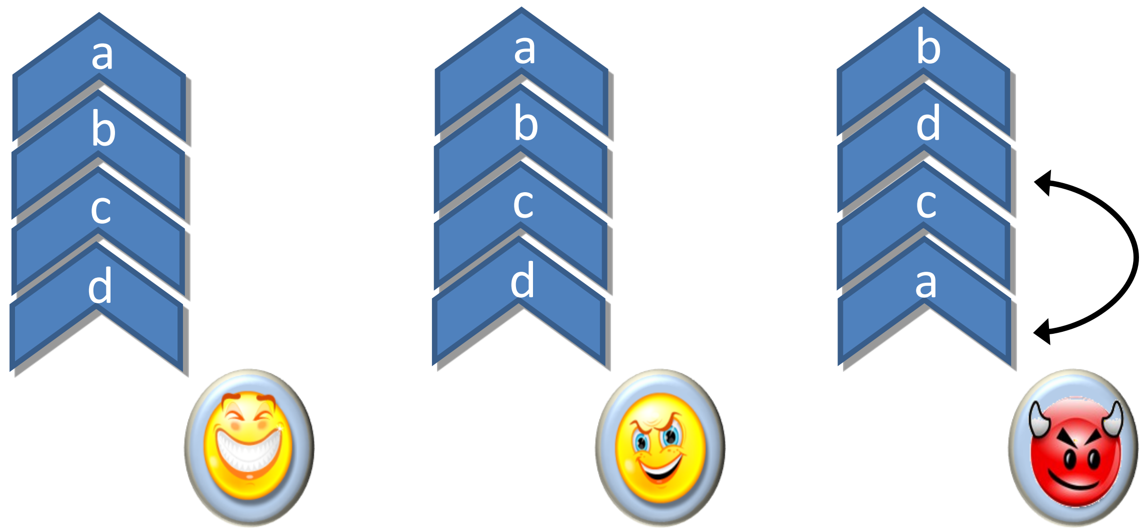

- But if voter 3 lies about his preferences…
  - Now a only has 6 points, and b wins!
- What would happen if we used Plurality?

Manipulation

- Neither Plurality nor Borda are immune to strategic voting
- We next see that under mild requirements, no voting rule is
  - The Gibbard-Satterthwaite theorem
- In the rest of the course we will consider the implications

Strategyproofness

<u>Definition</u>: *a voting rule f is* ***strategyproof (SP)****, if no (single) voter can ever benefit from lying about his preferences. Formally:*
$\forall \boldsymbol{R} \in \mathcal{L}(A)^n, \forall i \in N, \forall R_i' \in \mathcal{L}(A), f(\boldsymbol{R}_{-i}, R_i') \preccurlyeq_i f(\boldsymbol{R})$

<u>Claim</u>: If $|A| = 2$ (i.e. there are two candidates), then Majority is Strategyproof

– In this case **all** the standard voting rules are also SP

More axioms

<u>Definition</u>: *A voting rule f is **dictatorial** if there is an individual (the dictator) whose most preferred candidate is always chosen by f. formally:*

$$\exists i \in N, \forall \boldsymbol{R} \in \mathcal{L}(A)^n, f(\boldsymbol{R}) = top(R_i)$$

<u>Definition</u>: *A voting rule f is **onto** if it is possible for any of the candidates to win (given the right preference profile):*

$$\forall a \in A, \exists \boldsymbol{R} \in \mathcal{L}(A)^n, f(\boldsymbol{R}) = a$$

The Gibbard-Satterthwaite Theorem

contradiction!

(if |A|≥3)

$$\begin{cases} \forall \boldsymbol{R} \in \mathcal{L}(A)^n, \forall i \in N, \forall R_i' \in \mathcal{L}(A), f(\boldsymbol{R}_{-i}, R_i') \preccurlyeq_i f(\boldsymbol{R}) & \text{(SP)} \\ \neg \exists i \in N, \forall \boldsymbol{R} \in \mathcal{L}(A)^n, f(\boldsymbol{R}) = top(R_i) & \text{(no dictator)} \\ \forall a \in A, \exists \boldsymbol{R} \in \mathcal{L}(A)^n, f(\boldsymbol{R}) = a & \text{(onto)} \end{cases}$$

<u>The Gibbard-Satterthwaite theorem</u>*: If there are at least three candidates, any voting rule that is strategy-proof and onto is dictatorial.*

Proof outline for G-S theorem

- Every SP rule is (Maskin) Monotone
  - If $[f(\boldsymbol{R}) = a \text{ and } \forall i \in N \; \forall b \in A \; (b \prec_i a \Rightarrow b \prec'_i a)]$ *then* $f(\boldsymbol{R}') = a$
- Every onto + SP rule is Pareto
  - $!\exists b$ s.t. $\forall i \in N \quad b \succ_i f(\boldsymbol{R})$
- For $n = 2$, and any pair $a, b \in A$:
  - Either voter 1 can enforce a wins ("a-dictator")
  - Or voter 2 can enforce b wins ("b-dictator")

Consider a pair a, b

R_1 R_2

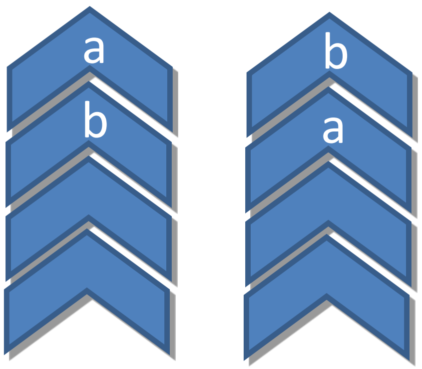

By Pareto, a or b wins

w.l.o.g. a wins

Consider a pair a, b

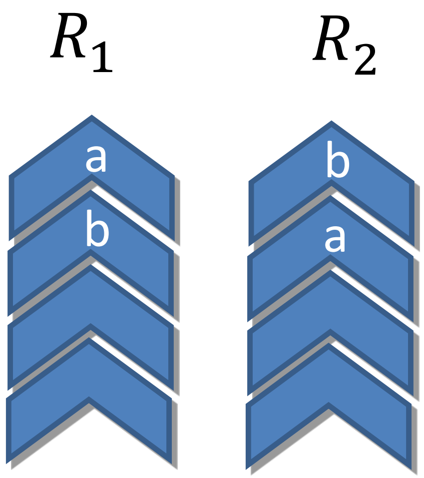

By Pareto, a or b wins
w.l.o.g. a wins

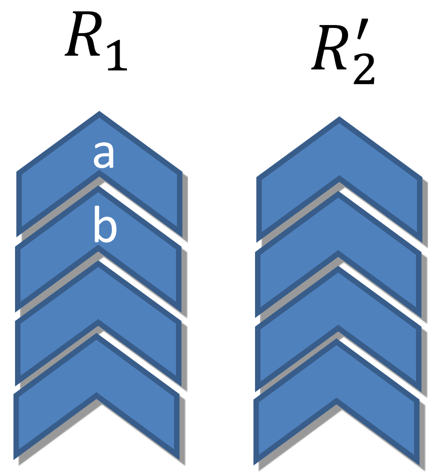

By SP for voter 2,
a still wins

Consider a pair a, b

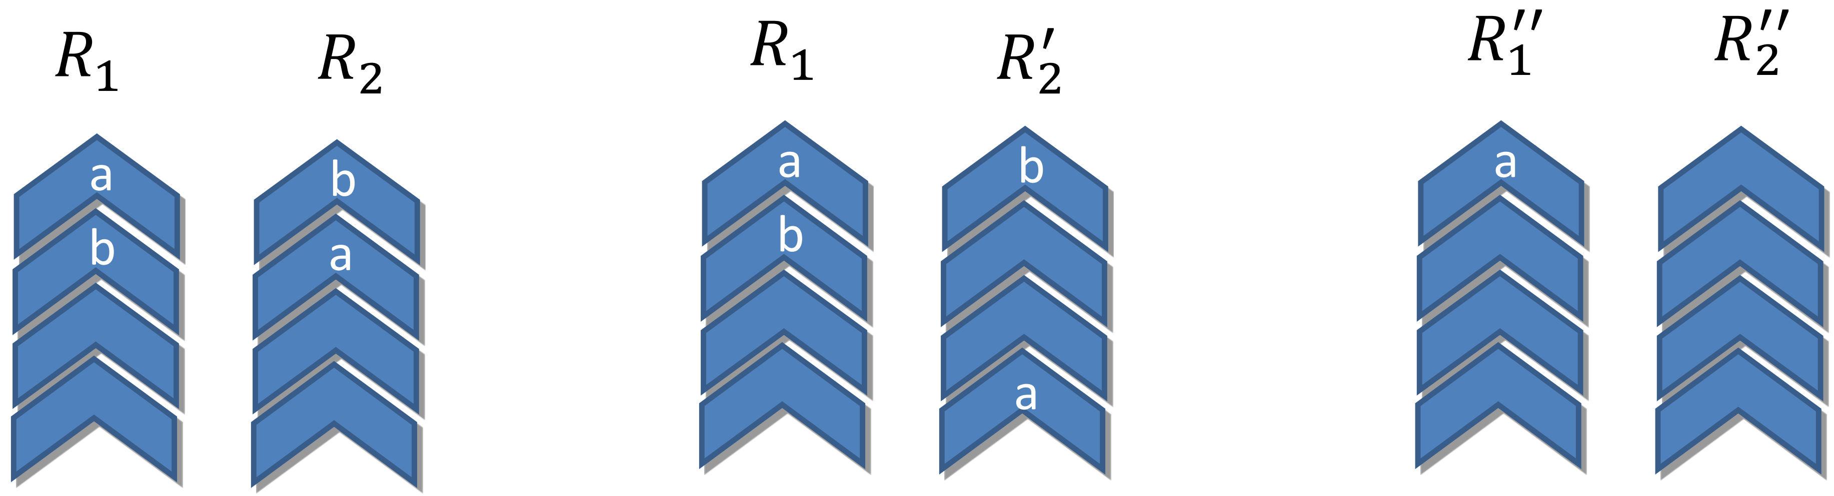

By Pareto, a or b wins
w.l.o.g. a wins

By SP for voter 2,
a still wins

By monotonicity,
a still wins

Thus voter 1 is an "a-dictator"

Proof outline for G-S theorem

- Every SP rule is Monotone
- Every onto + SP rule is Pareto
- For $n = 2$, and any pair $a, b \in A$:
  - Either voter 1 can enforce a wins ("a-dictator")
  - Or voter 2 can enforce b wins ("b-dictator")
- Conclude there is a dictator for all $c \in A$
- Extend to $n > 2$ by induction

For this and other simple proofs see [Svensson'99]

Course outline

The G-S theorem

More negative results

Achieving truthfulness
By additional assumptions
("workarounds" for G-S)

Relax truthfulness:
- Rational voting and equilibrium analysis
- Iterative voting and convergence

Relax rationality:
Heuristic voting

See lecture notes for more details and full references

More negative results:

- Manipulations occur often
- Randomization does not help (much)
- Strategyproofness entails dictatorship in other domains

Can manipulations occur often?

- Intuitively, a single voter is usually powerless
  - In particular, cannot manipulate
- How can we measure this formally?
- [Friedgut et al.'08] assume unbiased culture
  - (uniform distribution on all profiles)
- Define $M_i(f)$ as the probability that i has a manipulation R'_i, when the profile $\boldsymbol{R}$ is selected uniformly at random.
  - the "manipulation power" of voter i

“quantitative” G-S theorem

- Rephrasing the G-S theorem:

Either f is dictatorial, or a duple, or $\sum_i M_i(f) > 0$

Only 2 possible outcomes

Def.: f is ε-bad if there is some dictatorship/duple g s.t. $\Pr\left(f(\boldsymbol{R}) = g(\boldsymbol{R})\right) \geq 1 - \varepsilon$

Theorem [Friedgut et al.’08, Mossel and Rácz’12]:

Either f is ε-bad, or $\sum_i M_i(f) > \mathrm{poly}(\frac{1}{n}, \frac{1}{m}, \varepsilon)$

“manipulations occur often”

Randomized voting rules

- Suppose we allow our voting rule to use randomization
- We have more ways to define an SP rule:
  - A random fixed outcome
  - A random dictator
  - A random duple (select a pair of candidates at random, and use majority)
- Note that we have to define cardinal utilities for voters

Randomized voting rules

Theorem [Gibbard'77]: Any strategyproof randomized rule, is a lottery over dictatorships* and duples.

More negative results:

- Manipulations occur often
- Randomization does not help (much)
- SP means dictatorship in
  - **facility location**
  - Classification
  - Judgement aggregation

Strategyproofness on graphs

- Suppose agents report location on a graph
- Want a facility to placed as close as possible

Theorem [Schummer and Vohra'04]: if the graph has **cycles**, any SP+onto rule is a dictatorship.

- This assumes graph is continuous

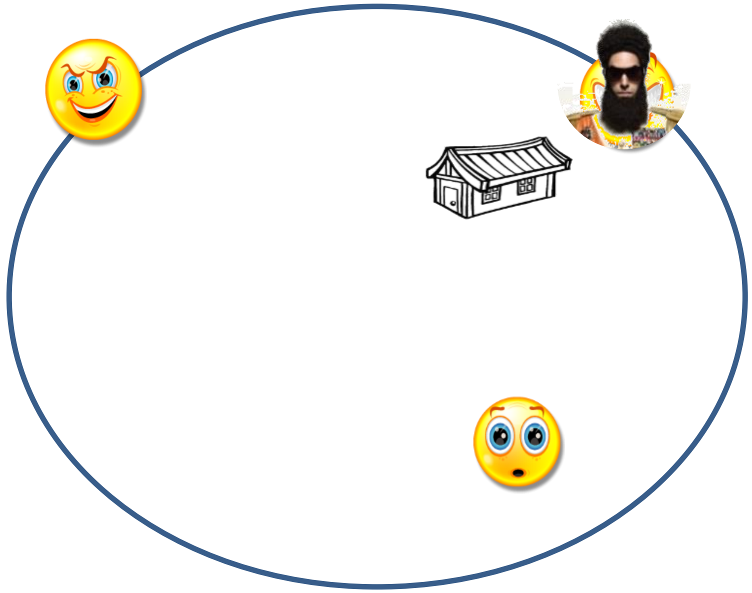

Strategyproofness on graphs

- What about discrete graphs?
- Agents and facility must be placed on vertices

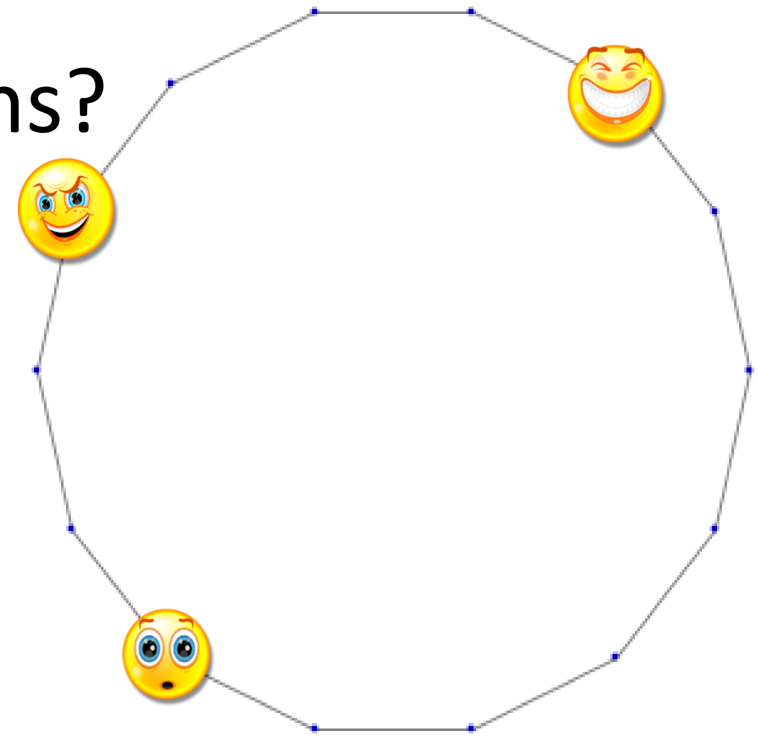

[Dokow et al.'12]: Still "almost dictatorial" for *large cycles*.

Not true for small cycles (at most 12 nodes)

Course outline

The G-S theorem

More negative results

Achieving truthfulness
By additional assumptions
("workarounds" for G-S)

Relax truthfulness:
- Rational voting and equilibrium analysis
- Iterative voting and convergence

Relax rationality:
Heuristic voting

Course outline

The G-S theorem

More negative results

5 surprising workarounds to the G-S theorem!!!

Relax truthfulness:

- Rational voting and equilibrium analysis
- Iterative voting and convergence

Relax rationality:

Heuristic voting

Achieving truthfulness under additional assumptions

1. **Domain restriction (e.g. single-peak)**
2. Complexity barriers
3. Approximation
4. Differential privacy
5. Payments

Were covered by Edith Elkind

Single-Peaked Preferences

- Definition: a preference profile is single-peaked (SP) wrt an ordering < of candidates (axis) if for each voter v:
  - if top(v) < D < E, v prefers D to E
  - if A < B < top(v), v prefers B to A
- Example:
  - voter 1: C > B > D > E > F > A
  - voter 2: A > B > C > D > E > F
  - voter 3: E > F > D > C > B > A

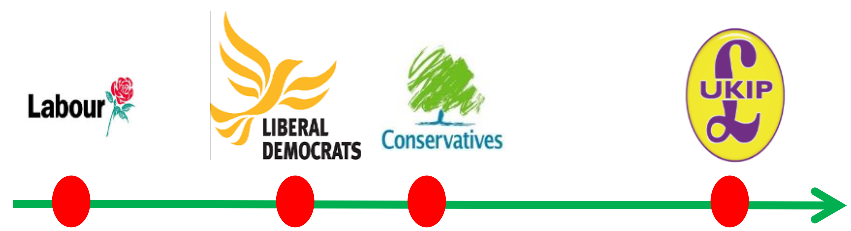

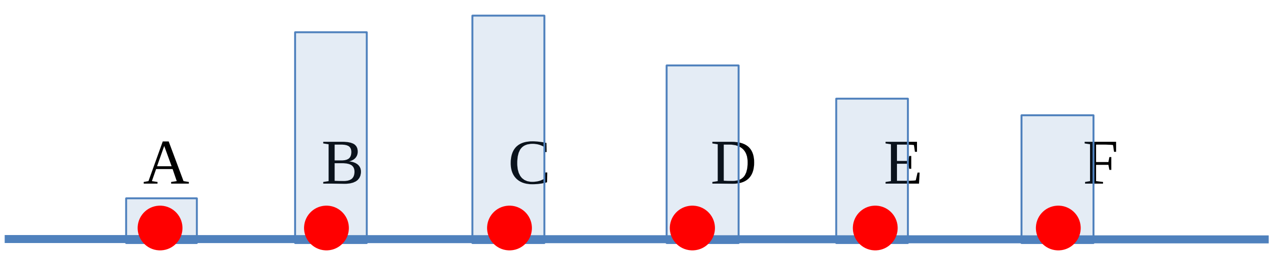

SP Preferences: Circumventing Gibbard-Satterthwaite

- Suppose we have $n = 2k+1$ voters
- Median voter rule:
  - consider an election that is single-peaked wrt $\boldsymbol{R}$
  - ask each voter v to vote for one candidate
    - let $C(v)$ denote the vote of voter v
  - order voters by $C(v)$, breaking ties arbitrarily
  - output $C^* = C(v_{k+1})$

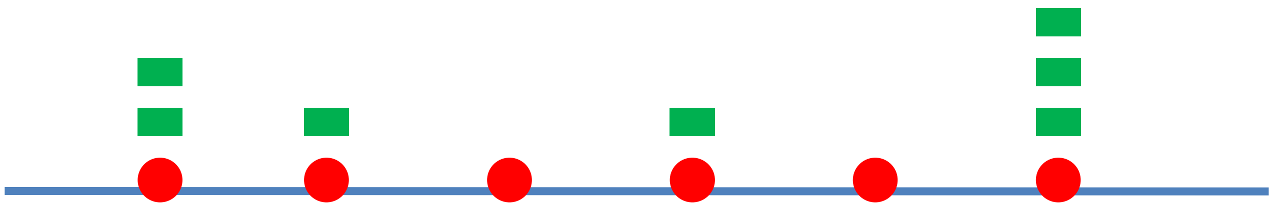

SP Preferences: Median Is Truthful

- Theorem: under the median voter rule, it is a dominant strategy to vote for one's top choice

- Still true for single-peaked preferences on a tree

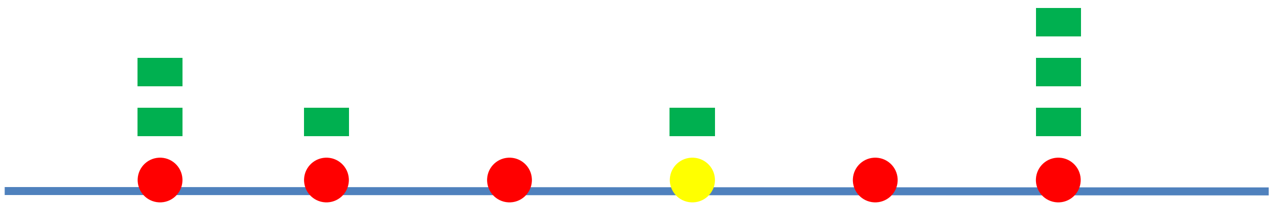

Achieving truthfulness

1. Domain restriction (e.g. single-peak)
2. **Complexity barriers**
3. Approximation
4. Differential privacy
5. Payments

MANIPULATION$_f$

Fix a voting rule f

Given:

- a set of candidates A
- a group of voters N
- a specific candidate p in A
- a manipulator $\mathrm{i} \in N$
- and a preference profile $\boldsymbol{R}_{-i}$ of all voters except i

- Answer whether the manipulator i can vote such that p will be chosen by f

A greedy algorithm for the (greedy) manipulator

- Rank p first
- While there are more candidates:
  - If there exists a “safe” candidate, rank that candidate in the next spot.
  - otherwise - declare that the desired preference does not exist.

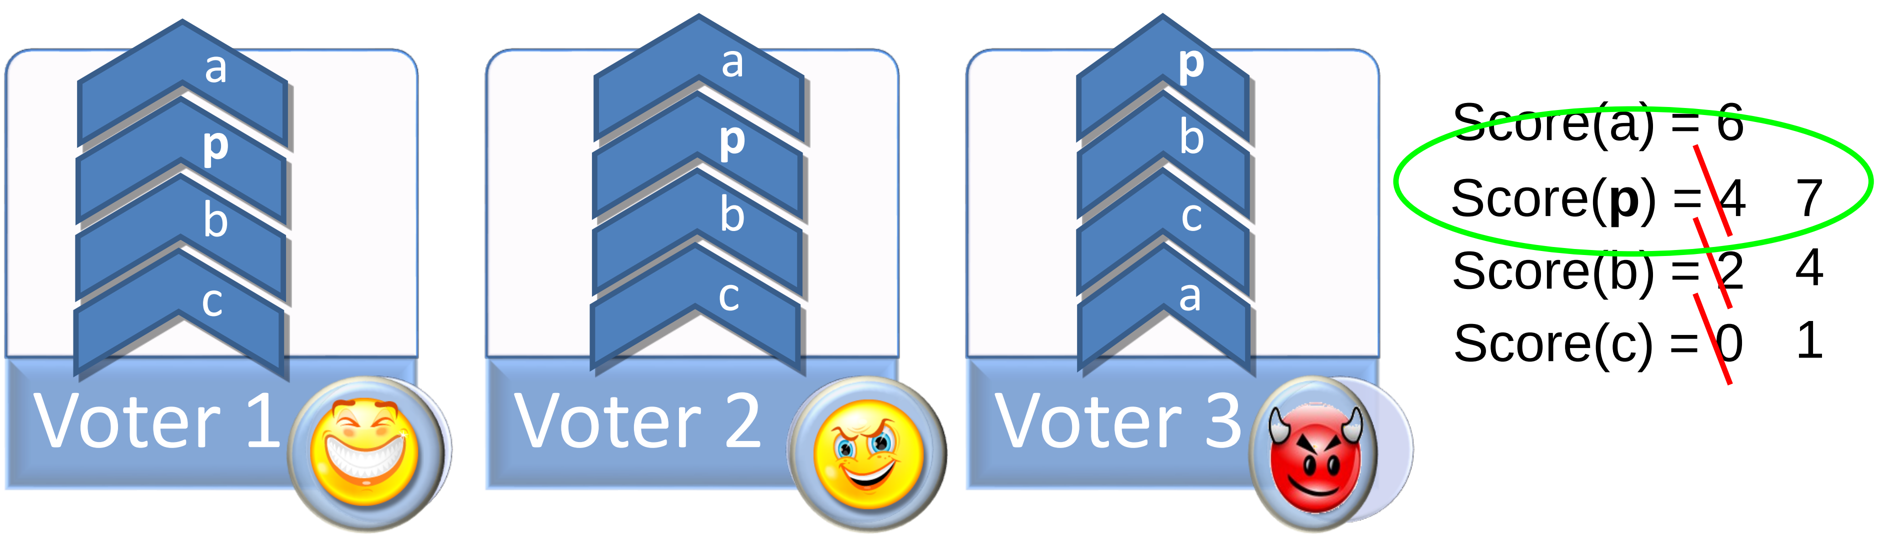

When will it work?

<u>Proposition:</u> The greedy algorithm works for every scoring-based rule:

- PSRs
- Copeland
- Maximin

Is there a similar algorithm for other rules?

What about other algorithms?

Theorem [BTT '89]: There is a voting rule f, for which MANIPULATION$_f$ is NP-hard

(believed that no efficient algorithm exists)

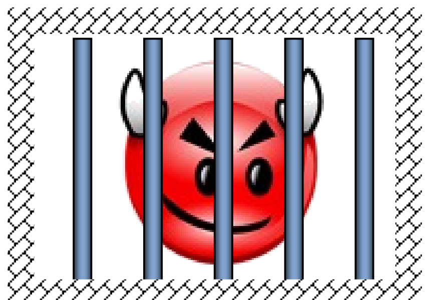

- Original proof used a variant of Copeland
- Also hard: Single Transferable Vote (STV)

Hardness of manipulation

- Proving that MANIPULATION_f is hard is a **positive result** – it means voters are likely to be truthful
- An argument in favor of some rules like STV

- But:
  - Only proves the worst-case
  - Very sensitive to small variations

Coalitional manipulation

- Suppose we use Borda
  - No single manipulator can gain
  - But if first three voters join forces…

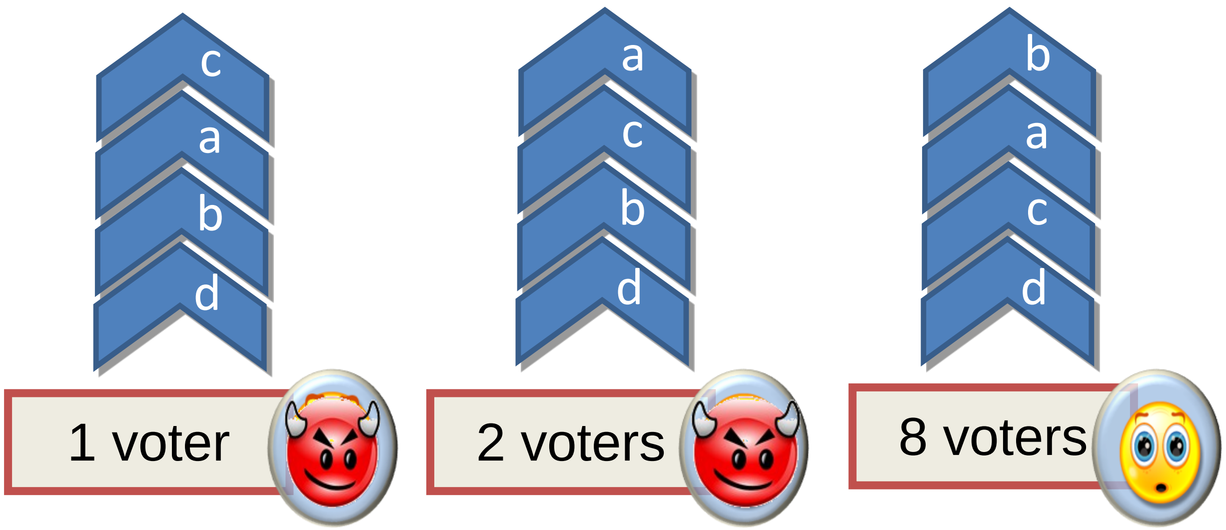

Score(a) = 2+6+16 = 24

Score(b) = 1+4+24 = 29

Score(c) = 3+0+8 = 11

Score(d) = 0+2+0 = 2

Coalitional manipulation

- Suppose we use Borda
  - No single manipulator can gain
  - But if first three voters join forces…

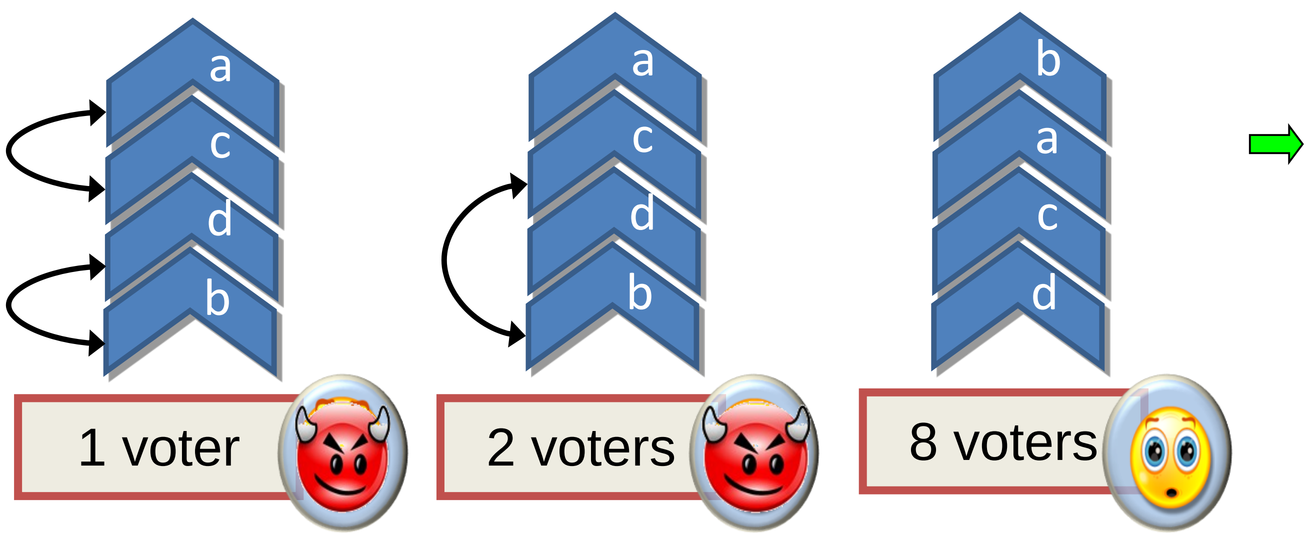

Score(a) = 3+6+16 = 25

Score(b) = 0+0+24 = 24

Score(c) = 2+3+8 = 13

Score(d) = 1+6+0 = 7

There is no efficient algorithm* for coalitional manipulation of Borda, even for 2 manipulators [DKNW11,BNW11]

Achieving truthfulness

1. Domain restriction (e.g. single-peak)
2. Complexity barriers
3. **Approximation**
4. Differential privacy
5. Payments

Approximation

- Suppose we allow randomization
- We saw that by [Gibbard’77] this only extends the class of SP rules to mixtures of:
  - Dictators (and monotone unilateral rules)
  - Duples
- Perhaps these rules are “good enough”?
- The winner is closed in expectation to the winner of another desired rule

Approximation

- Consider any scoring-based rule g
- A randomized rule f is a γ-approximation of g if for any profile $\boldsymbol{R}$,

$$E\left[score_g\left(f(\boldsymbol{R})\right)\right] \geq \gamma \cdot score_g\left(g(\boldsymbol{R})\right)$$

(f selects winners that have high g-score in expectation compared to the true winner of g)

<u>Theorem</u> [Procaccia'10]: For any PSR g there is a randomized SP rule f_g that is a $\Omega(\frac{1}{\sqrt{m}})$-approximation

- What rule approximates Plurality?

Achieving truthfulness

1. Domain restriction (e.g. single-peak)

2. Complexity barriers

3. Approximation

4. **Differential privacy**

5. Payments

Differentially private voting rules

- The main idea:
  - Take any voting rule f
  - **Add noise** to the voting profile (corrupt some votes randomly)
  - This induces a new randomized voting rule f'

  - f‘ is “almost” strategyproof
  - f‘ is an approximation of f

Differentially private voting rules

Theorem [Birrel&Pass'11]: For any deterministic voting rule f, any $\varepsilon > 0$ and any $\delta > \frac{m^2}{\varepsilon}$, there is a randomized voting rule f' s.t.

– f' is ε-strategyproof (an agent can gain at most ε)

– f' is a δ-approximation of f (we can always get $f(\boldsymbol{R})$ by modifying at most δ votes in f')

– **Not equivalent** to the definition of approximation by [Procaccia'10]

Achieving truthfulness

1. Domain restriction (e.g. single-peak)

2. Complexity barriers

3. approximation

4. Differential privacy

5. **Payments**

Voting with payments

- Suppose voters have **cardinal values**
- This means voter i is willing to pay v_{ia} if a is selected
- We can turn the voting process to an **auction**:
  - Each voter will report her valuations
  - The alternative a with the highest bids $\sum_{i \in N} v_{ia}$ wins
  - We charge payment from agents
  - How much?

Direct payment mechanism

- Initial attempt: if w wins, charge each voter v_{iw}
- This is not truthful
  - E.g. if w wins anyway, i can deviate by reporting $v'_{iw} = 0$

- Recall the second-price auction:
  - the payment of i should not depend on i's own bids

VCG voting

- Define $p_i = max_z(\sum_{k \neq i} v_{kz}) - \sum_{k \neq i} v_{kw}$
- Each agent pays p_i, gets utility $v_{iw} - p_i$

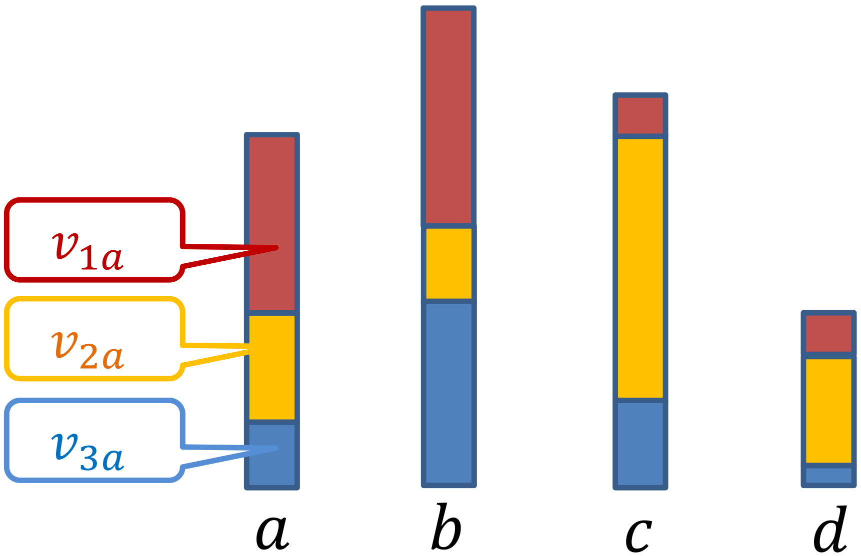

VCG voting

- Define $p_i = \boxed{max_z(\sum_{k \neq i} v_{kz})} - \sum_{k \neq i} v_{kw}$
- Each agent pays p_i, gets utility $v_{iw} - p_i$

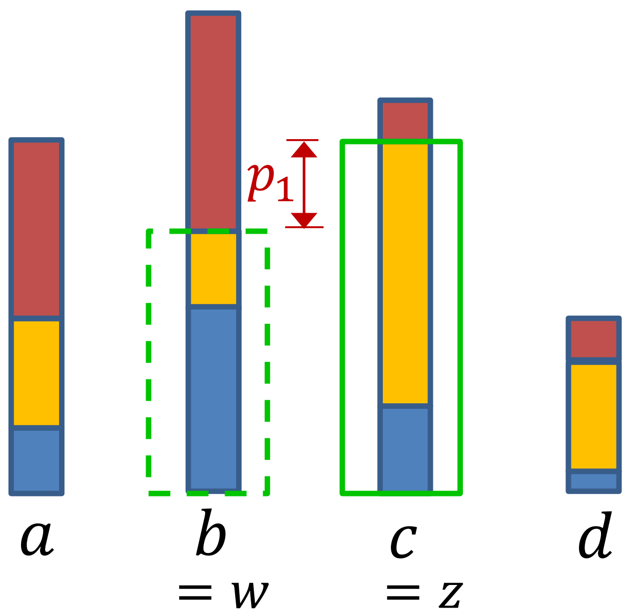

VCG voting

- Define $p_i = max_z(\sum_{k \neq i} v_{kz}) - \sum_{k \neq i} v_{kw}$
- Each agent pays p_i, gets utility $v_{iw} - p_i = u_i$

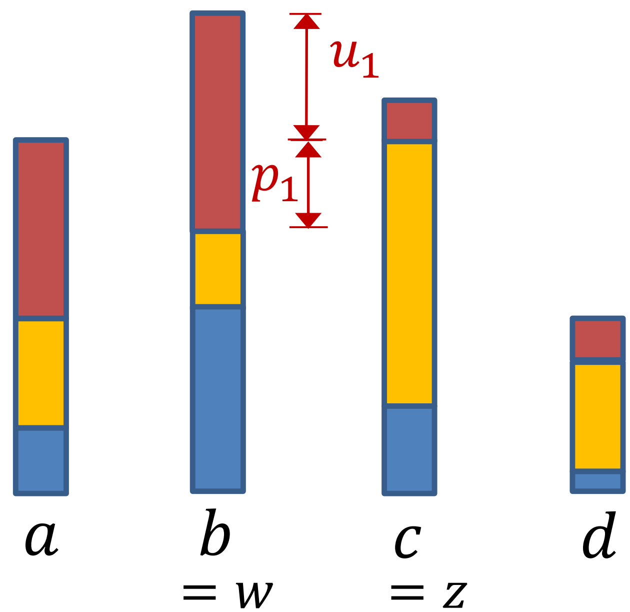

Theorem [Clarke’71]: VCG voting is strategyproof (for details see e.g. [Nissan’07, Section 9.3]).

VCG voting

- Define $p_i = max_z(\sum_{k \neq i} v_{kz}) - \sum_{k \neq i} v_{kw}$
- Each agent pays p_i, gets utility $v_{iw} - p_i = u_i$

Possible manipulations:

$v'_{1a} = 0 \ \forall a$

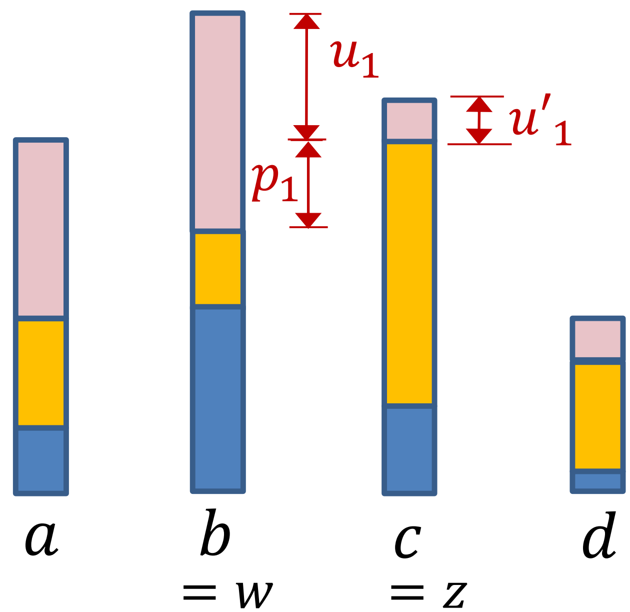

<u>Theorem</u>: VCG voting is strategyproof

VCG voting

- Define $p_i = max_z(\sum_{k \neq i} v_{kz}) - \sum_{k \neq i} v_{kw}$
- Each agent pays p_i, gets utility $v_{iw} - p_i = u_i$

Possible manipulations:

$v'_{1a} = 0 \ \forall a$

$v'_{1w} < v_{1w}$

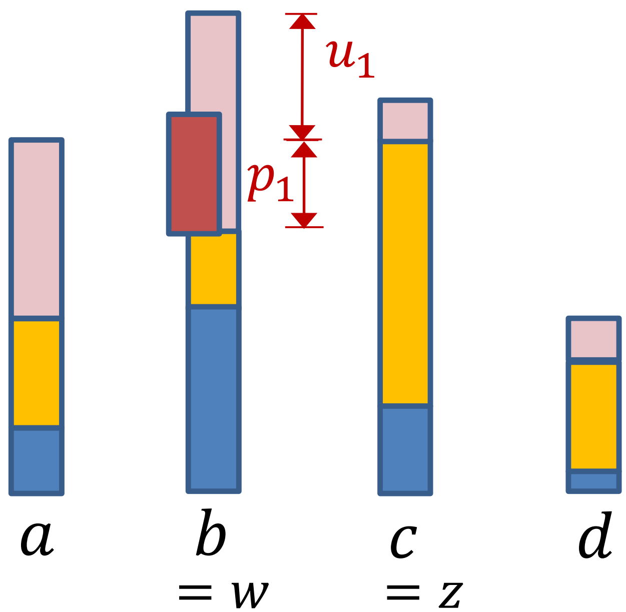

<u>Theorem</u>: VCG voting is strategyproof

VCG voting

- Define $p_i = max_z(\sum_{k \neq i} v_{kz}) - \sum_{k \neq i} v_{kw}$
- Each agent pays p_i, gets utility $v_{iw} - p_i = u_i$

Possible manipulations:

$v'_{1a} = 0 \;\; \forall a$

$v'_{1w} < v_{1w}$

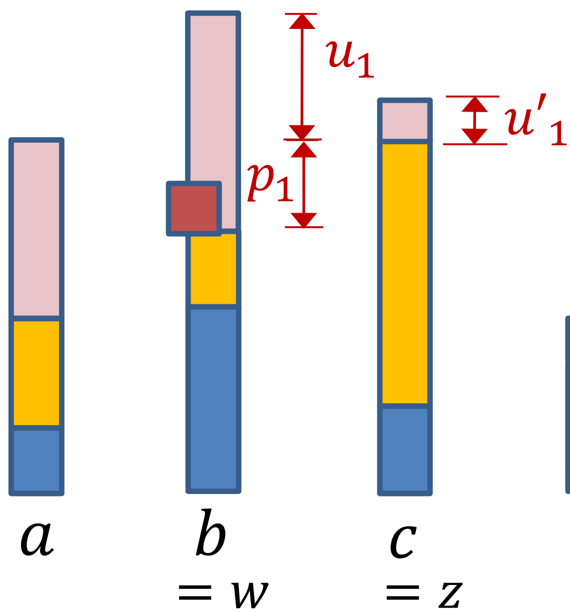

<u>Theorem</u>: VCG voting is strategyproof

VCG voting

- Define $p_i = max_z(\sum_{k \neq i} v_{kz}) - \sum_{k \neq i} v_{kw}$
- Each agent pays p_i, gets utility $v_{iw} - p_i = u_i$

Possible manipulations:

$v'_{1a} = 0 \;\; \forall a$

$v'_{1w} < v_{1w}$

$v'_{1y} > v_{1y}$

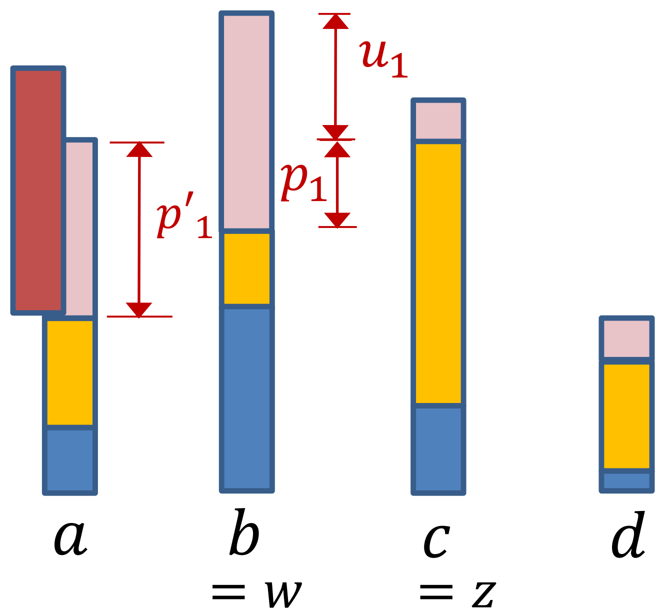

$\square$

<u>Theorem</u>: VCG voting is strategyproof

Course outline

The G-S theorem

More negative results

Achieving truthfulness
By additional assumptions
("workarounds" for G-S)

Relax truthfulness:
- Rational voting and equilibrium analysis
- Iterative voting and convergence

Relax rationality:
Heuristic voting

Voting as a game

- Instead of assuming truthfulness, we assume **rationality**
  - Voters vote the way that best suits their interests
  - Who wins when voters are rational?

- Every voting rule f defines a *game form*
- Together with a preference profile $\boldsymbol{R}$ we get a game $(f, \boldsymbol{R})$ with ordinal utilities
- The game-theory approach: analyze the **equilibria** of this game to predict the outcome

Voting as a normal-form game

$W_2=4$   $W_1=3$	a	b	c
a			
b			
c			

Initial score:	7	9	3
			

Voting as a normal-form game

$W_2=4$   $W_1=3$	a	b	c
a	(**14**,9,3)		
b	(11,**12**,3)		
c			

Initial score:	7	9	3
			

Voting as a normal-form game

W_2=4   W_1=3	a	b	c
a	(14,9,3)	(10,13,3)	(10,9,7)
b	(11,12,3)	(7,16,3)	(7,12,7)
c	(11,9,6)	(7,13,6)	(7,9,10)

Initial score:	7	9	3
			

Voting as a normal-form game

$W_2=4$ / $W_1=3$	a	b	c
a	(14,9,3)	(10,13,3)	(10,9,7)
b	(11,12,3)	(7,16,3)	(7,12,7)
c	(11,9,6)	(7,13,6)	(7,9,10)

Voters preferences:

- a > b > c
- c > a > b

Nash Equilibria

Voting equilibrium

1. Nash equilibrium
2. Strong equilibrium

Implementation theory

3. Truth-bias
4. Equilibrium under uncertainty

Nash equilibrium

- Let $NE_f(\boldsymbol{R}) \subseteq \mathrm{A}$ be all candidates that win in *some* Nash Equilibrium of the game $(f, \boldsymbol{R})$.
- Consider Plurality
  - Almost any state is a Nash
  - Thus $NE_f(\boldsymbol{R}) = \mathrm{A}$ almost always
  - Not informative at all!

- This seems to be true in all voting rules we have seen

NE Implementation

- Let $F: \mathcal{R}^n \rightarrow 2^A$ be some function that maps profiles to a winning subset
  - Examples: "All Plurality winners"; "All candidates"; "All Condorcet winners"; "All non-Condorcet losers"
- A voting rule f ***implements*** F in NE if

$$NE_f(\boldsymbol{R}) = F(\boldsymbol{R}) \text{ for all } \boldsymbol{R}$$

- What does Plurality implement in NE?

NE Implementation

- Question 1: Can a voting rule f implement f itself in NE?

- Question 2: Can a voting rule f be implemented in NE by *some* mechanism M?

NE Implementation

- Question 1: Can a voting rule f implement f itself in NE?

- Question 2: Can a voting rule f be implemented in NE by *some* mechanism M?
  - No, except for dictator/duple [Maskin'85]
  - Possible for some non-resolute rules (SCCs)
    - E.g. 1 "All outcomes"
    - E.g. 2 "Pareto outcomes"

NE Implementation

- Question 1: Can a voting rule f implement f itself in NE? No (except trivial rules)

- Question 2: Can a voting rule f be implemented in NE by some mechanism M?
  – No, except for dictator/duple [Maskin'85]

NE Implementation

- Question 1: Can a voting rule f implement f itself in NE? No (except trivial rules)

- Question 1*: Can a voting rule f implement f itself in Dominant Strategies?
  - No (Except trivial rules)
  - Due to the G-S theorem

Strong Equilibrium Implementation

- A strong equilibrium (SE): no **coalition** has a reply where all members gain
- Let $F: \mathcal{R}^n \rightarrow 2^A$ be some function that maps profiles to a winning subset
- A voting rule f ***implements*** F in SE if

$$SE_f(\boldsymbol{R}) = F(\boldsymbol{R}) \text{ for all } \boldsymbol{R}$$

Theorem [Sertel & Sanver'04]:
Plurality implements Condorcet in SE.

(proof is almost trivial!)

Other notions of implementation

- Protective equilibria [Barbera&Dutta'86]
  - Veto implements itself
- Demanding equilibria [Merlin&Naeve'01]
  - Plurality implements itself
- Scoring rules [Falik, M., Tennenholz'12]
  - Plurality implements Maximin

Voting equilibrium

1. Nash equilibrium

2. Strong equilibrium

3. **Truth-bias**

4. Equilibrium under uncertainty

Truth bias

- Suppose that voters have some very weak preference to be truthful
  - Will be strategic if it helps them even slightly
  - If they have no effect at all, will remain truthful
- This assumption "kills" many weird equilibria like "all vote for candidate x"
- Let $TNE_f(\boldsymbol{R}) \subseteq \mathrm{A}$ be all winners in some NE of $(f, \boldsymbol{R})$ under truth-bias.

Truth biased equilibrium

- $TNE_f(\boldsymbol{R})$ may be empty.

Example: Plurality with 2 voters:

- $R_1: c \succ a \succ b \succ d$
- $R_2: d \succ b \succ a \succ c$

1 \ 2	a	b	c	d
a	a	a	a	a
b	a	b	b	b
c	a	b	c	c
d	a	b	c	d

Truth biased equilibrium

- $TNE_f(\boldsymbol{R})$ may be empty.

A characterization of TNEs in Plurality voting games:

- Easy for the truthful winner
- NP-Hard otherwise
- [Obraztsova et al.'13]

Example: Plurality with 2 voters:

- $R_1: c \succ a \succ b \succ d$
- $R_2: d \succ b \succ a \succ c$

1 \ 2	a	b	c	d
a	a	a	a	a
b	a	b	b	b
c	a	b	c	c
d	a	b	c	d

TNEs on average [Thompson et al.'13]

- Some TNE almost always exist
- Truthful TNEs are common

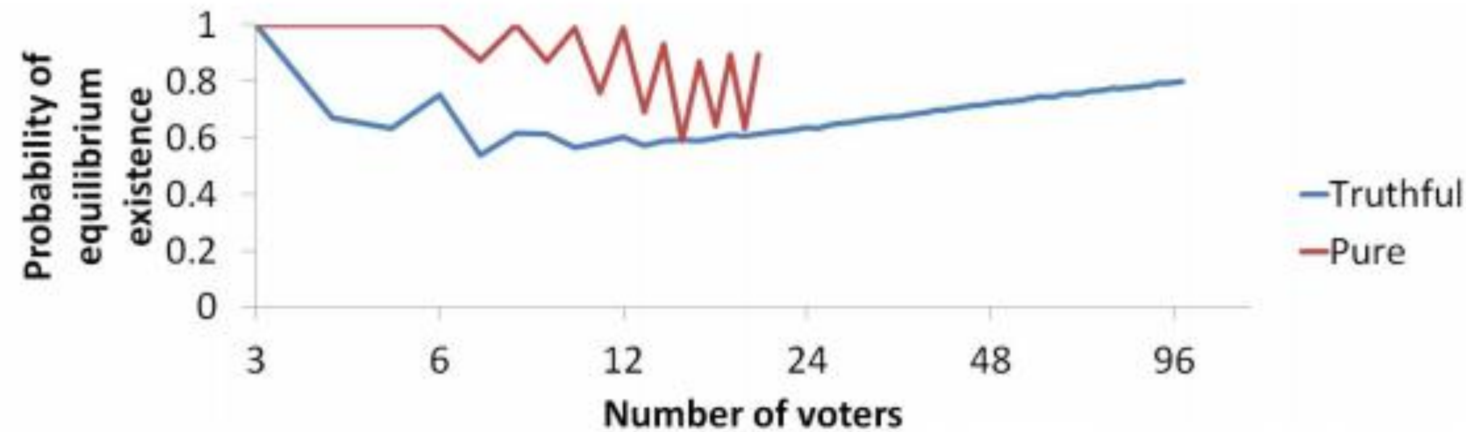

- On the other hand, instead of millions of NEs, there are typically just a few TNEs

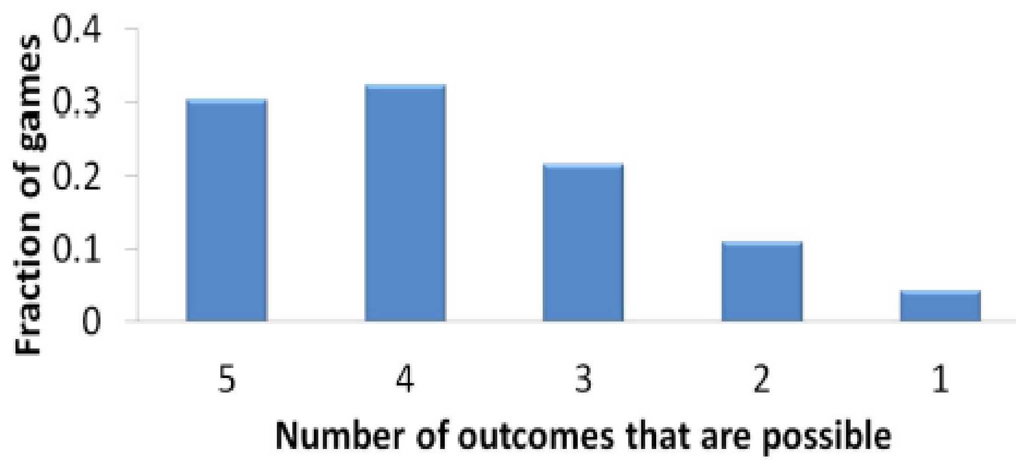

TNEs on average

- NEs are often really bad (e.g. when all vote to a bad candidate)
- How about TNEs?

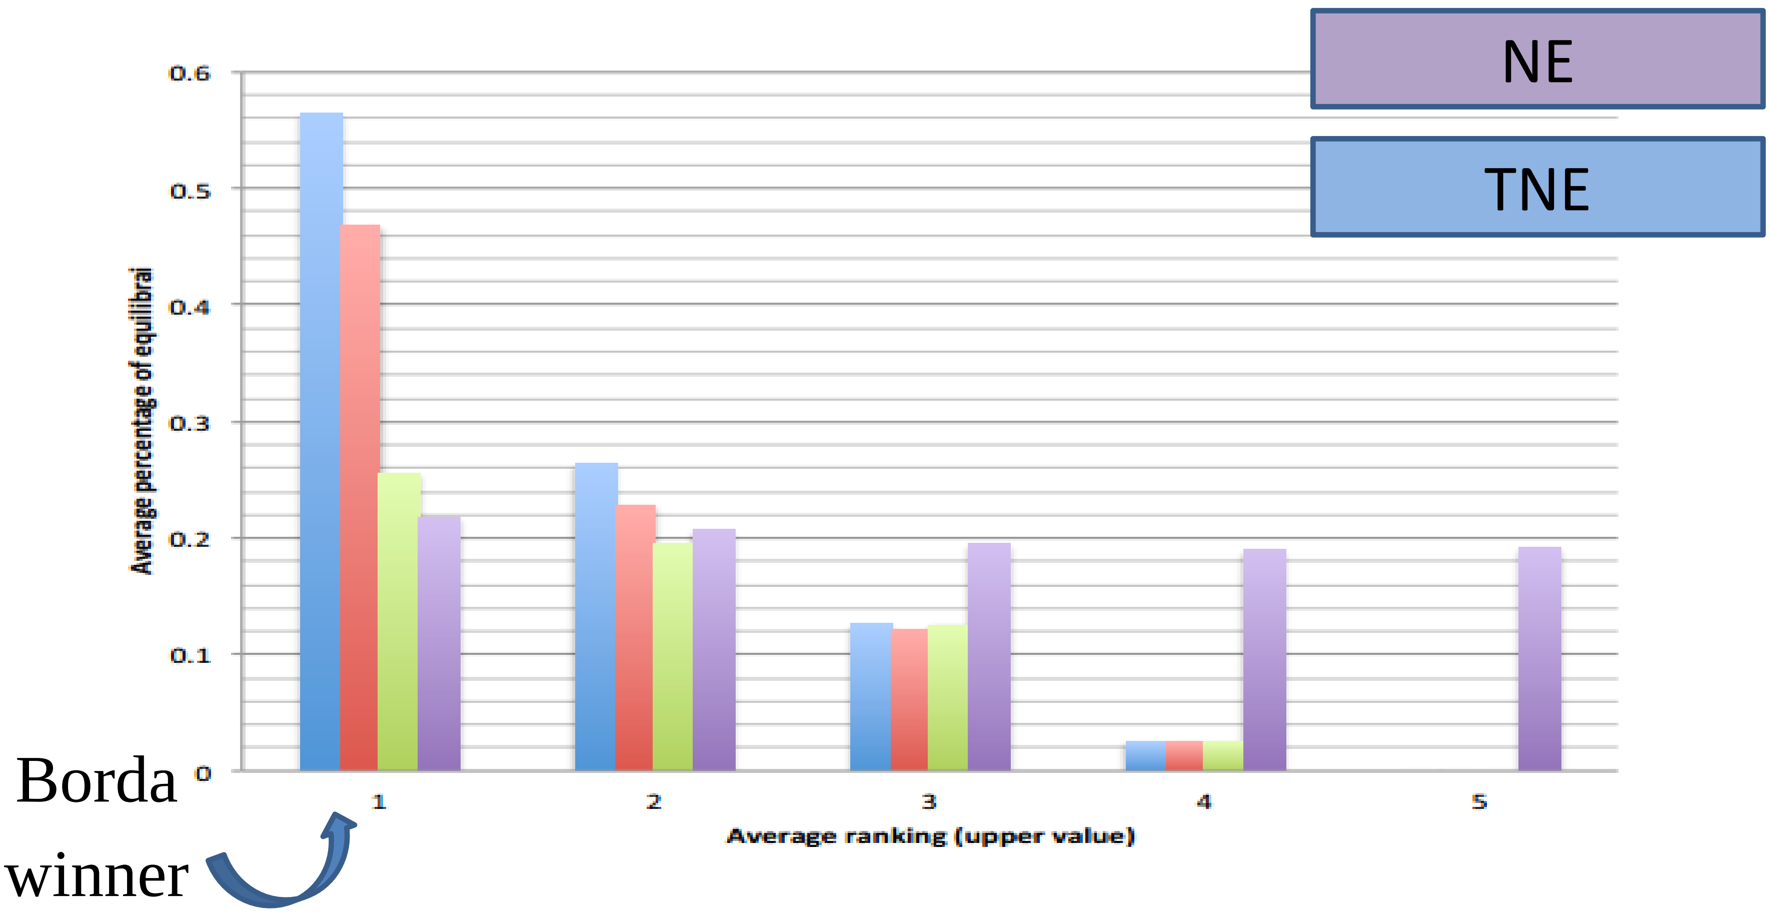

Voting equilibrium

1. Nash equilibrium
2. Strong equilibrium
3. Truth-bias
4. **Equilibrium under uncertainty**

Voters with Bayesian reasoning

- Typically voters do not know $\boldsymbol{R}$ exactly
- Suppose voters' utility $u_i \in \mathbb{R}^m$ is sampled from a (known) distribution over all types
- Each voter predicts the probability p_{xy} that x, y are tied, for any $x, y \in A$
- Then the expected utility of voting for x is

$$E[x|\boldsymbol{p}] = \sum_{y \neq x} p_{xy} \left(u_x - u_y\right)$$

- (this is for Plurality but can be generalized)

Voters with Bayesian reasoning

$$E[x|\boldsymbol{p}] = \sum_{y \neq x} p_{xy}\,(u_x - u_y)$$

- Each voter is assumed to vote for the candidate x that maximizes $E[x|\boldsymbol{p}]$
- If we know how a voter of type u votes, we can estimate candidates' scores $\boldsymbol{s}$
- Then estimate pivot probabilities $\boldsymbol{p}$:

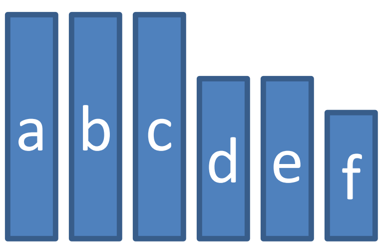

Case 1: $p_{ab} = p_{bc} = p_{ac} \sim \frac{1}{3}$

$p_{xy} < \varepsilon$ for all others

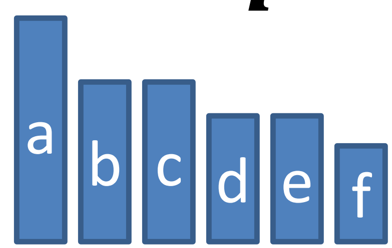

Case 2: $p_{ab} = p_{ac} \sim \frac{1}{2}$

$p_{xy} < \varepsilon$ for all others

A voting equilibrium

- A voting equilibrium for profile $\boldsymbol{u}$ is $\boldsymbol{s}$ and $\boldsymbol{p}$ such that
  - Pivot probabilities $\boldsymbol{p}$ are consistent with $\boldsymbol{s}$
  - If all voter types maximize their expected utility according to $\boldsymbol{p}$, scores are $\boldsymbol{s}$

- Theorem [Myerson&Weber'93]: A voting equilibrium exists for any scoring rule

Trembling hand perfection

- Suppose each vote is miscounted with some small probability ε
- Thus every voter has some chance to be pivotal
- A TH-equilibrium is a voting profile that has no deviation when $\varepsilon \rightarrow 0$

[Messner & Polborn'04] show that in any TH-equilibrium in Plurality, at most 2 candidates get votes.

- This phenomenon is known as "Duverger's law"

Course outline

The G-S theorem

More negative results

Achieving truthfulness
By additional assumptions
(“workarounds” for G-S)

Relax truthfulness:
- Rational voting and equilibrium analysis
- Iterative voting and convergence

Relax rationality:
Heuristic voting

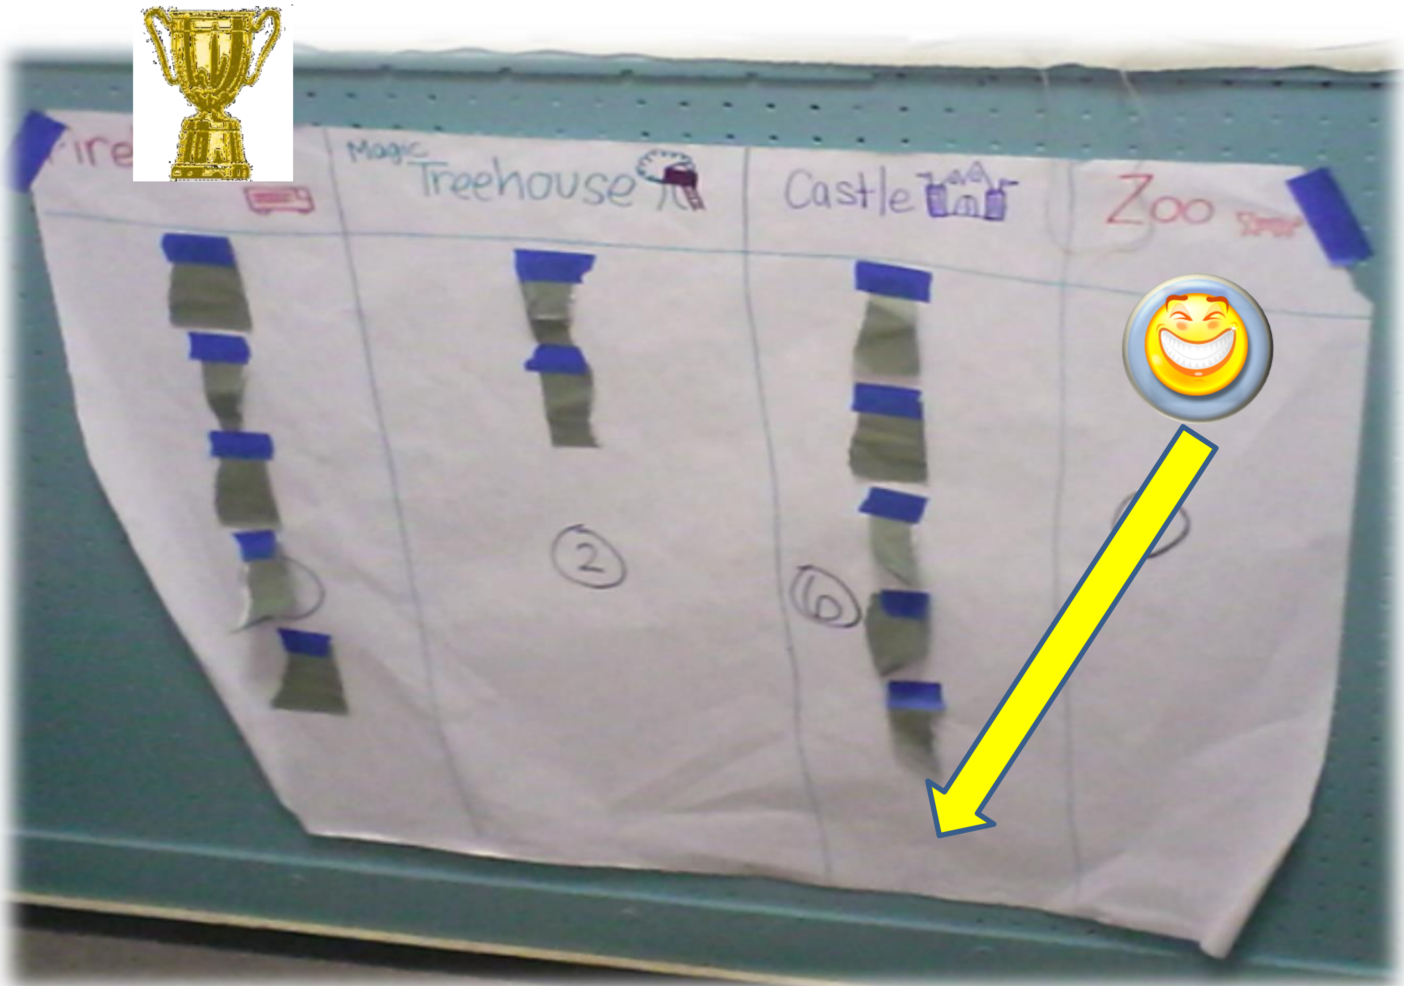

Least preferred

Most preferred

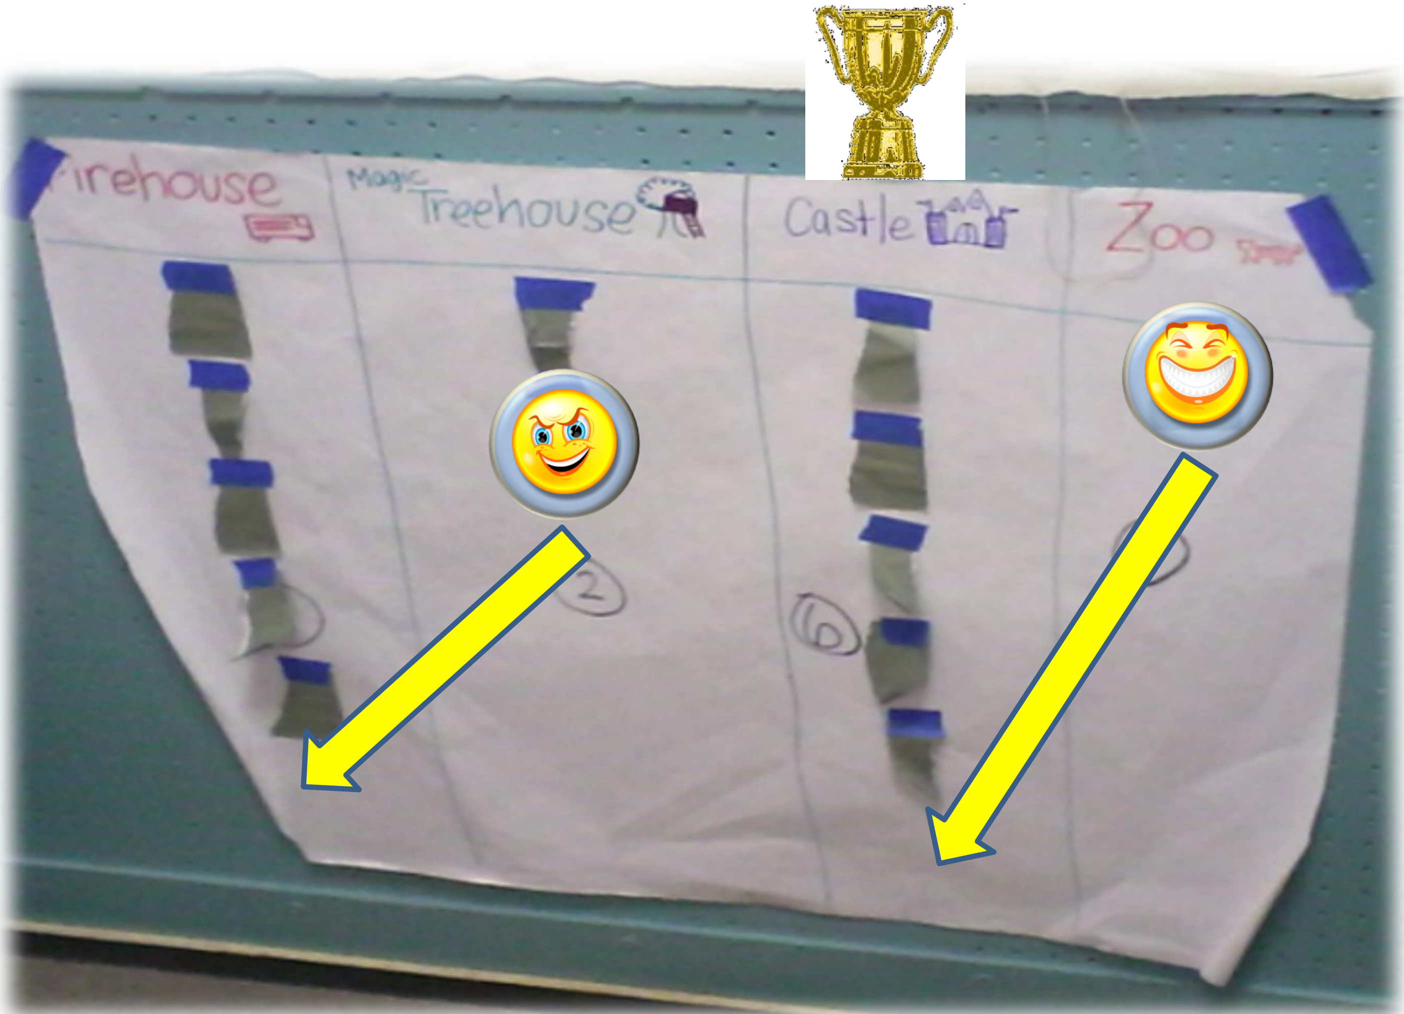

Most preferred

Least preferred

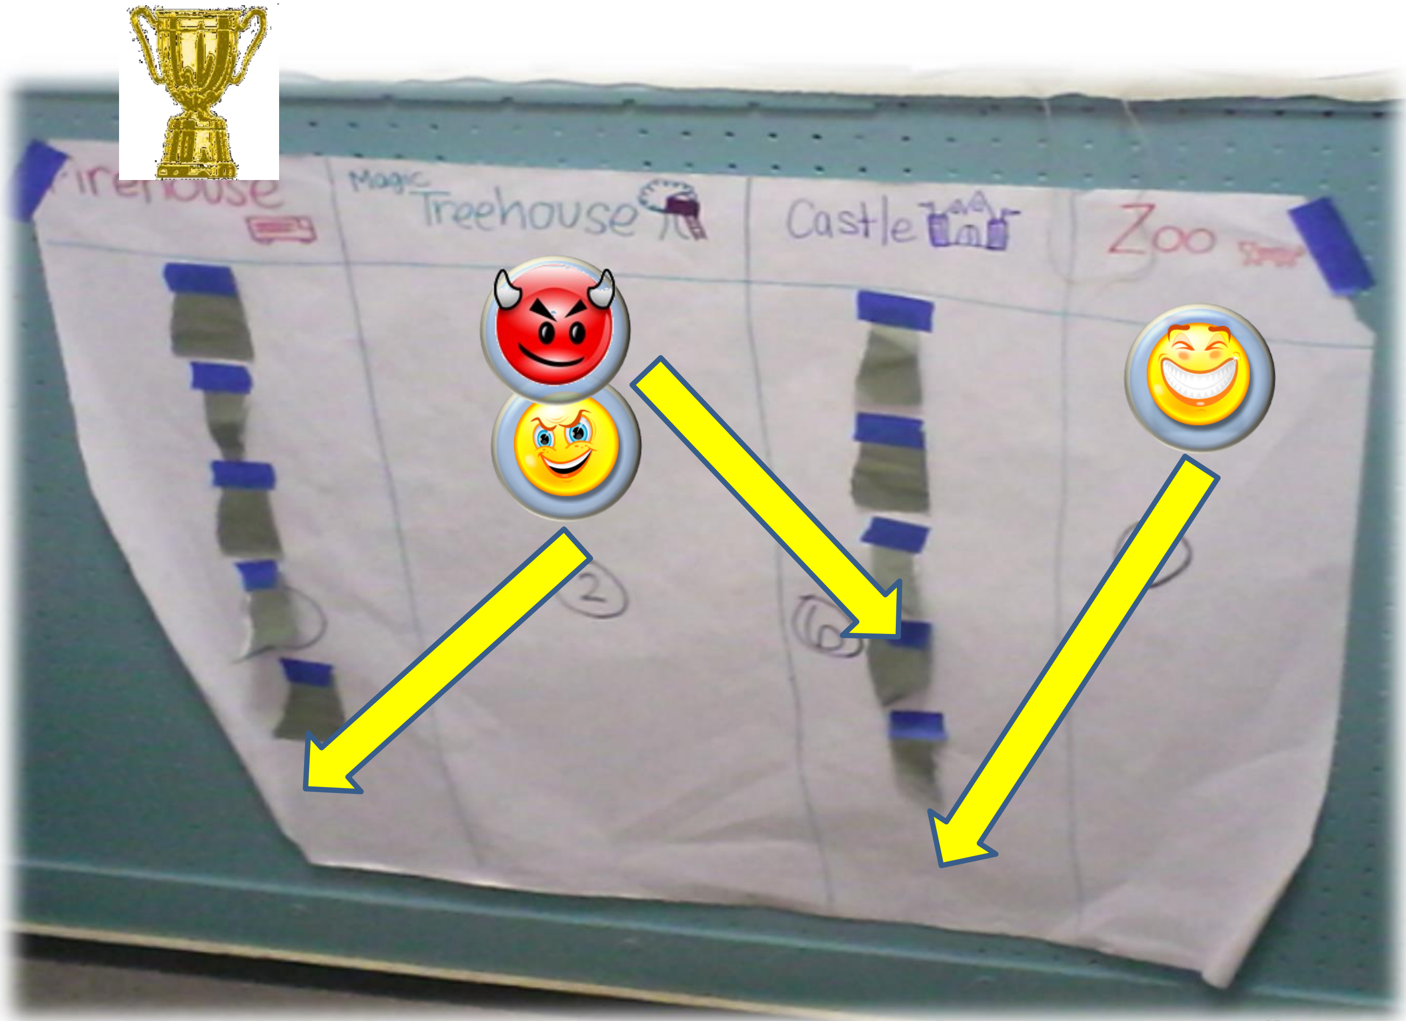

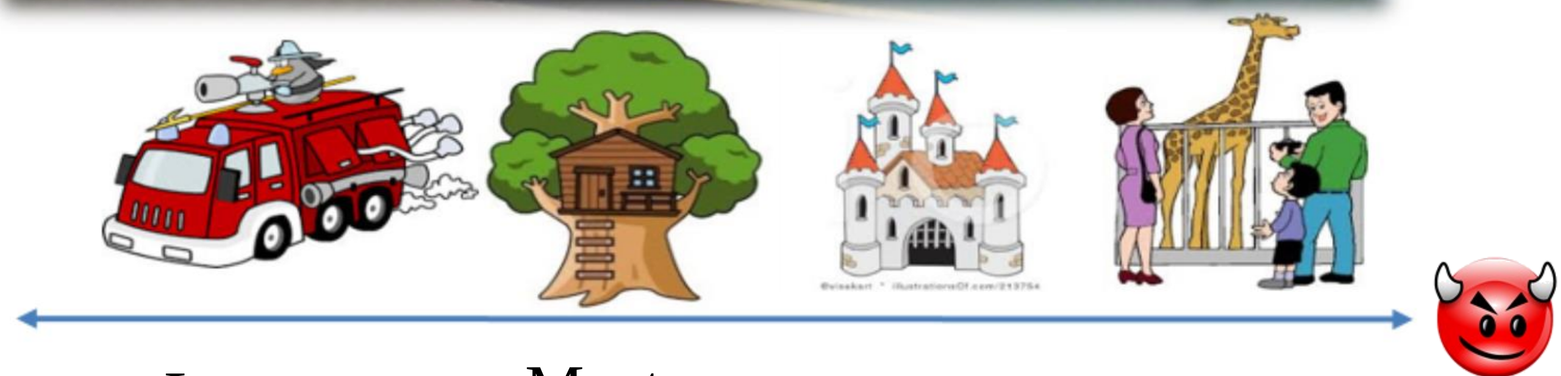

Least preferred

Most preferred

Voting in turns

- We allow each voter to change his vote
- Only one voter may act at each step
- The game ends when there are no objections

This mechanism is implemented in some on-line voting systems, e.g. in Facebook, Doodle, etc.

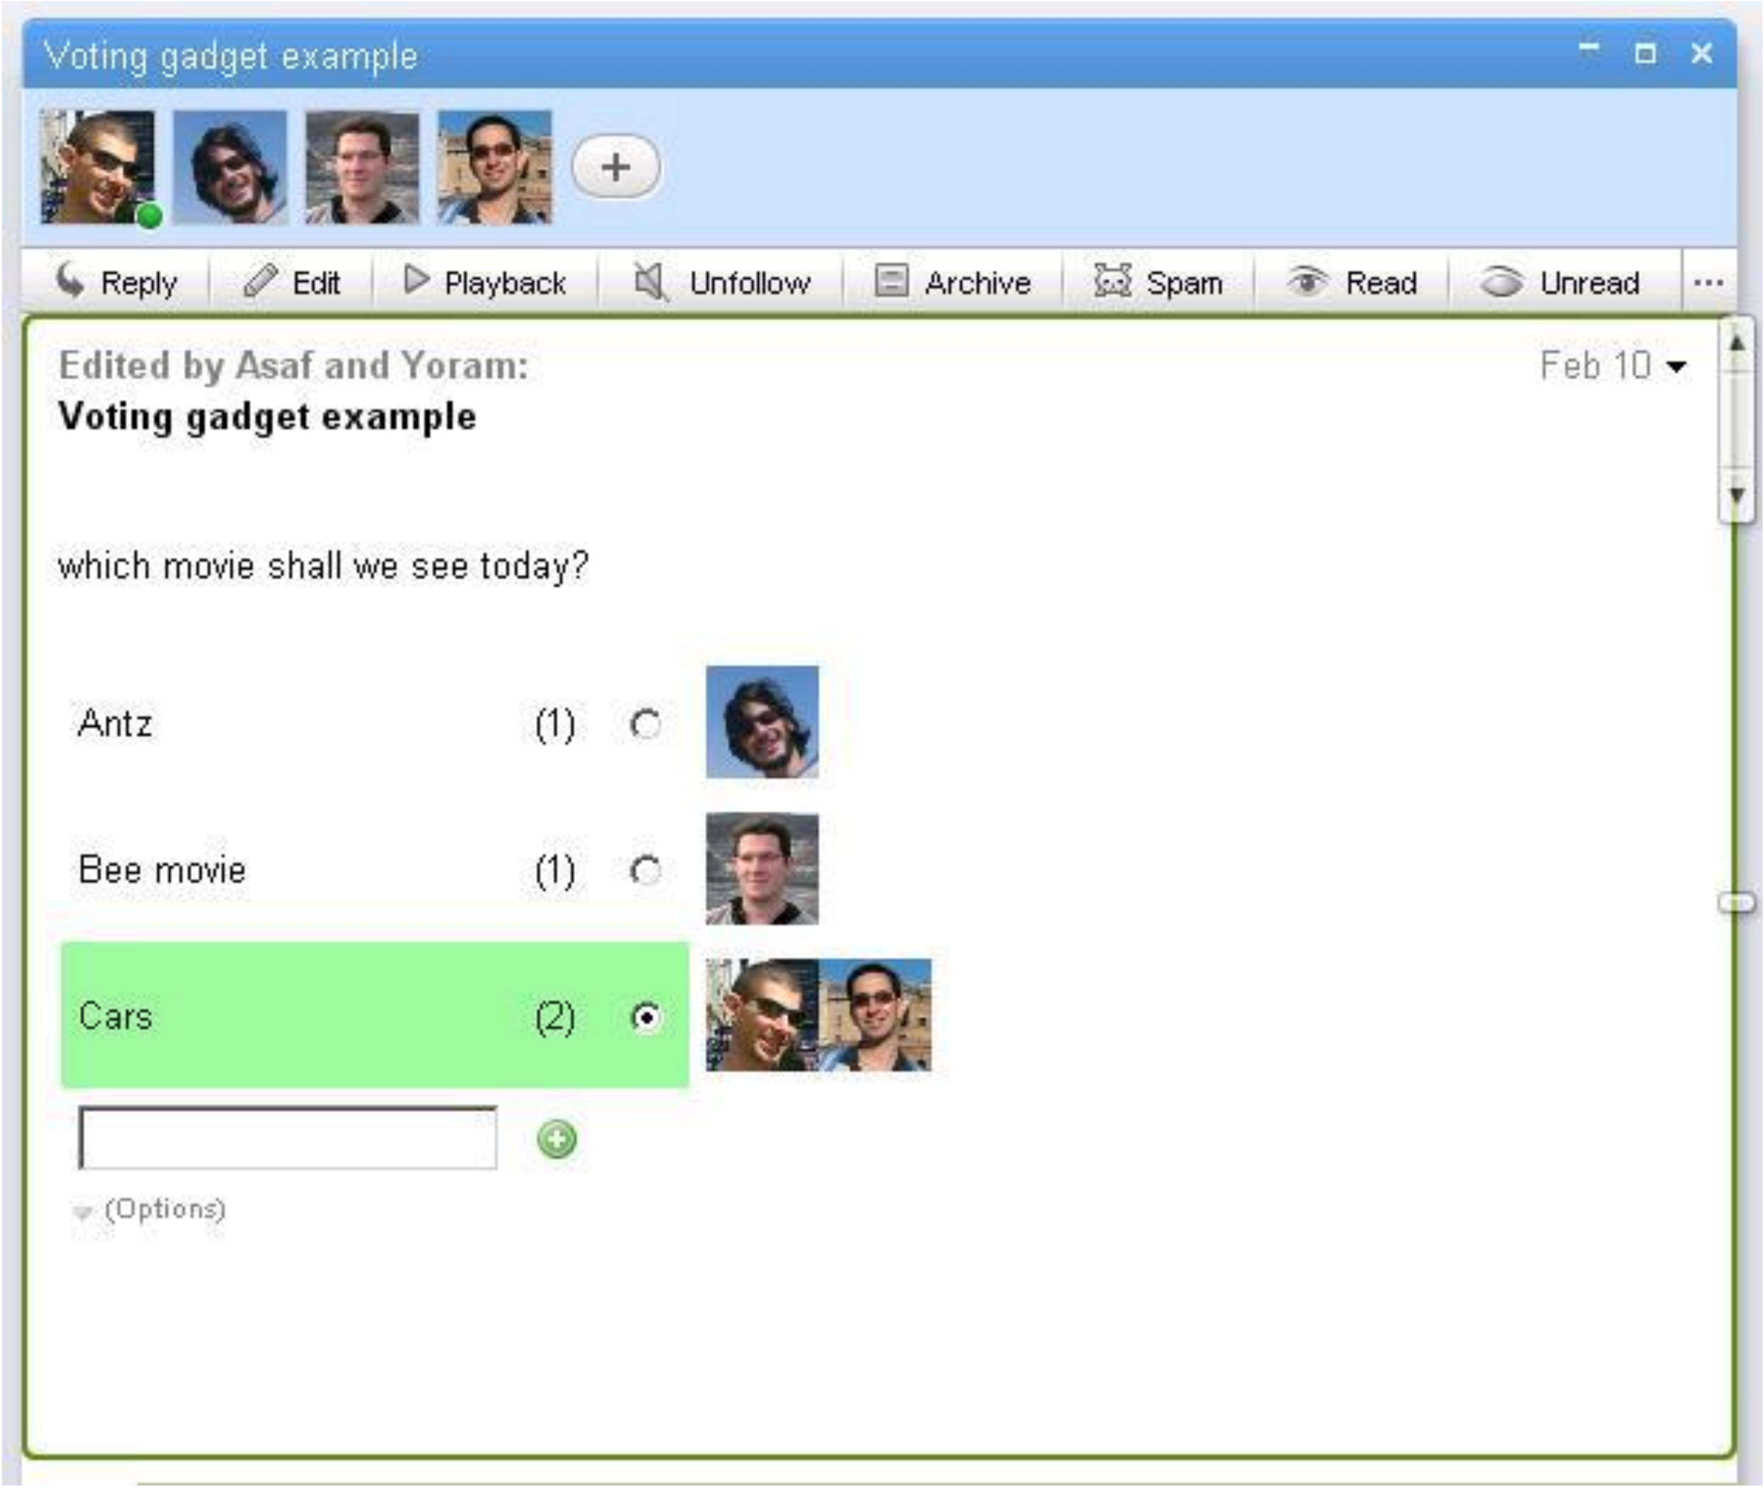
Voting gadget example
Reply
Edit
Playback
Unfollow
Archive
Spam
Read
Unread
Edited by Asaf and Yoram:
Feb 10
Voting gadget example
which movie shall we see today?
Antz (1)
Bee movie (1)
Cars (2)
(Options)

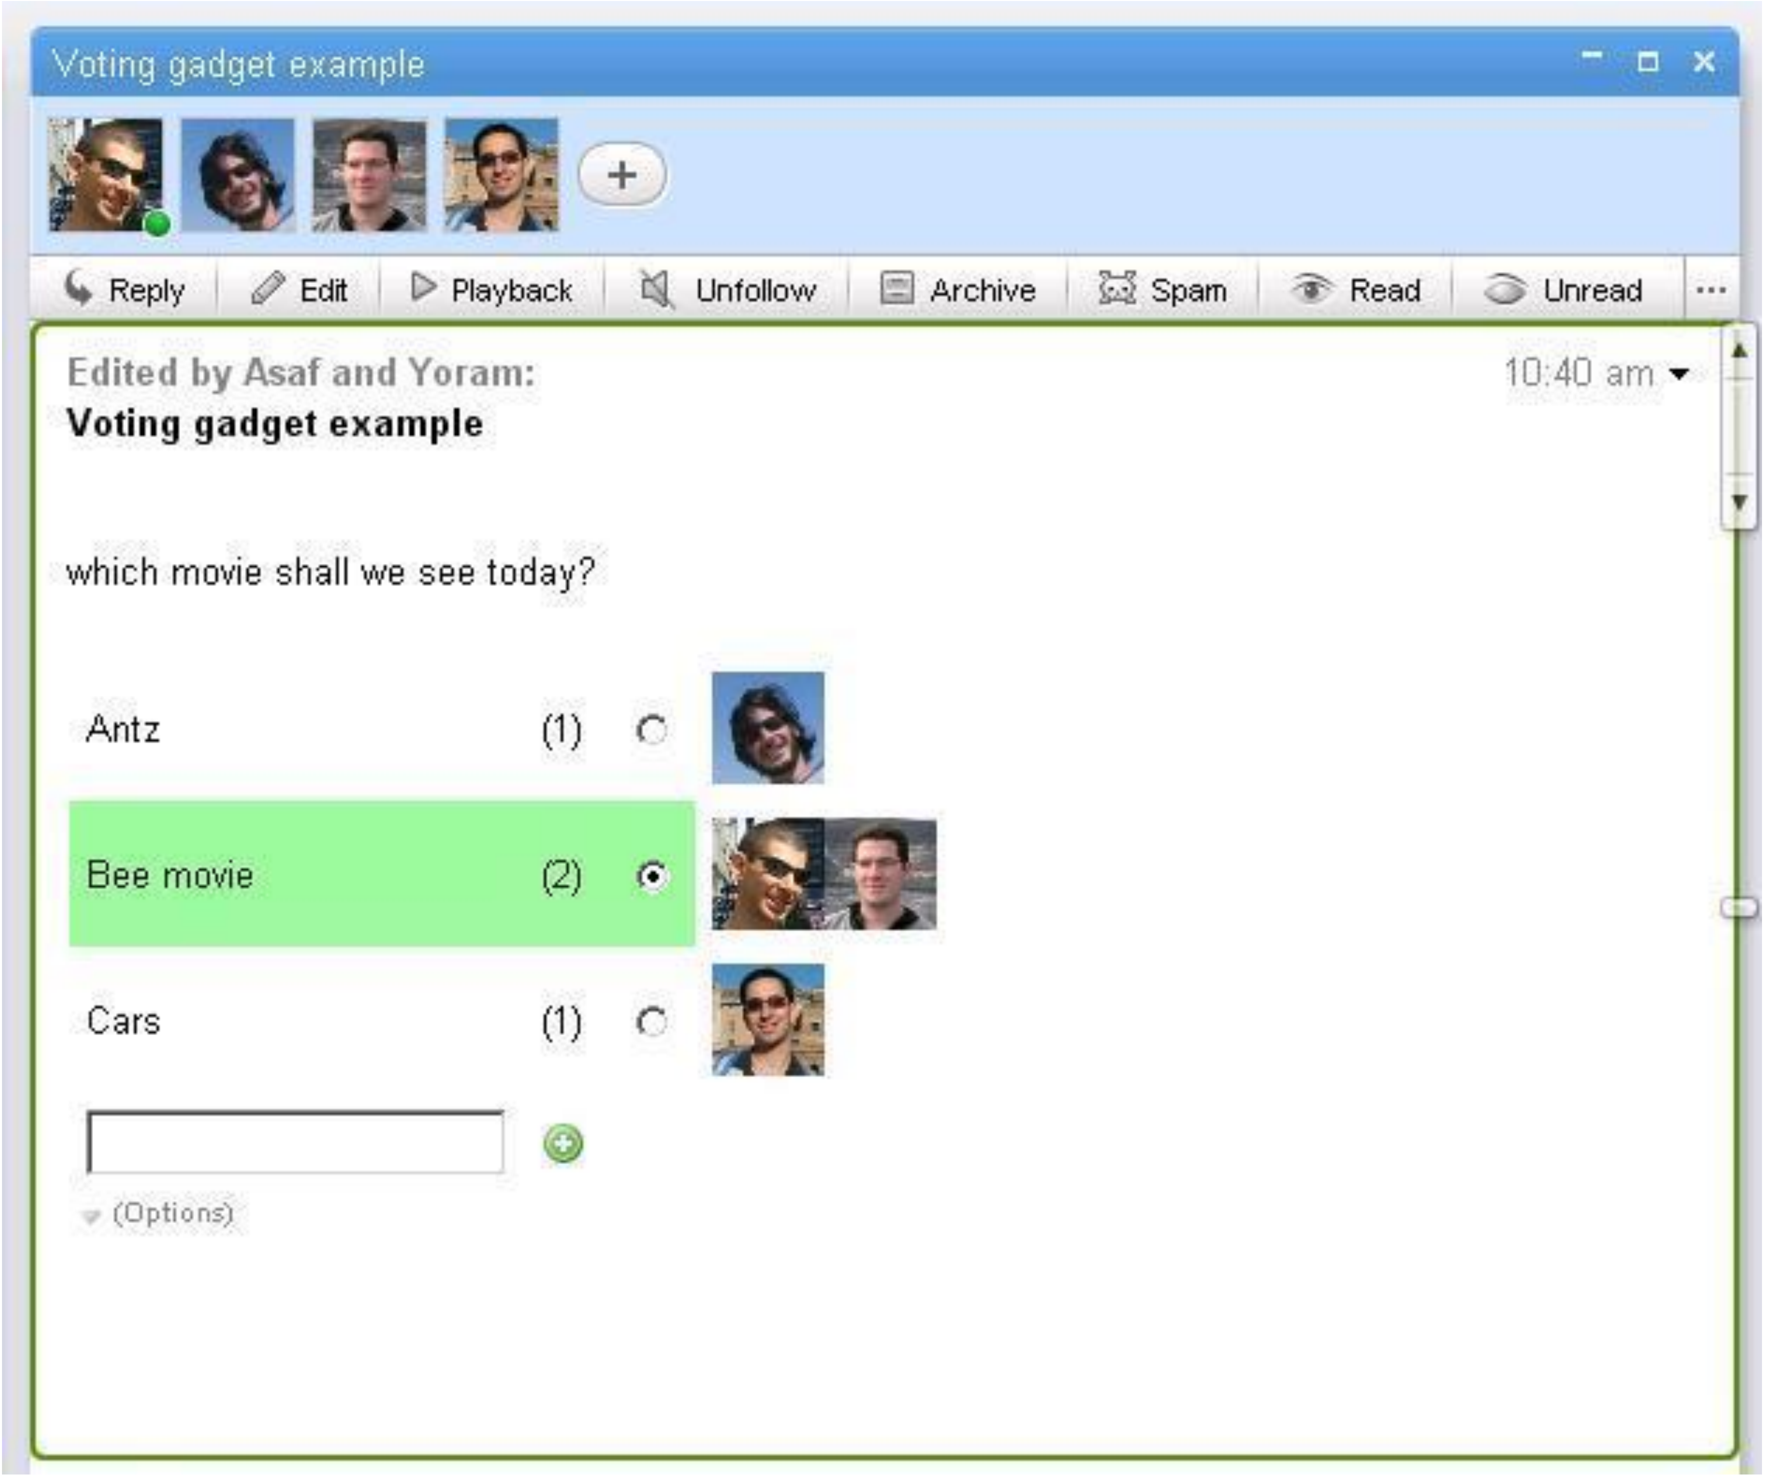
Voting gadget example
Reply
Edit
Playback
Unfollow
Archive
Spam
Read
Unread
Edited by Asaf and Yoram:
Voting gadget example
10:40 am
which movie shall we see today?
Antz (1)
Bee movie (2)
Cars (1)
(Options)

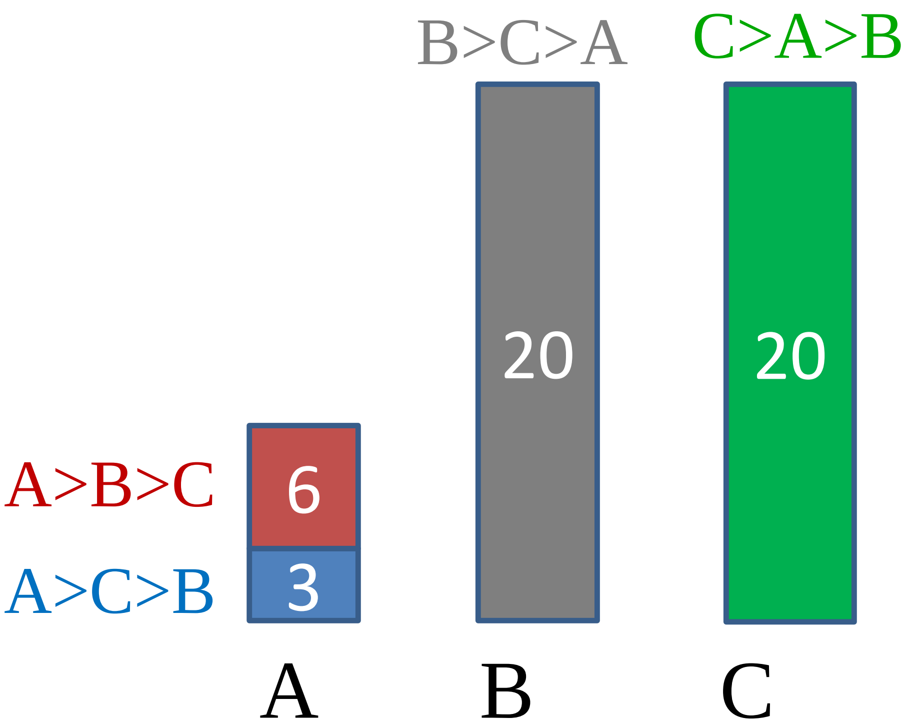

(Lexicographic tie braking)

Some games **always** converge

Theorem [M. et al.'10]: Plurality games converge from *any initial state*.

Assumes: all voters have *equal weights* and always use *direct-reply.*

Not true otherwise.

Other voting rules

- Studied by [Lev & Rosenschein’12, Reyhani & Wilson ‘12] and others.
- Veto also converges
- For many other voting rules there are counter examples (cycles)
  - Weighted Plurality
  - Borda
  - Minimax
  - Copeland
  - …

Implications

- What are the implications of convergence?
  - Will voters reach a “better” outcome?
  - Note that we still have all Nash equilibria, including all the weird ones
  - Fewer if we assume voters start by being truthful
- We want to compare the equilibrium outcomes to the truthful outcomes

Dynamic Price of Anarchy

- Approach #1: for score-based rules (where score is a measure to candidate's quality), compare the scores of equilibrium and truthful winners.

Worst case approach

- $DPoA(f,R) = \min_{R} \min_{R' \in EQ^T(f,R)} \frac{score_f(f(R'),R)}{score_f(f(R),R)}$, where $EQ^T(f,R)$ contains all equilibrium states reachable from the truth

- Results [Branzei et al.'13]:
  - In Plurality, $DPoA$ close to 1
  - In Borda, $DPoA = \Omega(n)$

Objective quality measures

- Approach #2: use external quality measures independent of f:
  - Social welfare
  - Condorcet consistency
  - Distance from ground truth (when exists)

- Study the **average effect** using simulations

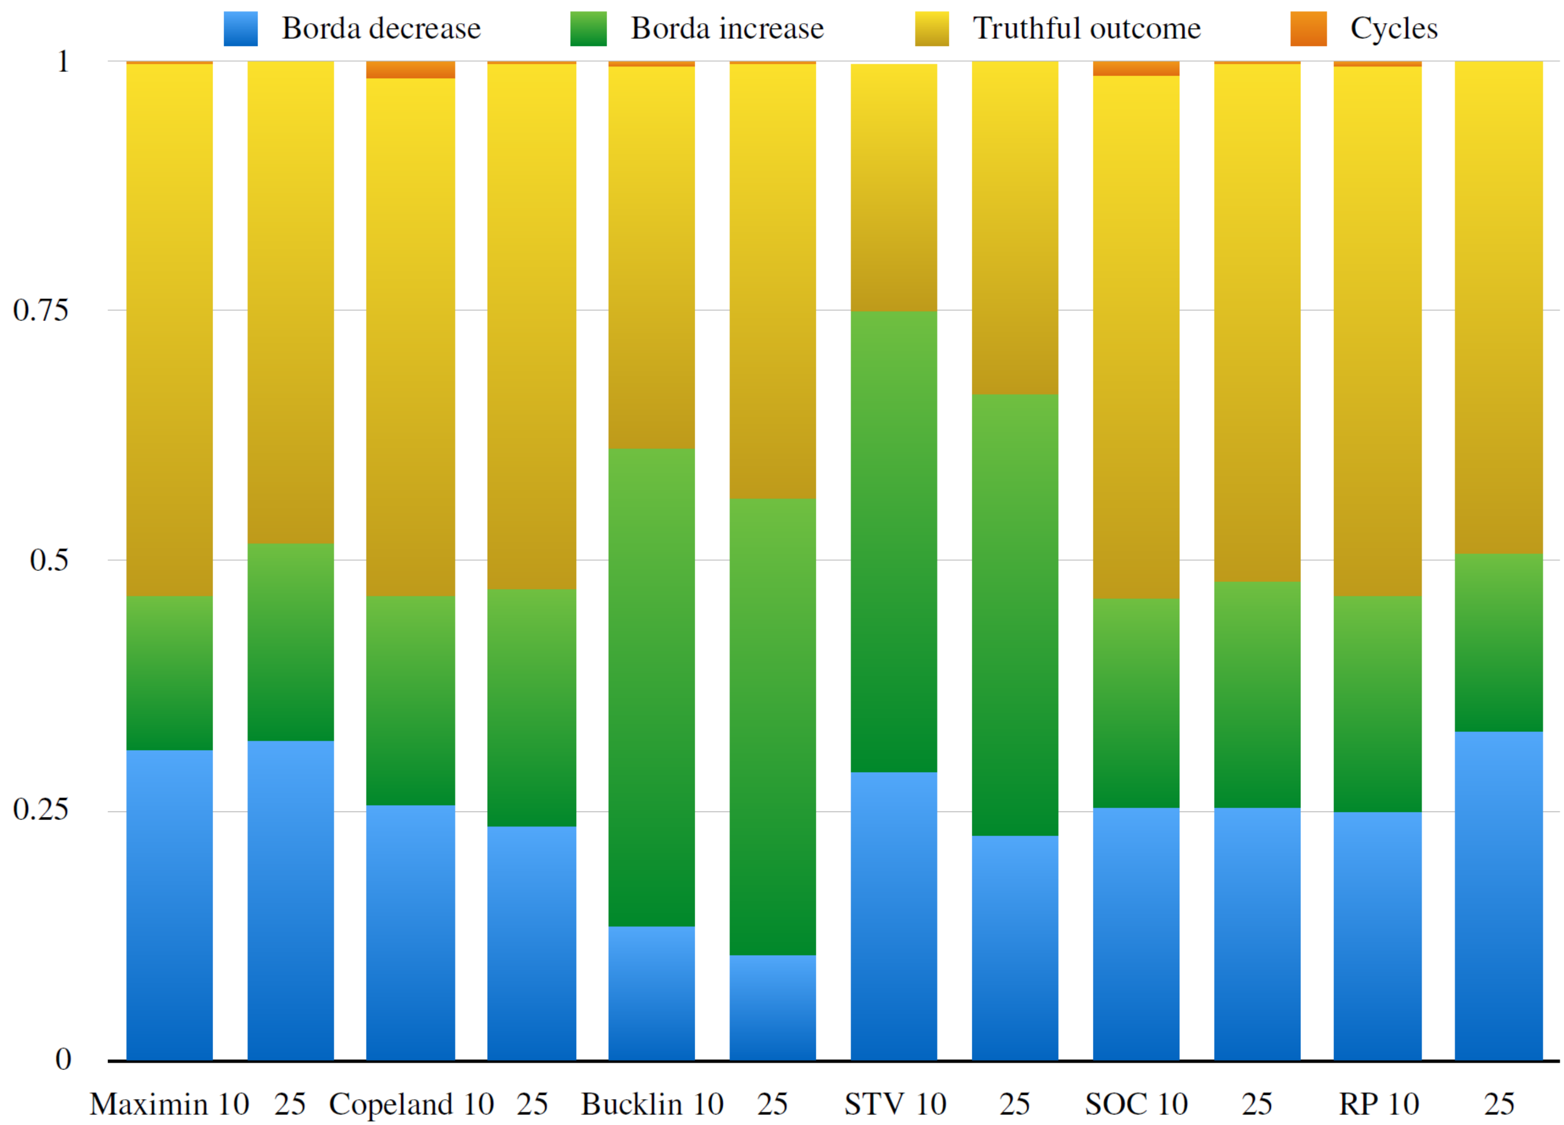

[Koolyk et al.'16]: mixed effect on social welfare, Condorcet consistency improves.

Course outline

The G-S theorem

More negative results

Achieving truthfulness
By additional assumptions
("workarounds" for G-S)

Relax truthfulness:
- Rational voting and equilibrium analysis
- Iterative voting and convergence

Relax rationality:
Heuristic voting

Taking a step back

- “best response” is a myopic heuristic

Does not look forward

“rule of thumb”

Taking a step back

- "best response" is a myopic heuristic
- Other heuristics were suggested :
  - "Second chance": promote the second best candidate [Grandi et al.'13]
  - "Best Upgrade": look at all candidates preferred over the winner, do best-reply to one of them if possible [Grandi et al.'13]
  - "k-pragmatist": look at the leading k candidates, vote for best one [Reijngoud&Endriss'12]
  - "T-threshold": look at all candidates above some threshold T, vote for best one
  - "far-sighted": best-reply assuming k more voters will change their vote [Obraztsova et al.'15]
  - "leader rule" (Approval only): approve everyone you prefer to the current leader [Laslier'09]

Scoring rules only

Properties

- Different heuristics require different levels of information [Reijngoud&Endriss'12], e.g:
  - All votes
  - Only order of candidates
  - Only the winner
  - No information

Convergence

- Various results for combinations of

Voting rule X **heuristic**

Voting rule	k-pragmatist	Second Chance	Best upgrade	Upgrade	Unit Upgrade
PSRs	V [RE12]	V	V [GLR+13]	?	V* [OMM+15]
Maximin	V [RE12]	V	V [GLR+13]	V [OMM+15]	V [OMM+15]
Copland	V [RE12]	V	V [GLR+13]	?	?
Bucklin	-	V	?	?	V [OMM+15]
all rules	-	V [GLR+13]	?	?	?

- Some general guidelines in [Obraztsova et al.'15]
  - Works for many such combinations
  - Assumes voters start from the truth

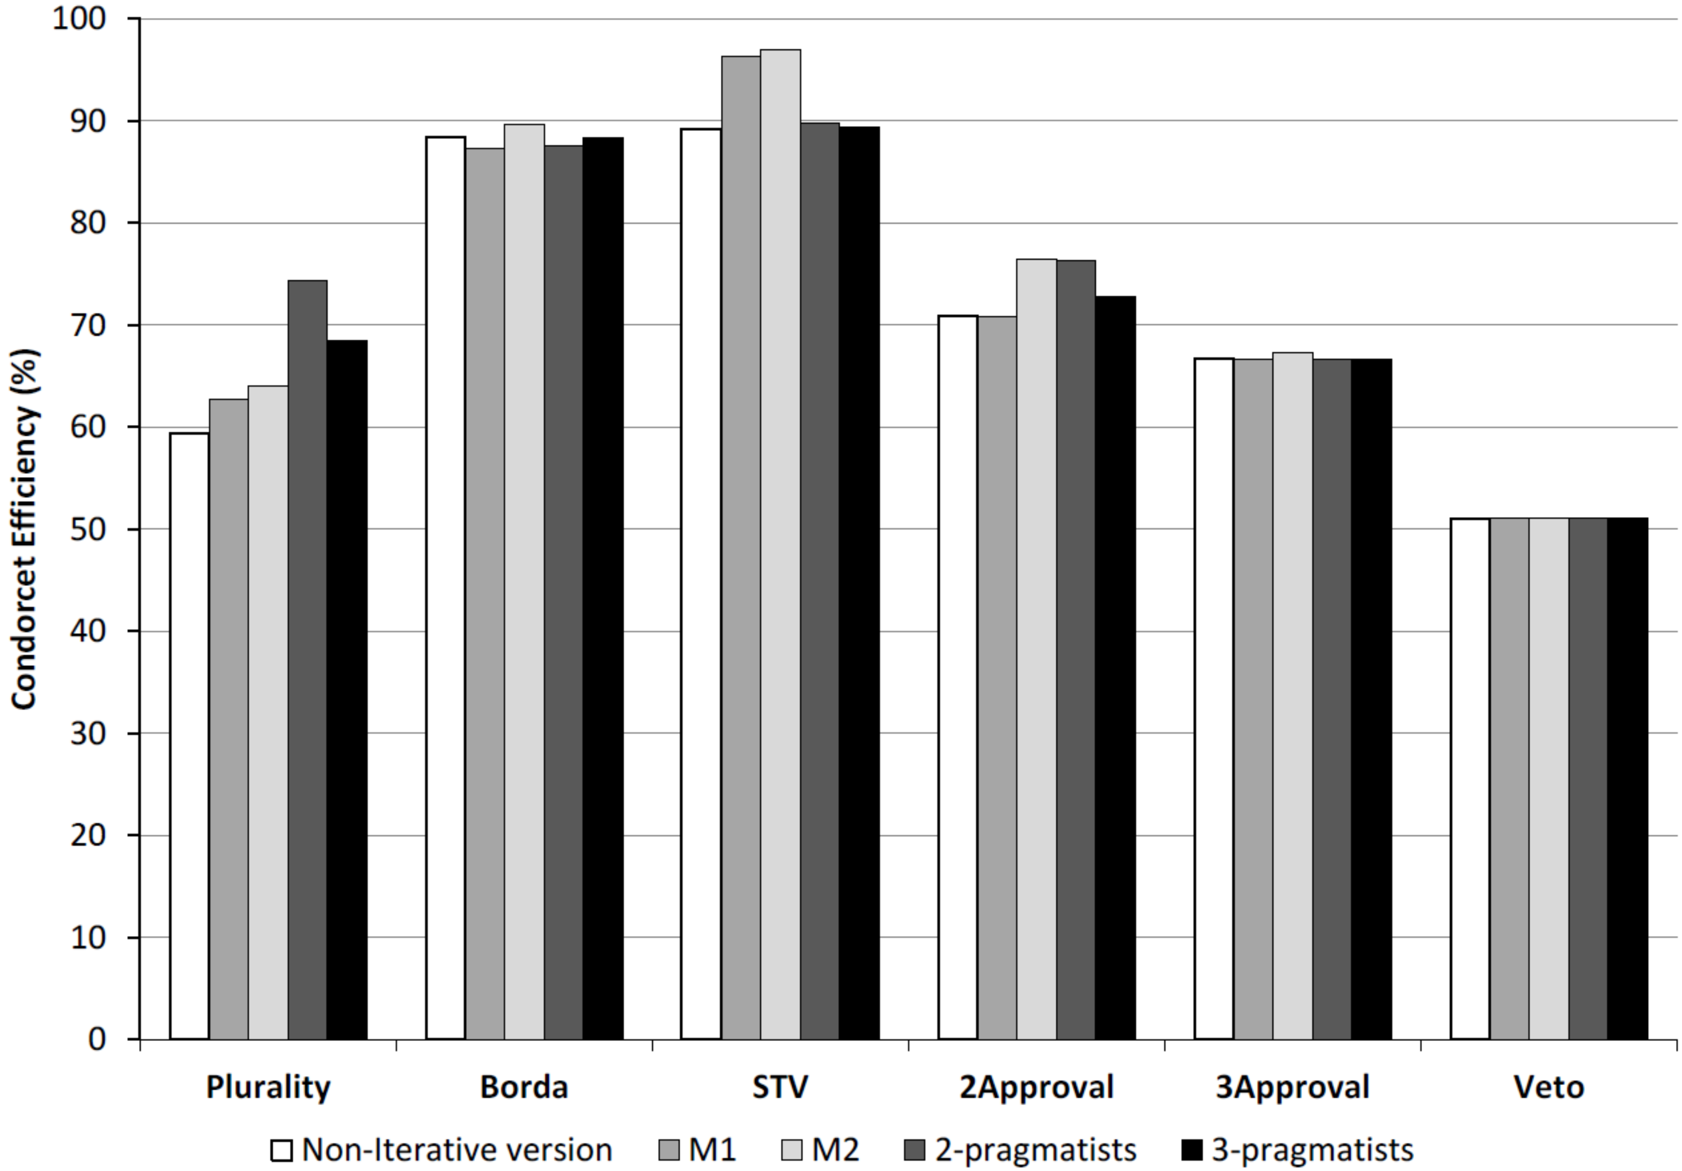

- Effect of heuristic voting on Condorcet efficiency [Grandi et al. '13]

Local-Dominance

- Some heuristics seem kind of arbitrary
- We want to derive the behavior from basic (game-theoretic?) principles
  - Attempt 0: best-response
- The key idea: add **uncertainty**
  - The voters are unsure about the exact outcome
  - Unlike [Myerson&Weber'93]:
    - No distributions
    - No cardinal utilities
    - Care about dynamics rather than just equilibrium

Epistemic model

Prospective scores $\boldsymbol{s}$

- E.g. from a poll
- "world state"

Uncertainty level $r_i \geq 0$

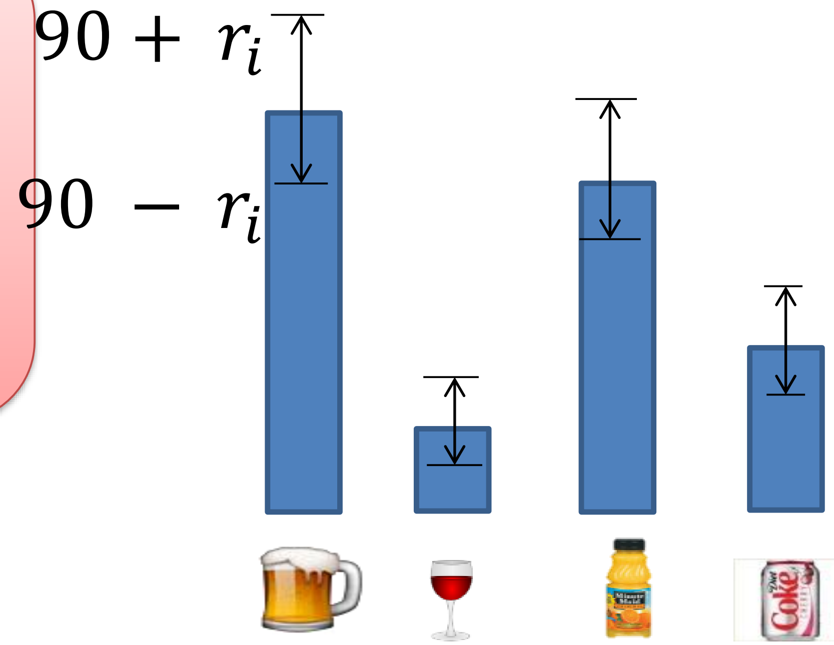

Voter i considers as "possible" all states close enough to $\boldsymbol{s}$. $S(\boldsymbol{s}, r_i) = \{\boldsymbol{s}' : \|\boldsymbol{s}' - \boldsymbol{s}\| \leq r_i\}$

– Example I: "*additive uncertainty*"

Behavioral model

Rational agents avoid dominated strategies!

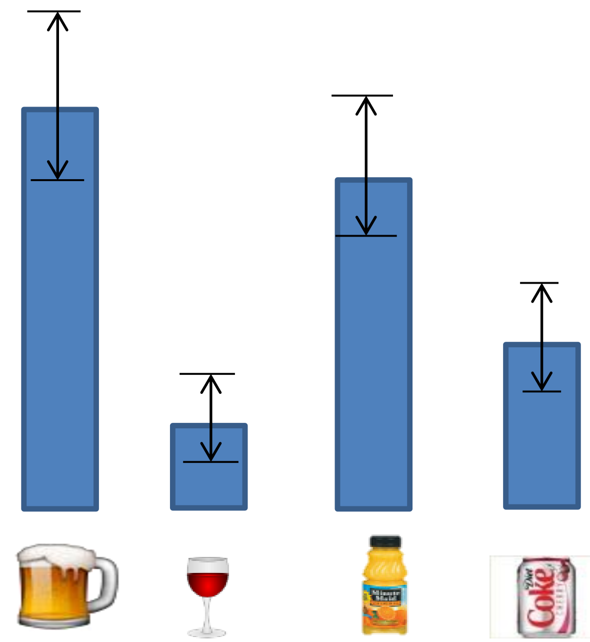

<u>Def. I (*Local dominance*)</u>: A candidate c' *S-dominates* candidate c if it is always weakly better for i to vote for c'.

in every state $\boldsymbol{s}' \in S$

Behavioral model

Rational agents avoid dominated strategies!

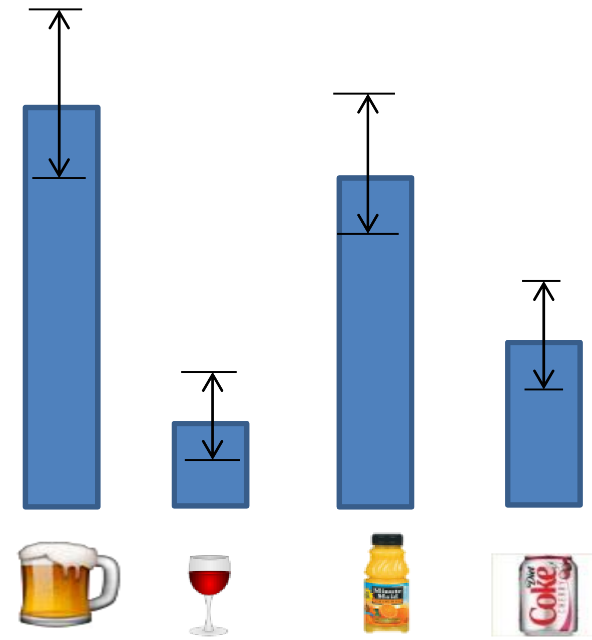

<u>Iterative setting</u>: As long as your vote is locally dominated, switch to a candidate that dominates it. **Otherwise – stay.**

Local dominance move

Lemma: All dominance relations in state $\boldsymbol{s}$ are characterized by a single threshold $T(\boldsymbol{s}, r_i)$: (depends on winner's score)

c **is dominated** iff **below the threshold** *or* **least preferred**.*

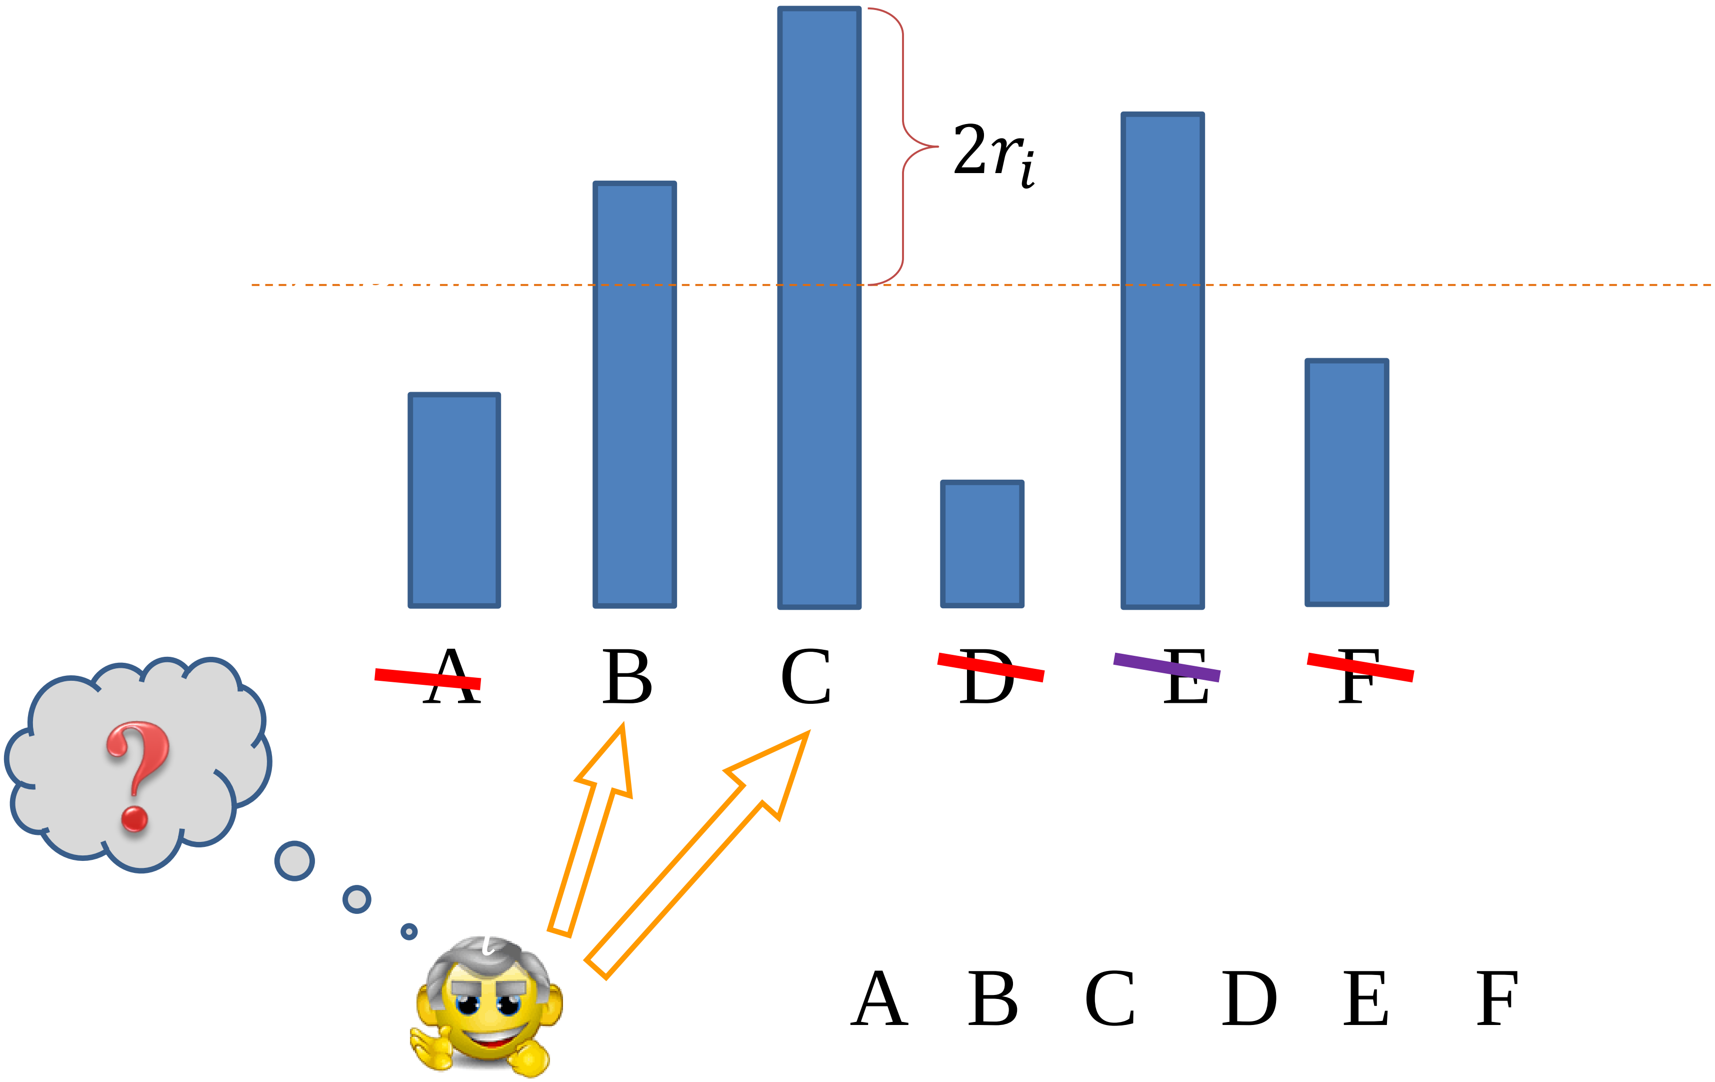

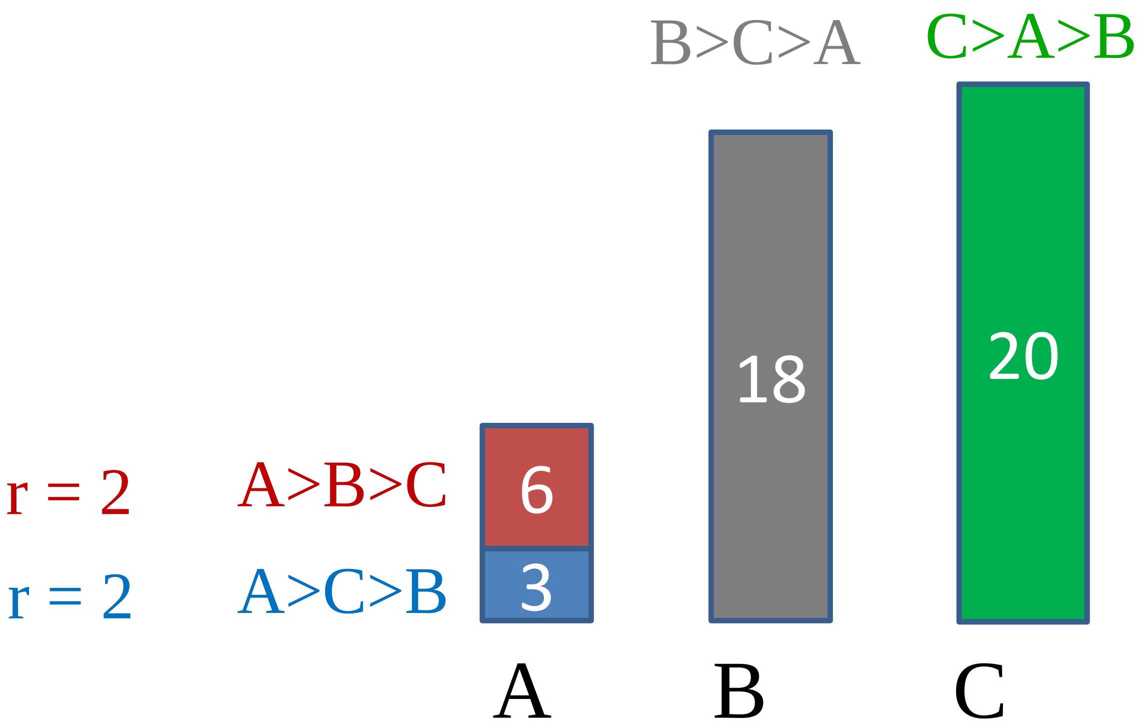

(Lexicographic tie braking)

Results

<u>Theorem</u> *[M., AAAI'15]*:

Any sequence $\boldsymbol{s}^0 \rightarrow \boldsymbol{s}^1 \rightarrow \boldsymbol{s}^2 \rightarrow \cdots$ of Local-dominance moves is acyclic (must converge).

In particular, a voting equilibrium always exists.

Still true for:

- Arbitrary initial (non-truthful) profile
- Arbitrary order of players
- Diverse uncertainty levels r_i

Results

<u>Theorem</u> *[M., AAAI'15]*:

Any sequence $\boldsymbol{s}^0 \rightarrow \boldsymbol{s}^1 \rightarrow \boldsymbol{s}^2 \rightarrow \cdots$ of Local-dominance moves is acyclic (must converge).

In particular, a voting equilibrium always exists.

Prop. *[M. et al., AAAI'10]:*

"best-response converges to a Nash equilibrium."

Follows as a special case!

Proof sketch: $r_i = 0$ for all i $\Rightarrow$ $S(\boldsymbol{s}, r_i) = \{\boldsymbol{s}\}$

$\Rightarrow$ Local-dominance ≡ Best response

$\Rightarrow$ Voting equilibrium ≡ Nash equilibrium ■

Equilibrium properties
(computer simulations)

- Decisiveness

- Duverger Law

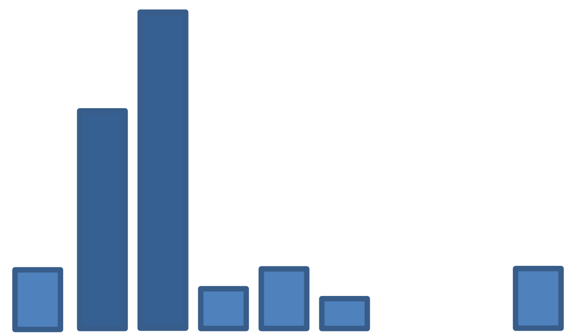

- Participation

- Welfare

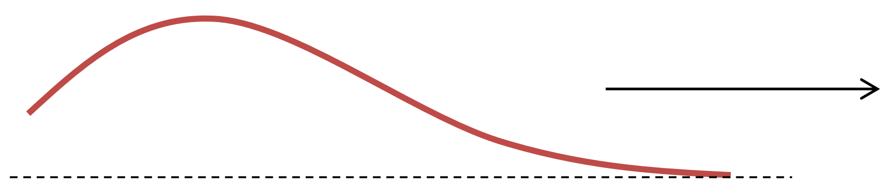

[M., Lev, Rosenschein, EC'14]

Not covered

- Behavioral voting experiments
- How do people really vote?
- See lecture notes for some references