Voting in Metric Spaces

Prasanna Ramakrishnan

Stanford University

Includes joint work with:



Moses Charikar
Stanford University

Kangning Wang
Stanford University

Hongxun Wu
UC Berkeley

Elections and Voting

- Voters choose from candidates
- Voters express *preferences* over candidates
- Preferences are aggregated, winner is chosen
- Often studied: ranked preference lists
- Voting rule: algorithm that maps ranked preferences to winning candidate

A central question in Social Choice Theory:

Can we design ***effective* voting rules**?

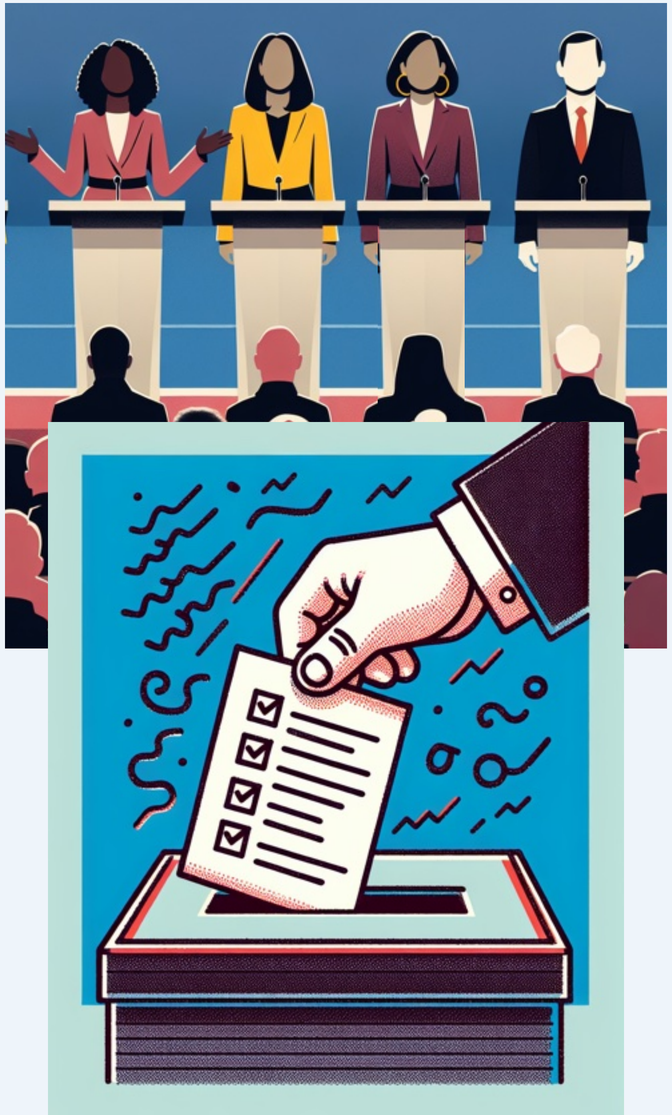

Elections can model...

Voters choosing

a representative

Community members choosing

a location for a public facility

An organization deciding

who to hire

Friends deciding

where to eat

Friends deciding

what to play

Emerging applications?

Google DeepMind 2023-12-11

Evaluating Agents using Social Choice Theory

Marc Lanctot[1], Kate Larson[1,2], Yoram Bachrach[1], Luke Marris[1], Zun l
Anthony[1], Brian Tanner[4] and Anna Koop[1]
[1]Google DeepMind, [2]University of Waterloo, [3]University of Michigan, [4]Artificial.Ag

We argue that many general evaluation problems can be viewe Each
task is interpreted as a separate voter which requires only or isons
of agents to produce an overall evaluation. By viewing the ag n, we
leverage centuries of research in social choice theory to derive principled evaluation frameworks with

Tasks = Voters, Agents = Candidates

What’s the ***best*** voting rule?

Early Social Choice Theory

- From the middle ages through the 19th century:
  - **Llull** (c.1235–1315)
  - **Cusanus** (1401–1464)
  - **von Pufendorf** (1632–1694)
  - **Borda** (1733–1799)
  - **Condorcet** (1743-1794)
  - **Dodgson** (1832–1898)

Can we design ***effective*** **voting rules**?

Majority support

Ramon Llull

Copeland Rule (Llull 1299)

Winner of most pairwise majority votes

See: "Social Choice Theory", *The Stanford Encyclopedia of Philosophy*

Early Social Choice Theory

- From the middle ages through the 19th century:
  - **Llull** (c.1235–1315)
  - **Cusanus** (1401–1464)
  - **von Pufendorf** (1632–1694)
  - **Borda** (1733–1799)
  - **Condorcet** (1743-1794)
  - **Dodgson** (1832–1898)

Can we design ***effective*** **voting rules**?

→ **Majority support**

Nicholas de Condorcet

Condorcet's Paradox

Majority preferences can be inconsistent!

(e.g. can prefer A over B, B over C, C over A)

Voter	Ranking
1	ABC
2	CAB
3	BCA

See: "Social Choice Theory", *The Stanford Encyclopedia of Philosophy*

The Prominent "Axiomatic" Approach

- Black (1948), Arrow (1950): define *necessary properties* ***(axioms)***, find rules satisfying them

Can we design ***effective*** **voting rules**?

Satisfies axioms

Arrow, Gibbard–Satterthwaite

No rules simultaneously satisfy **basic axioms**

- Unfortunate downsides
  - Little practical guidance on what rules to use

Kenneth Arrow

See: "Social Choice Theory", *The Stanford Encyclopedia of Philosophy*

Comparison of preferential electoral systems [hide]

System	Monotonic	Condorcet winner	Majority	Condorcet loser	Majority loser	Mutual majority	Smith	ISDA	LIIA	Independence of clones	Reversal symmetry	Participation, consistency	Later-no-harm	Later-no-help	Polynomial time	Resolvability
Schulze	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	No	No	No	Yes	Yes
Ranked pairs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	Yes	Yes
Tideman's Alternative	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	Yes	No	No	No	No	Yes	Yes
Kemeny–Young	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	Yes	No	No	No	No	Yes
Copeland	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	Yes	No	No	No	Yes	No
Nanson	No	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	Yes	No	No	No	Yes	Yes
Black	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	Yes	No	No	No	Yes	Yes
Instant-runoff voting	No	No	Yes	Yes	Yes	Yes	No	No	No	Yes	No	No	Yes	Yes	Yes	Yes
Smith/IRV	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	Yes	No	No	No	No	Yes	Yes
Borda	Yes	No	No	Yes	Yes	No	No	No	No	No	Yes	Yes	No	Yes	Yes	Yes
Baldwin	No	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No	No	Yes	Yes
Bucklin	Yes	No	Yes	No	Yes	Yes	No	No	No	No	No	No	No	Yes	Yes	Yes
Plurality	Yes	No	Yes	No	No	No	No	No	No	No	No	Yes	Yes	Yes	Yes	Yes
Contingent voting	No	No	Yes	Yes	Yes	No	No	No	No	No	No	No	Yes	Yes	Yes	Yes
Coombs[37]	No	No	Yes	Yes	Yes	Yes	No	No	No	No	No	No	No	No	Yes	Yes
Mini-Max[specify]	Yes	Yes	Yes	No	No	No	No	No	No	No	No	No	No	No	Yes	Yes
Anti-plurality[37]	Yes	No	No	No	Yes	No	No	No	No	No	No	Yes	No	No	Yes	Yes
Sri Lankan contingent voting	No	No	Yes	No	No	No	No	No	No	No	No	No	Yes	Yes	Yes	Yes
Supplementary voting	No	No	Yes	No	No	No	No	No	No	No	No	No	Yes	Yes	Yes	Yes
Dodgson[37]	No	Yes	Yes	No	No	No	No	No	No	No	No	No	No	No	No	Yes

The Adoption of Ranked Choice Voting Raised Turnout 10 Points

An expansive new study by University of Missouri-St. Louis Professor, David Kimball, and Ph.D. candidate, Joseph Anthony, examines the impact of ranked choice voting (RCV) on voter turnout in 26 American cities across 79 elections.

MINIMIZES STRATEGIC VOTING

Ideally, voters vote for candidates they support, not against those they oppose most. In most cases with our current election system, voters often feel the need to vote for the "lesser of two evils" because they believe their favorite candidate is less likely to win.

Source: FairVote

The Prominent "Axiomatic" Approach

- Black (1948), Arrow (1950): define *necessary properties (**axioms**)*, find rules satisfying them

Can we design ***effective* voting rules**?

→ **Satisfies axioms**

Arrow, Gibbard–Satterthwaite

No rules simultaneously satisfy **basic axioms**

- Unfortunate downsides
  - Little practical guidance on what rules to use
  - Weak motivation for new rules

Kenneth Arrow

See: "Social Choice Theory", *The Stanford Encyclopedia of Philosophy*

Alternative: **quantitative approach**, leverage ***spatial structure***

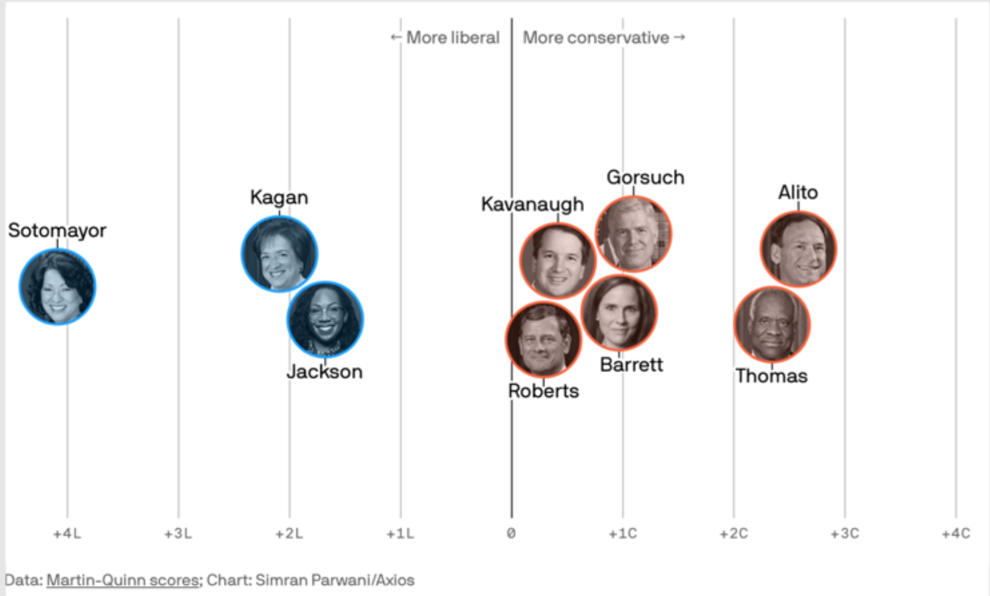

Axios

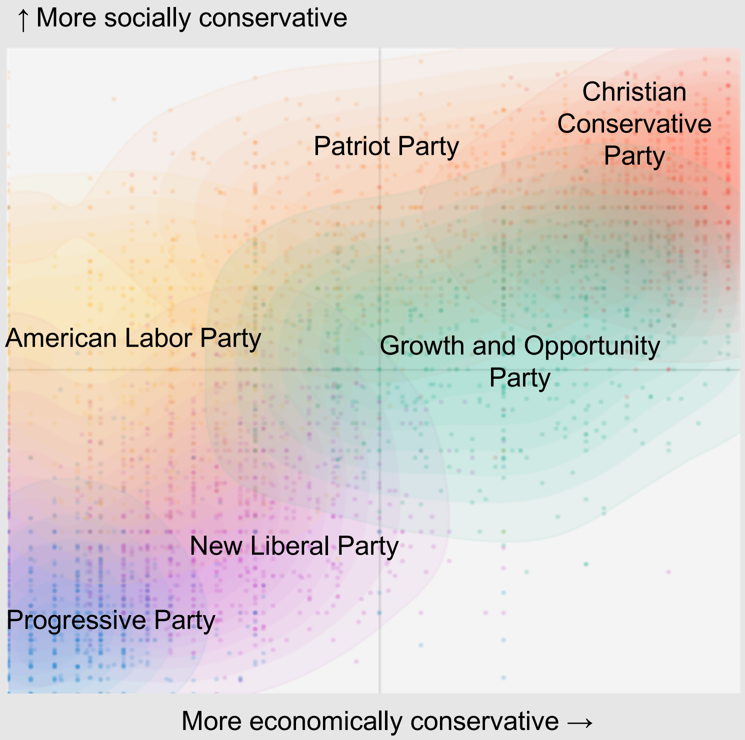

The New York Times

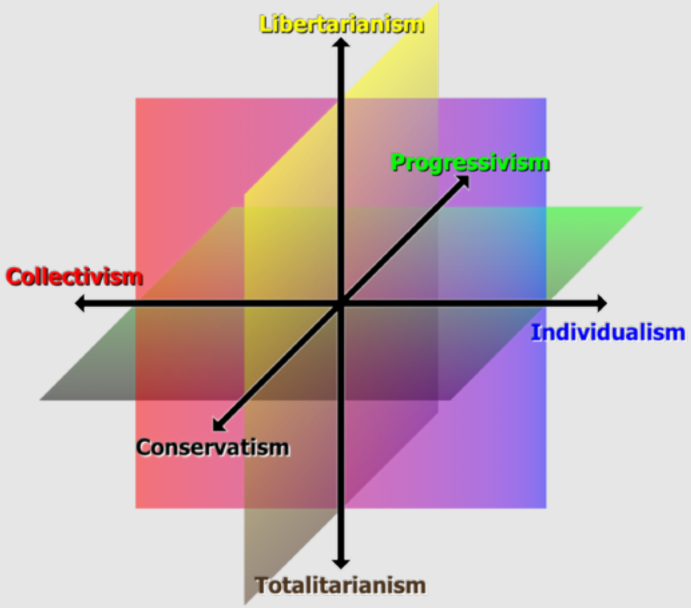

Wikipedia

The "Political Spectrum"

Alternative: **quantitative approach**, leverage ***spatial structure***

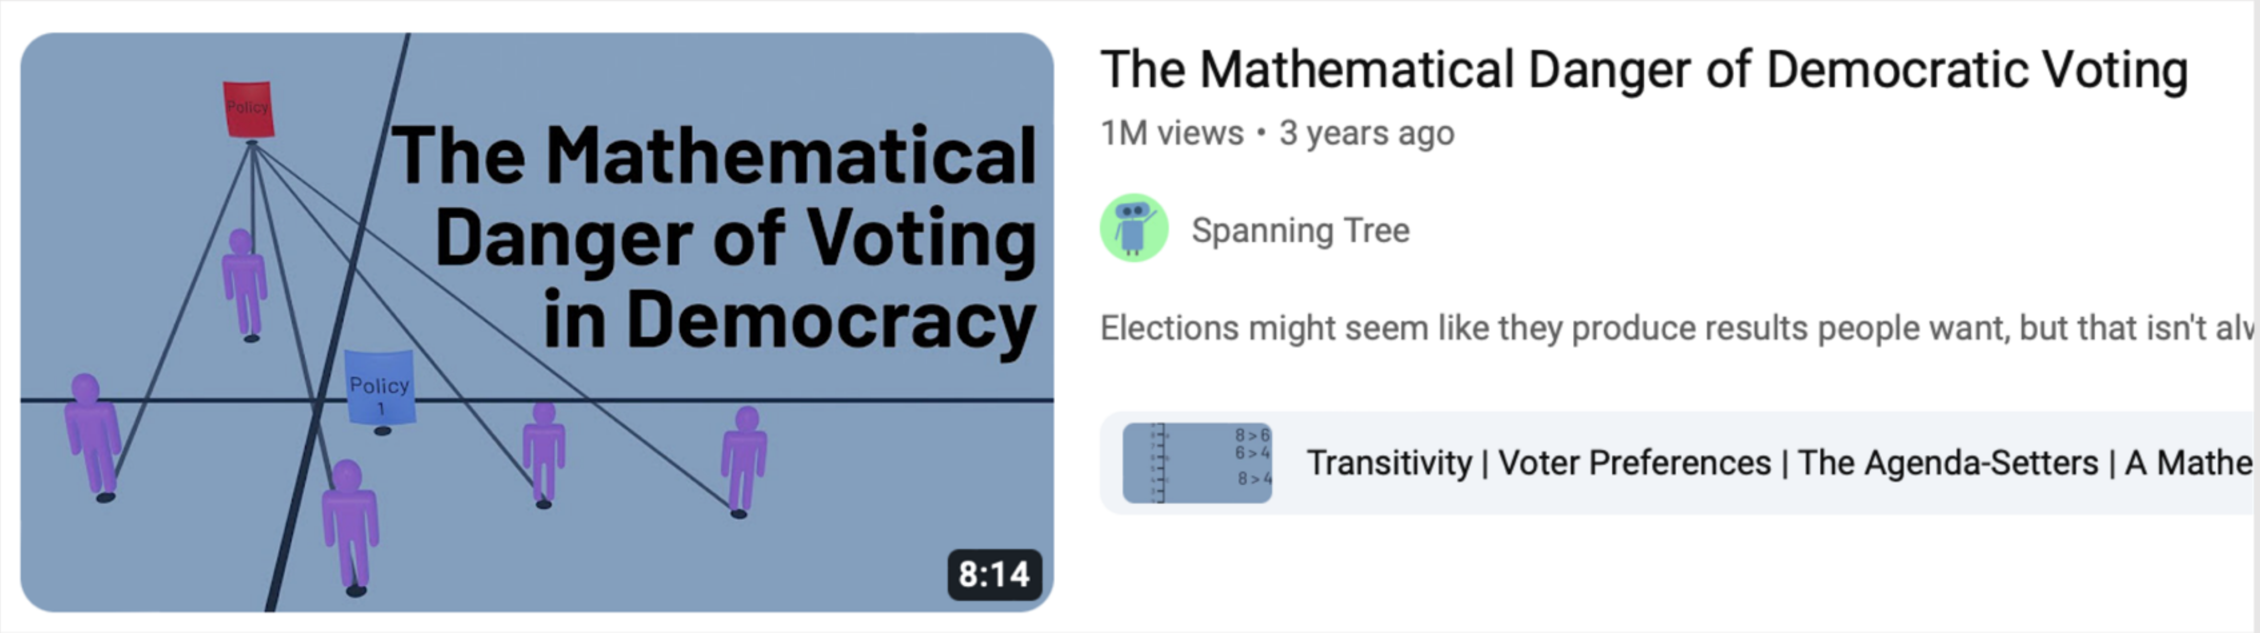

(McKelvey–Schofield chaos theorem)

Metric Distortion

- Voters and candidates lie in a metric space

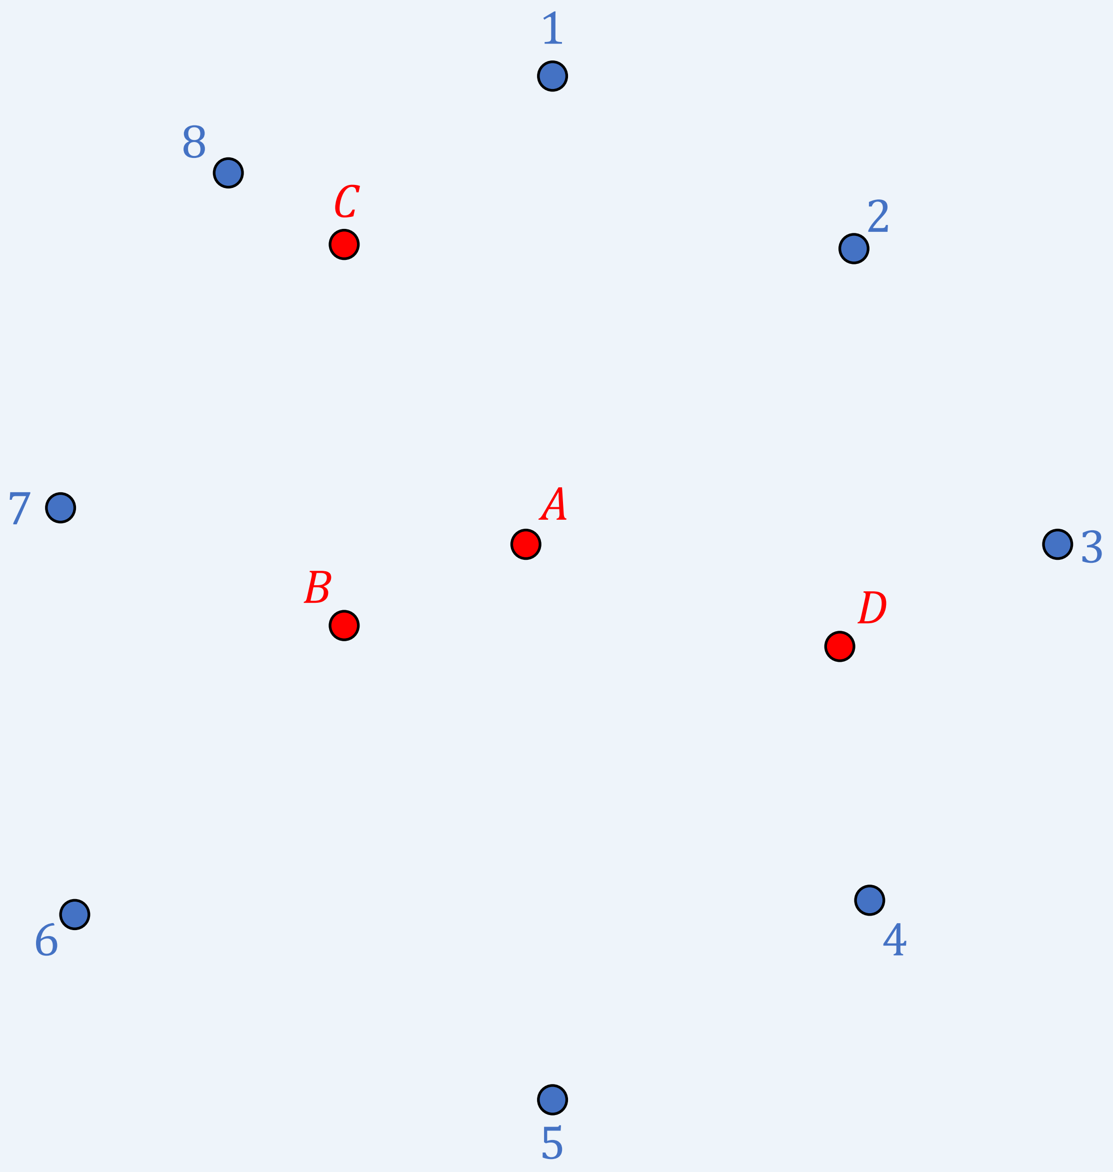

Metric Distortion

- Voters and candidates lie in a metric space
- Voter’s cost of candidate: distance
- Goal: minimize total cost

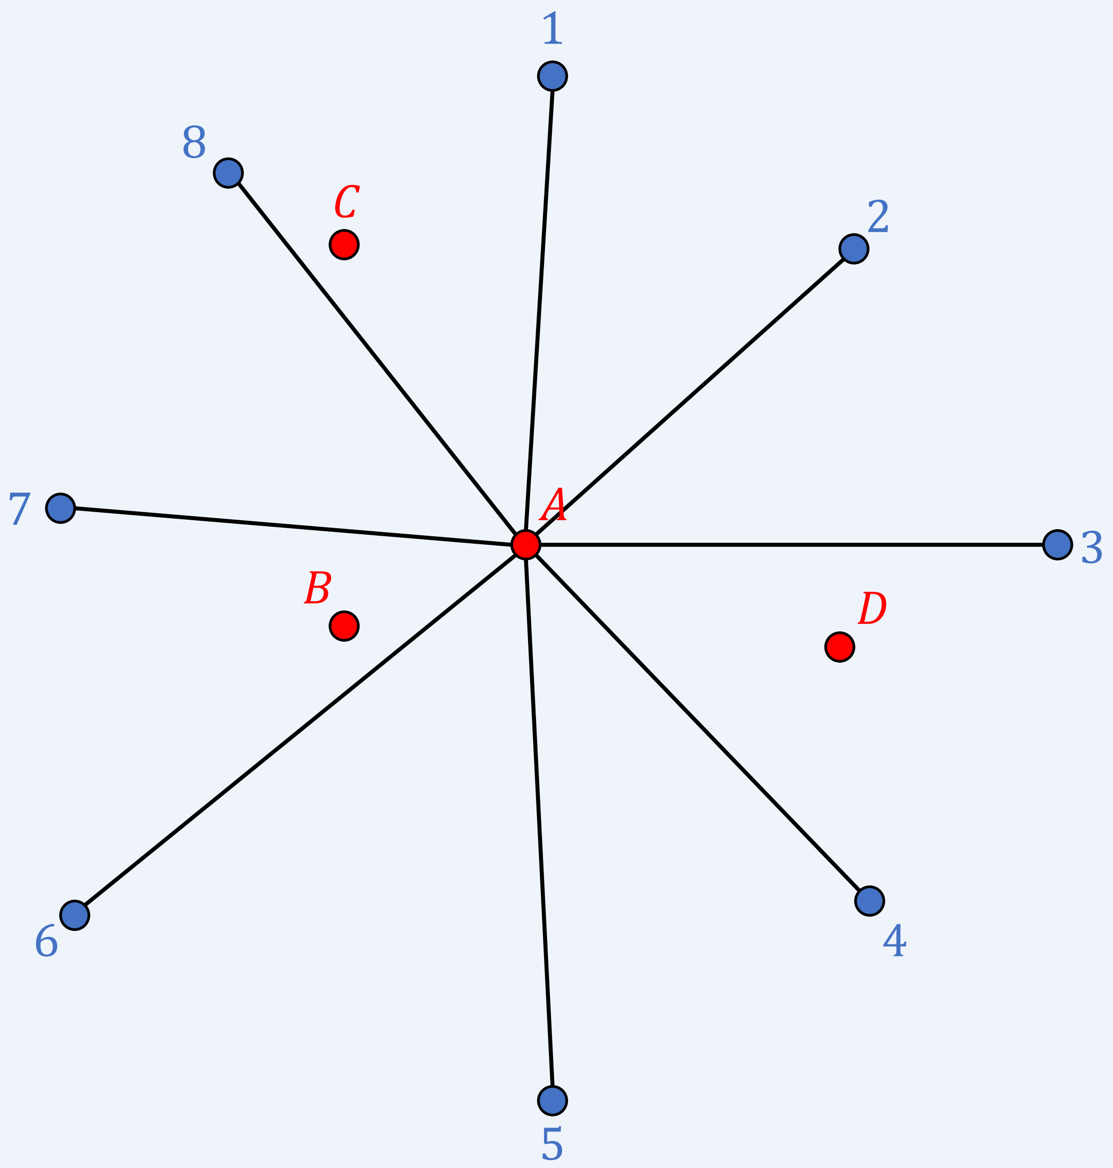

Good candidate

Metric Distortion

- Voters and candidates lie in a metric space
- Voter's cost of candidate: distance
- Goal: minimize total cost

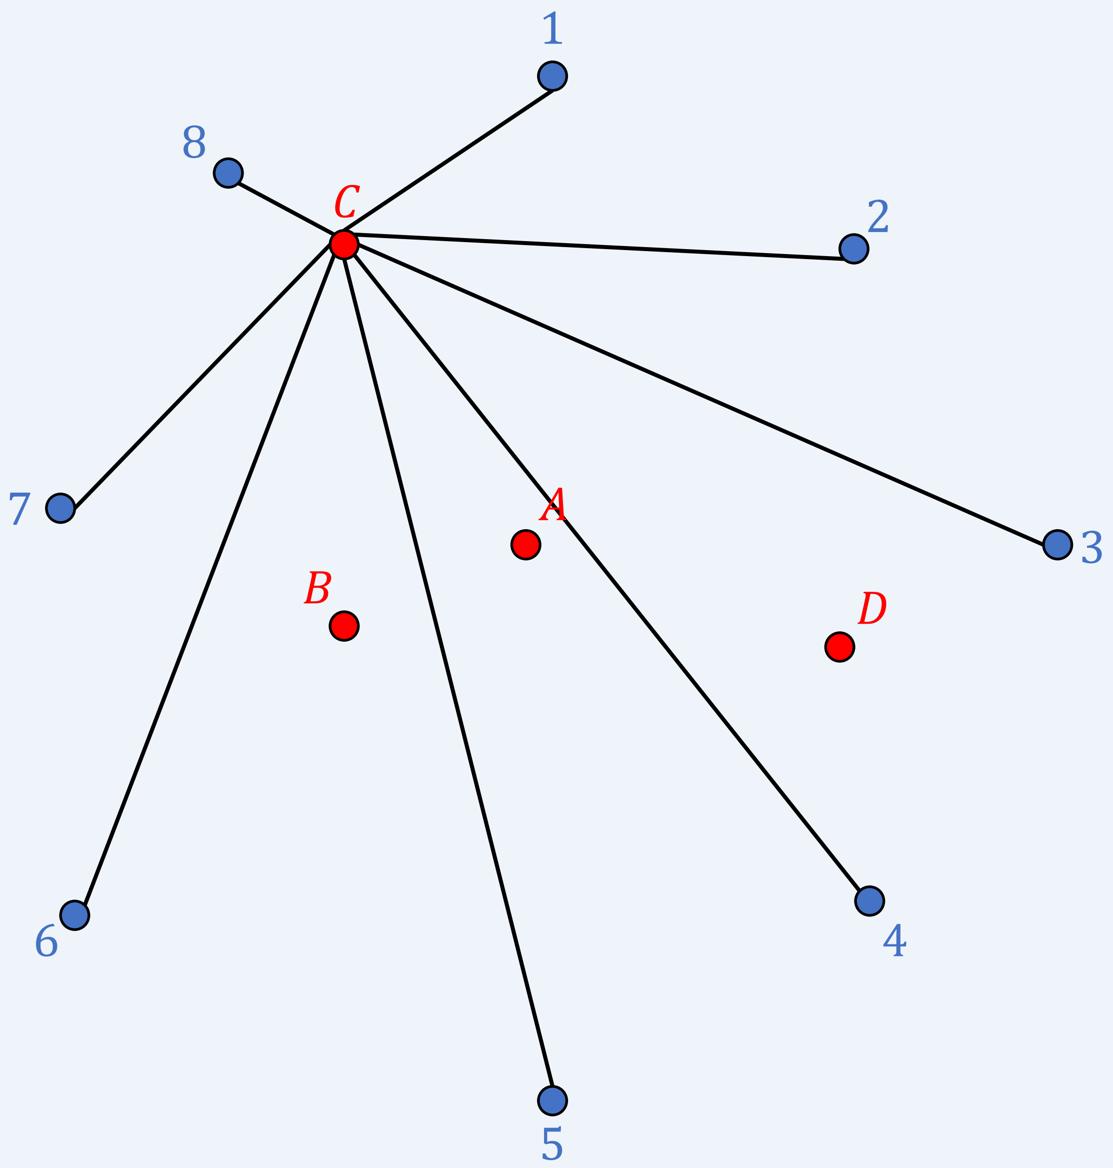

Bad candidate

Metric Distortion

- Voters and candidates lie in a metric space
- Voter's cost of candidate: distance
- Goal: minimize total cost
- Catch: don't know metric space, only have voters' **ranking** of candidates **by distance**

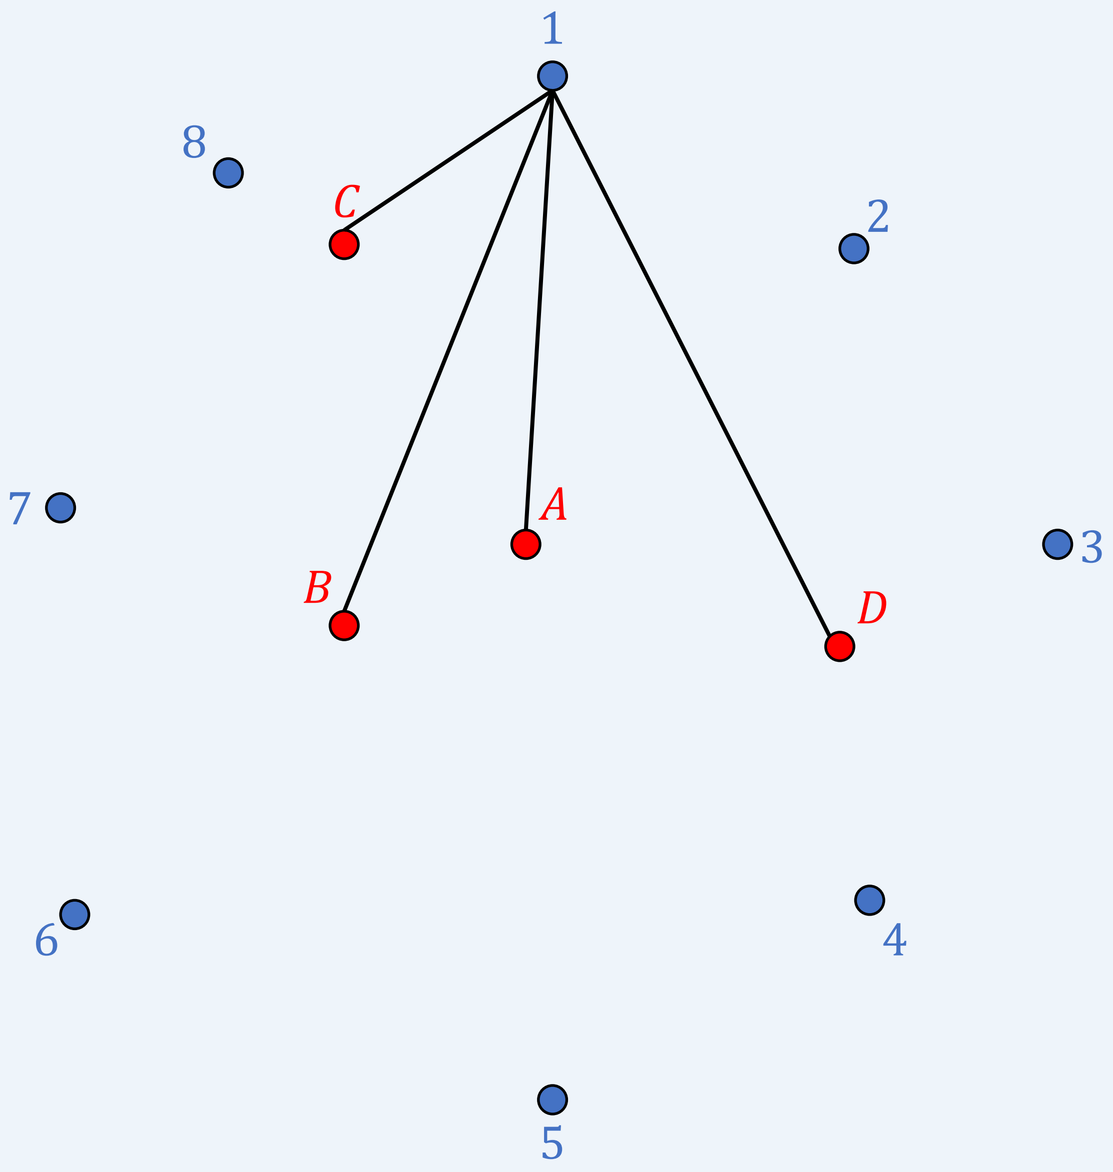

Metric Distortion

- Voters and candidates lie in a metric space
- Voter's cost of candidate: distance
- Goal: minimize total cost
- Catch: don't know metric space, only have voters' **ranking** of candidates **by distance**
- Can we find a good candidate? Cost within small ***factor*** of true optimum

"distortion"

Input:

Voter	Ranking
1	$CABD$
2	$DACB$
3	$DABC$
4	$DABC$
5	$BDAC$
6	$BACD$
7	$BCAD$
8	$CBAD$

Output: B

Problem Summary

- Input: voters' ranking of candidates by distance
- Output: single candidate
- Cost of candidate: total distance to voters
- Goal: regardless of underlying metric space, cost of chosen candidate only small factor worse than true OPT

Can we design *effective* **voting rules**?

→ **Low distortion**

Input:

Voter	Ranking
1	$CABD$
2	$DACB$
3	$DABC$
4	$DABC$
5	$BDAC$
6	$BACD$
7	$BCAD$
8	$CBAD$

Output: B

Easy lower bound

- Two candidates, two disagreeing voters:

Voter	Ranking
1	AB
2	BA

If rule picks A…

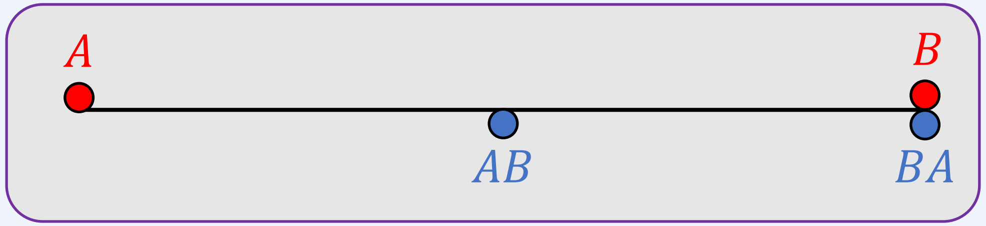

$\rightarrow cost(A) = 3 \cdot cost(B)$

If rule picks B…

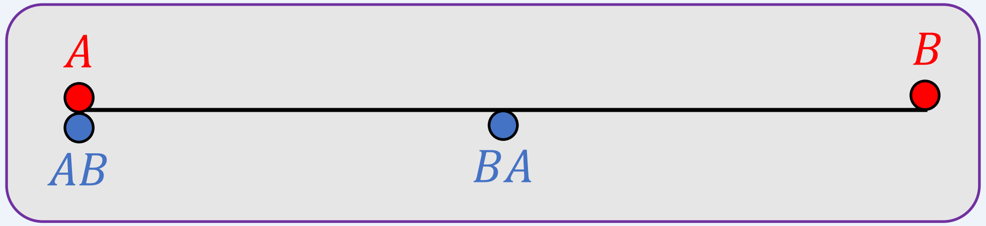

$\rightarrow cost(B) = 3 \cdot cost(A)$

- All deterministic rules: distortion $\geq \mathbf{3}$
- All randomized rules: distortion $\geq \mathbf{2}$

What’s the ***optimal*** distortion?

Deterministic Rules

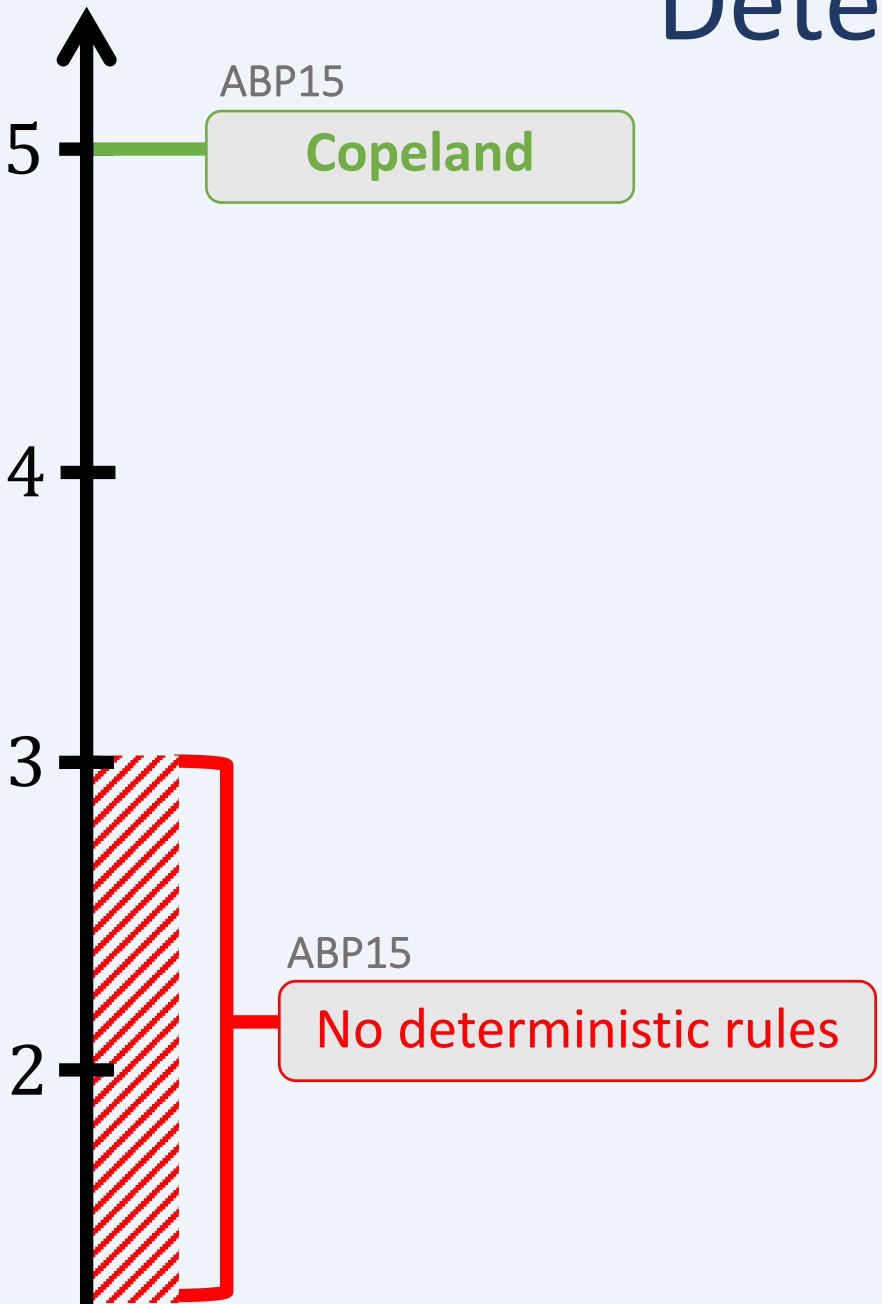

Anshelevich–Bhardwaj–Postl 2015

Optimal deterministic distortion is ≥ 3

Copeland has distortion 5

- Winner of most **pairwise** majority votes
- **Key property**: winner **beats** or **beats-someone-who-beats** every other candidate

Ramon Llull

Copeland winner $\boldsymbol{A}$: for all $\boldsymbol{B}$

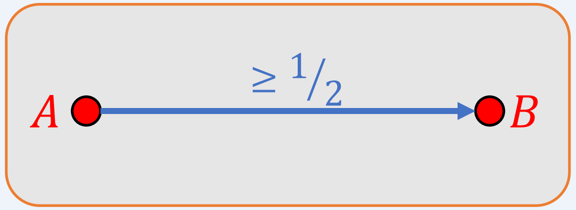

or

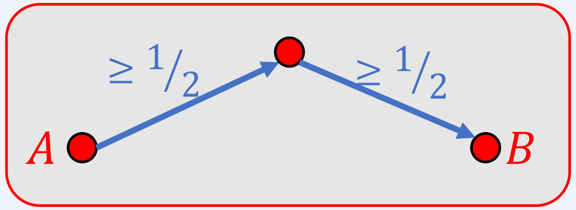

Deterministic Rules

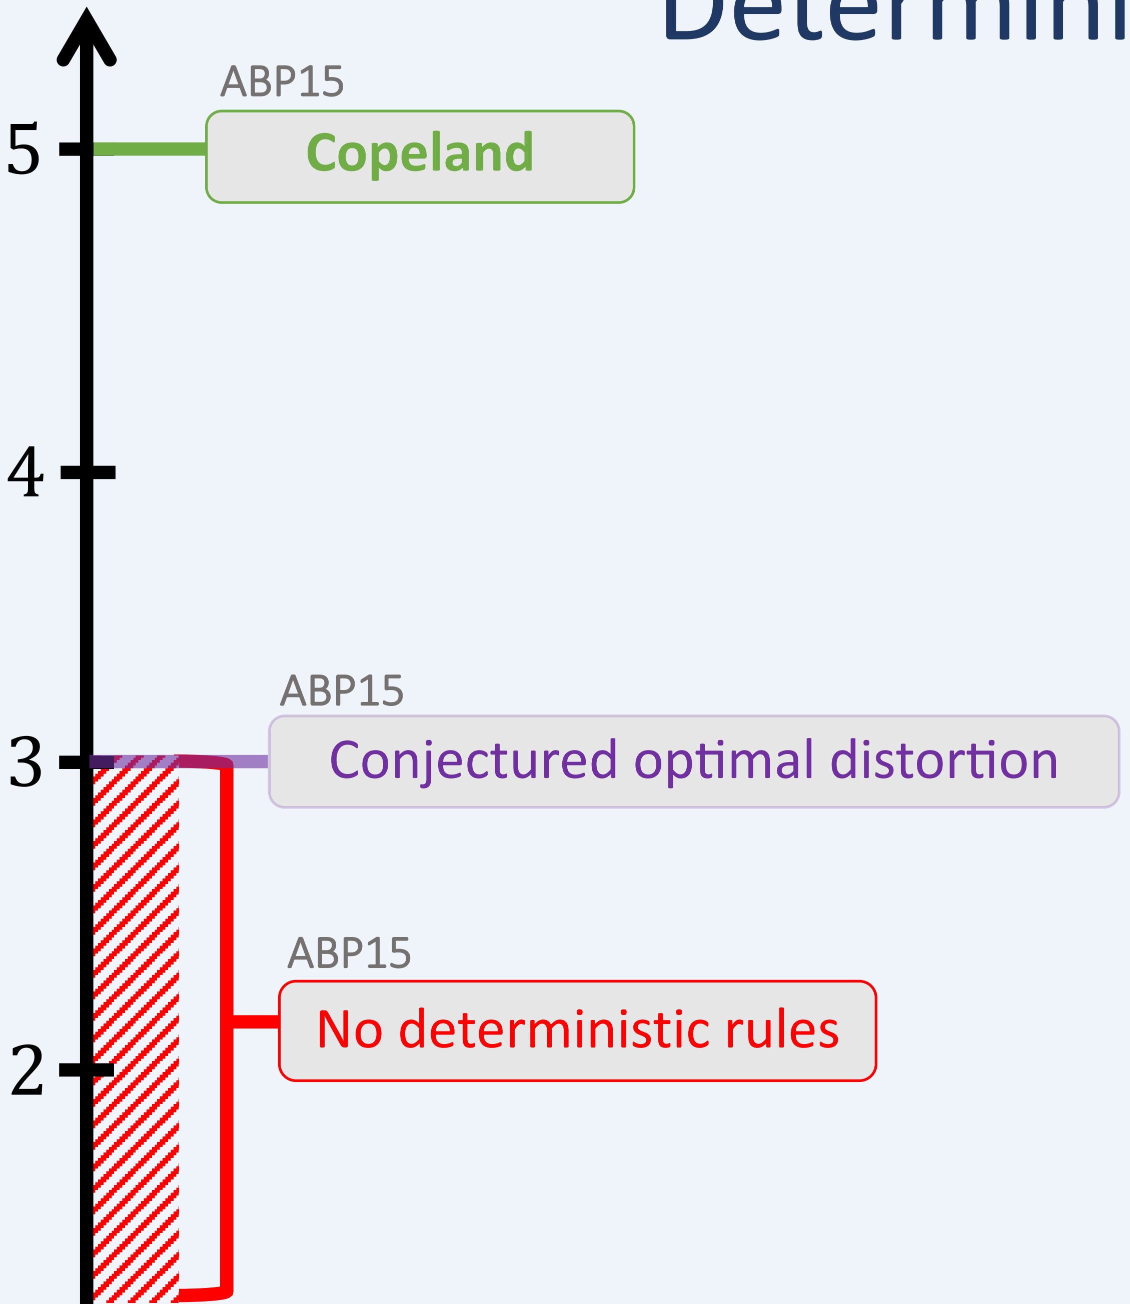

Anshelevich–Bhardwaj–Postl 2015

Optimal deterministic distortion is ≥ 3

Copeland has distortion 5

Conjecture: **Ranked Pairs** has distortion 3

Deterministic Rules

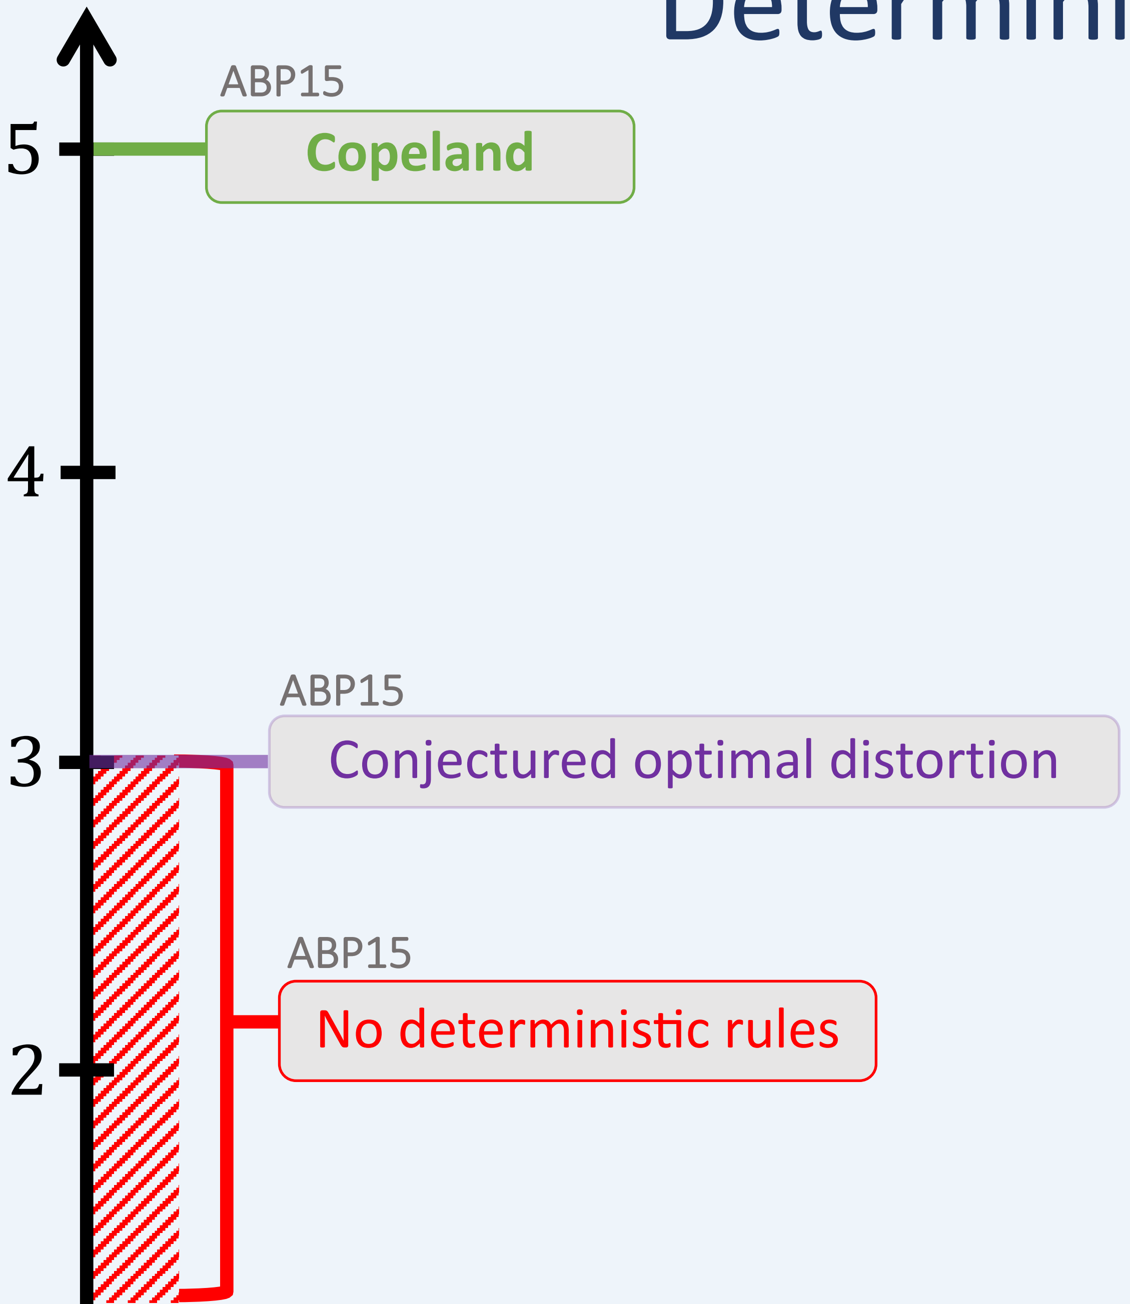

Goel–Krishnaswamy–Munagala 2017

Ranked Pairs has distortion ≥ 5

Conjecture: opt deterministic distortion is 3

Deterministic Rules

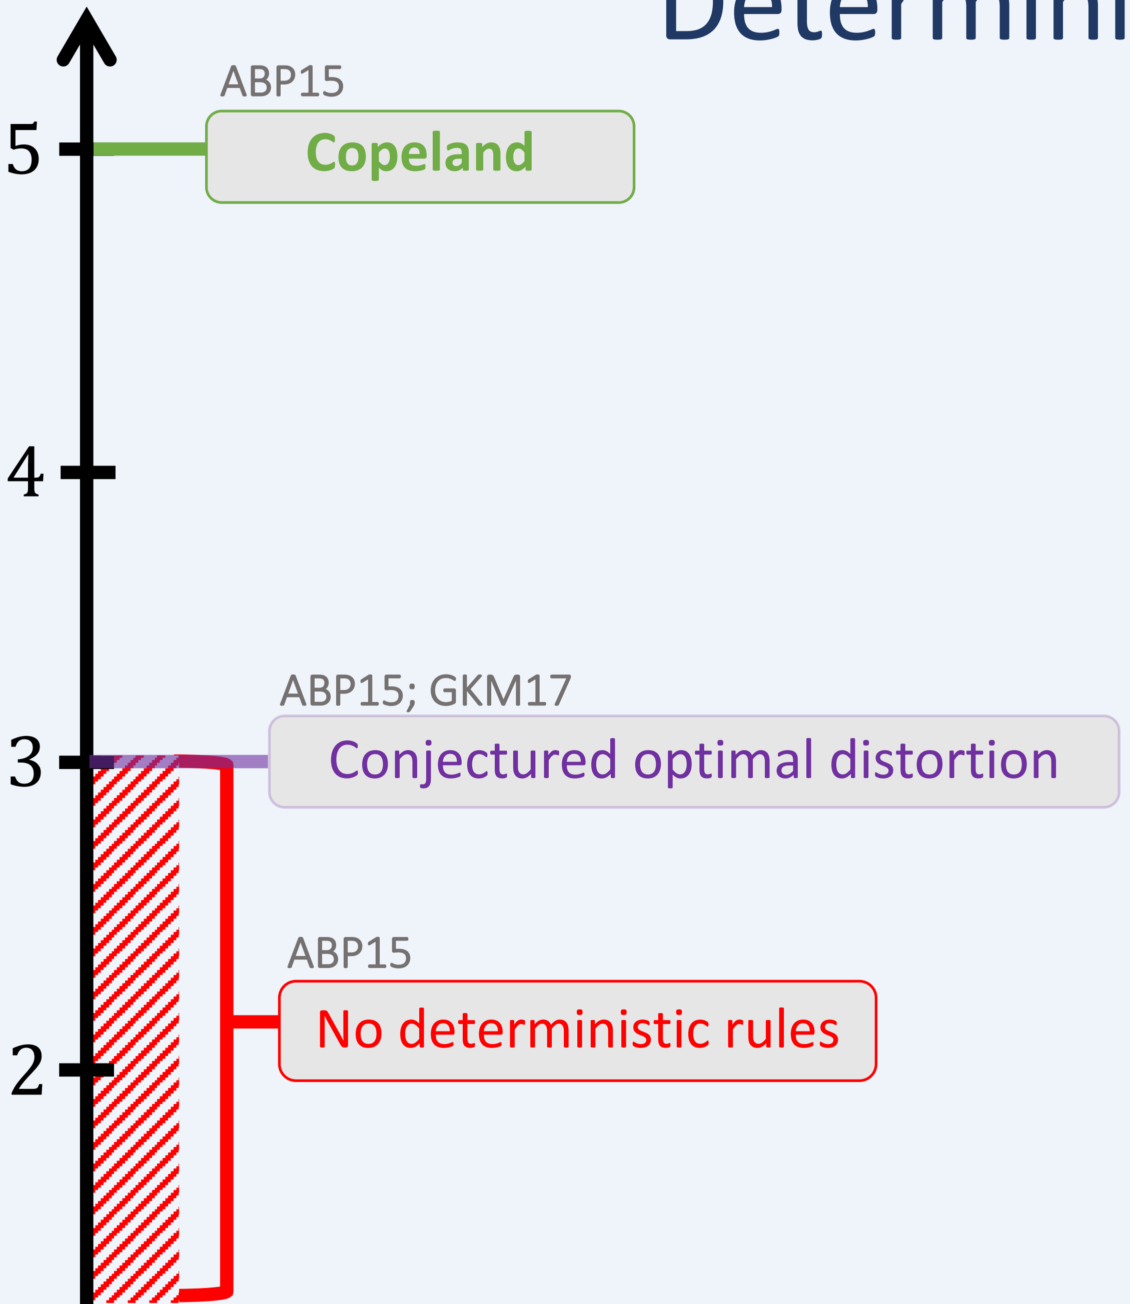

Goel–Krishnaswamy–Munagala 2017

Ranked Pairs has distortion ≥ 5

Conjecture: opt deterministic distortion is 3

Conjecture: opt *randomized* distortion is 2

Deterministic Rules

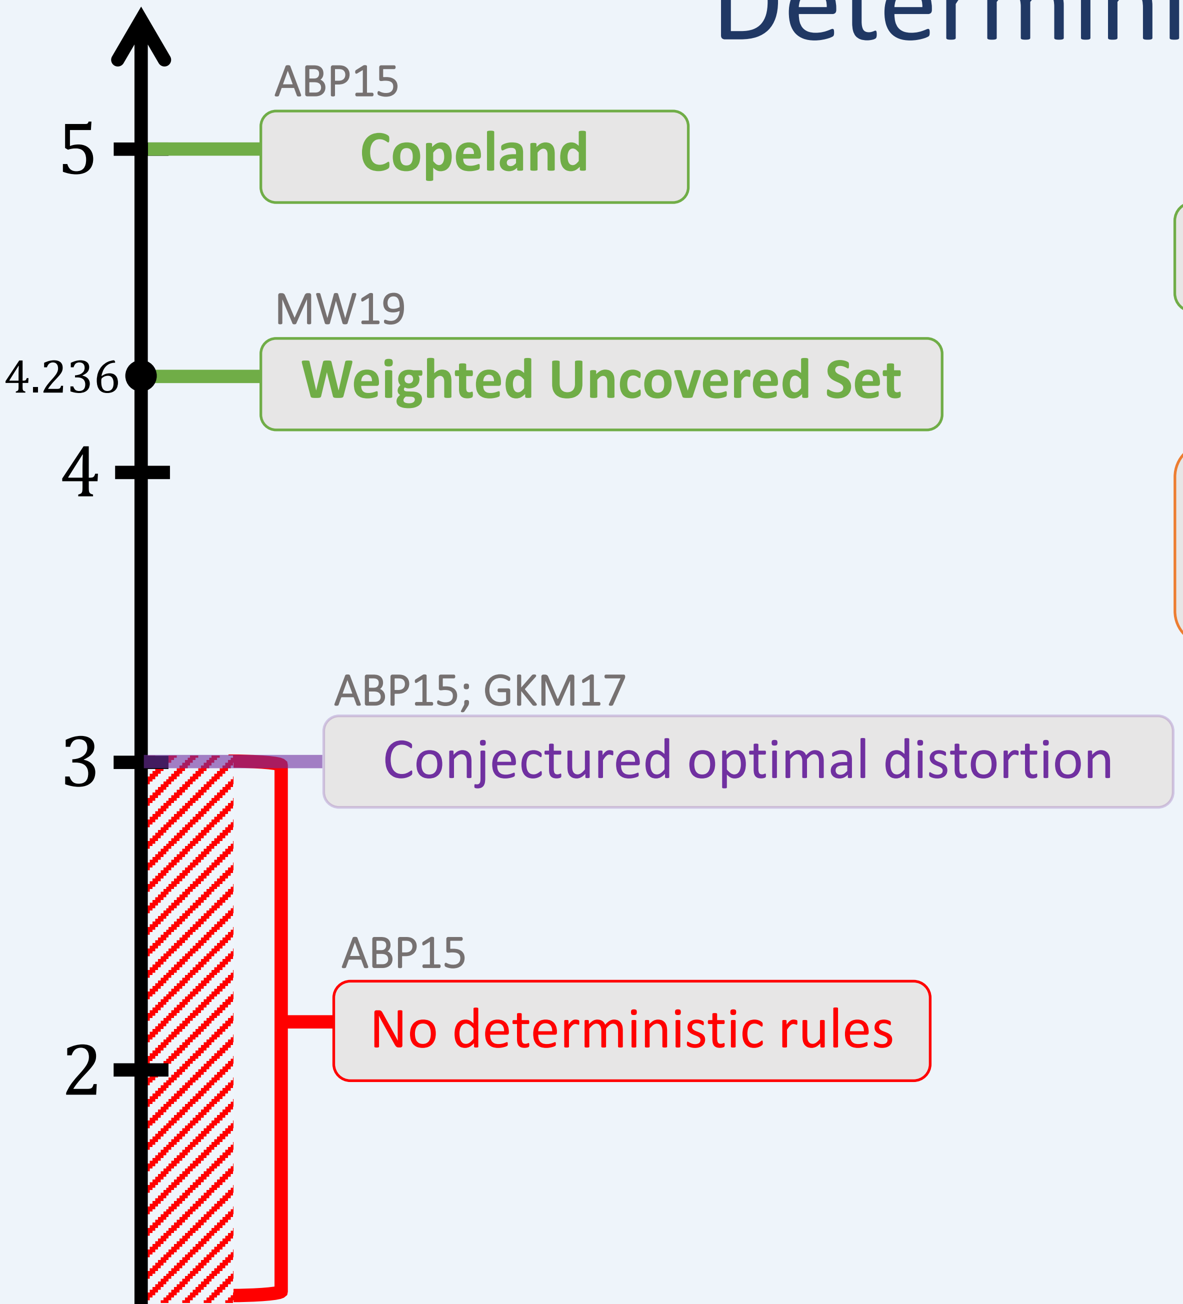

Munagala–Wang 2019

Novel Rule with distortion $\leq 2 + \sqrt{5} \approx 4.236$

Copeland winner $\boldsymbol{A}$: for all $\boldsymbol{B}$

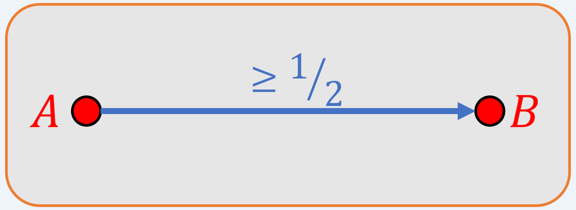

$\Rightarrow cost(A) \leq 3 \cdot cost(B)$

or

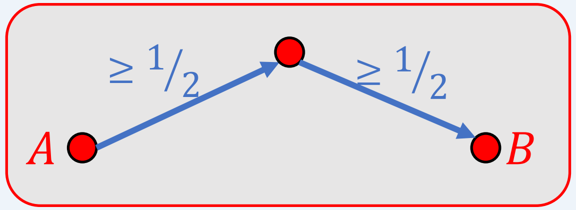

$\Rightarrow cost(A) \leq 5 \cdot cost(B)$

Idea: weaken left, strengthen right

Deterministic Rules

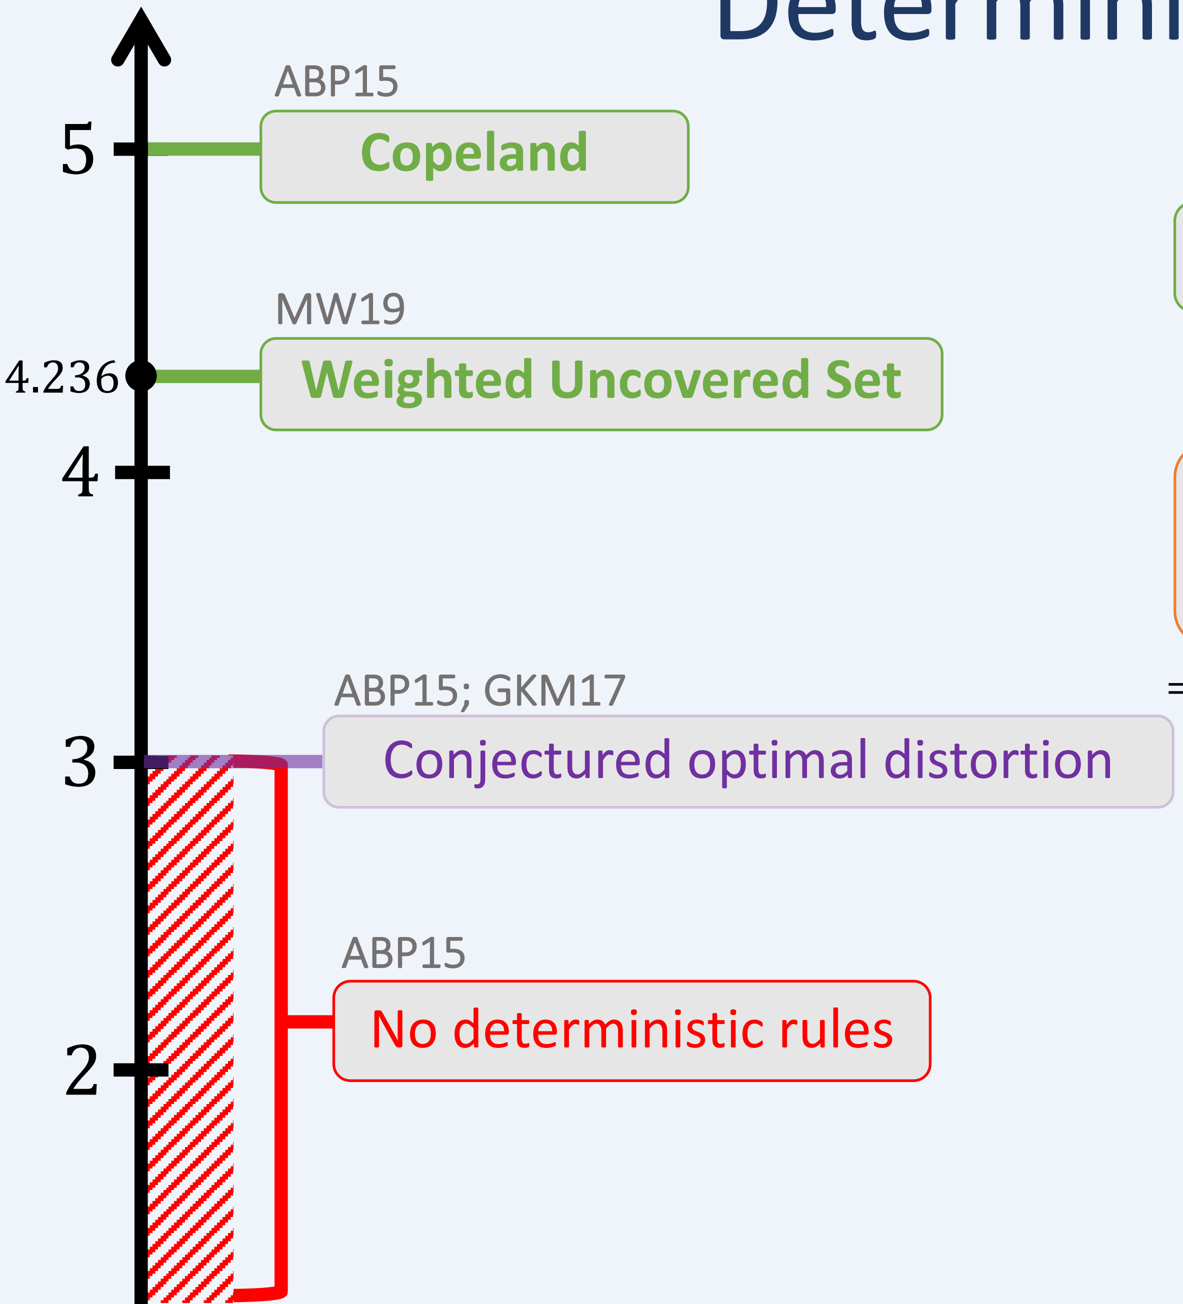

Munagala–Wang 2019

Novel Rule with distortion $\leq 2+\sqrt{5} \approx 4.236$

$\beta \in [\frac{1}{2}, 1]$. Exists? ***A***: for all ***B***

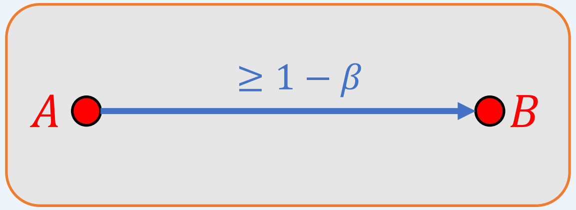

or

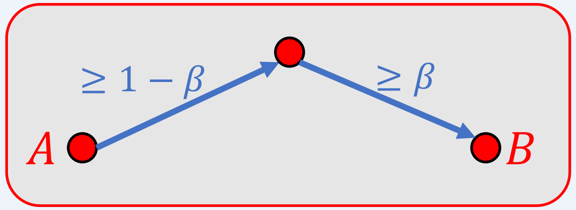

$\Rightarrow cost(A) \leq (1 + 2\frac{\beta}{1-\beta}) \cdot cost(B)$

$\Rightarrow^* cost(A) \leq (1 + \frac{2}{\beta}) \cdot cost(B)$

Exists for all β! Best choice $\beta = \varphi^{-1} \approx 0.618$

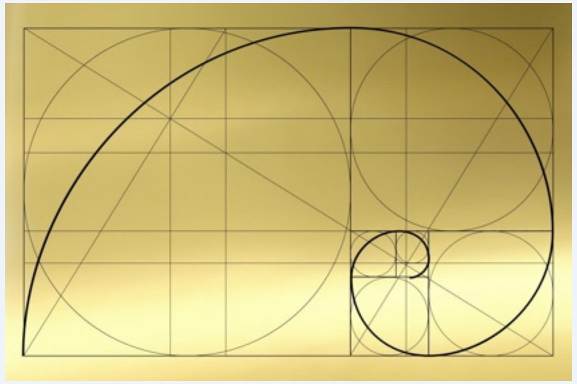

Deterministic Rules

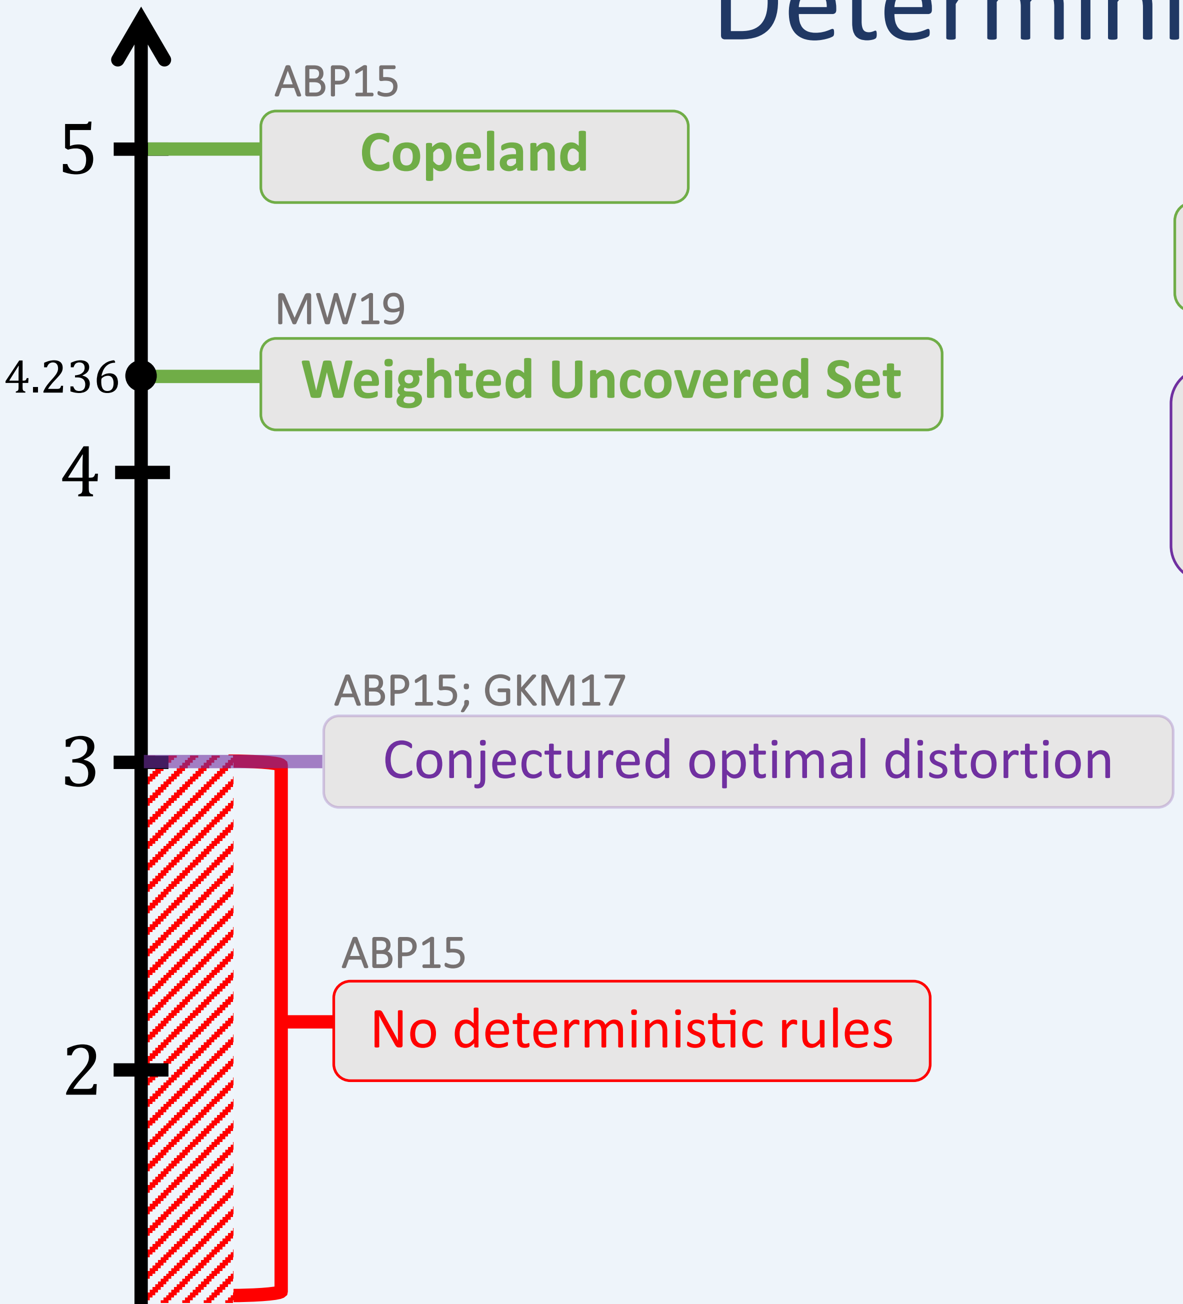

Munagala–Wang 2019

Novel Rule with distortion $\leq 2 + \sqrt{5} \approx 4.236$

Formulated **combinatorial conjecture,** implying distortion 3

Deterministic Rules

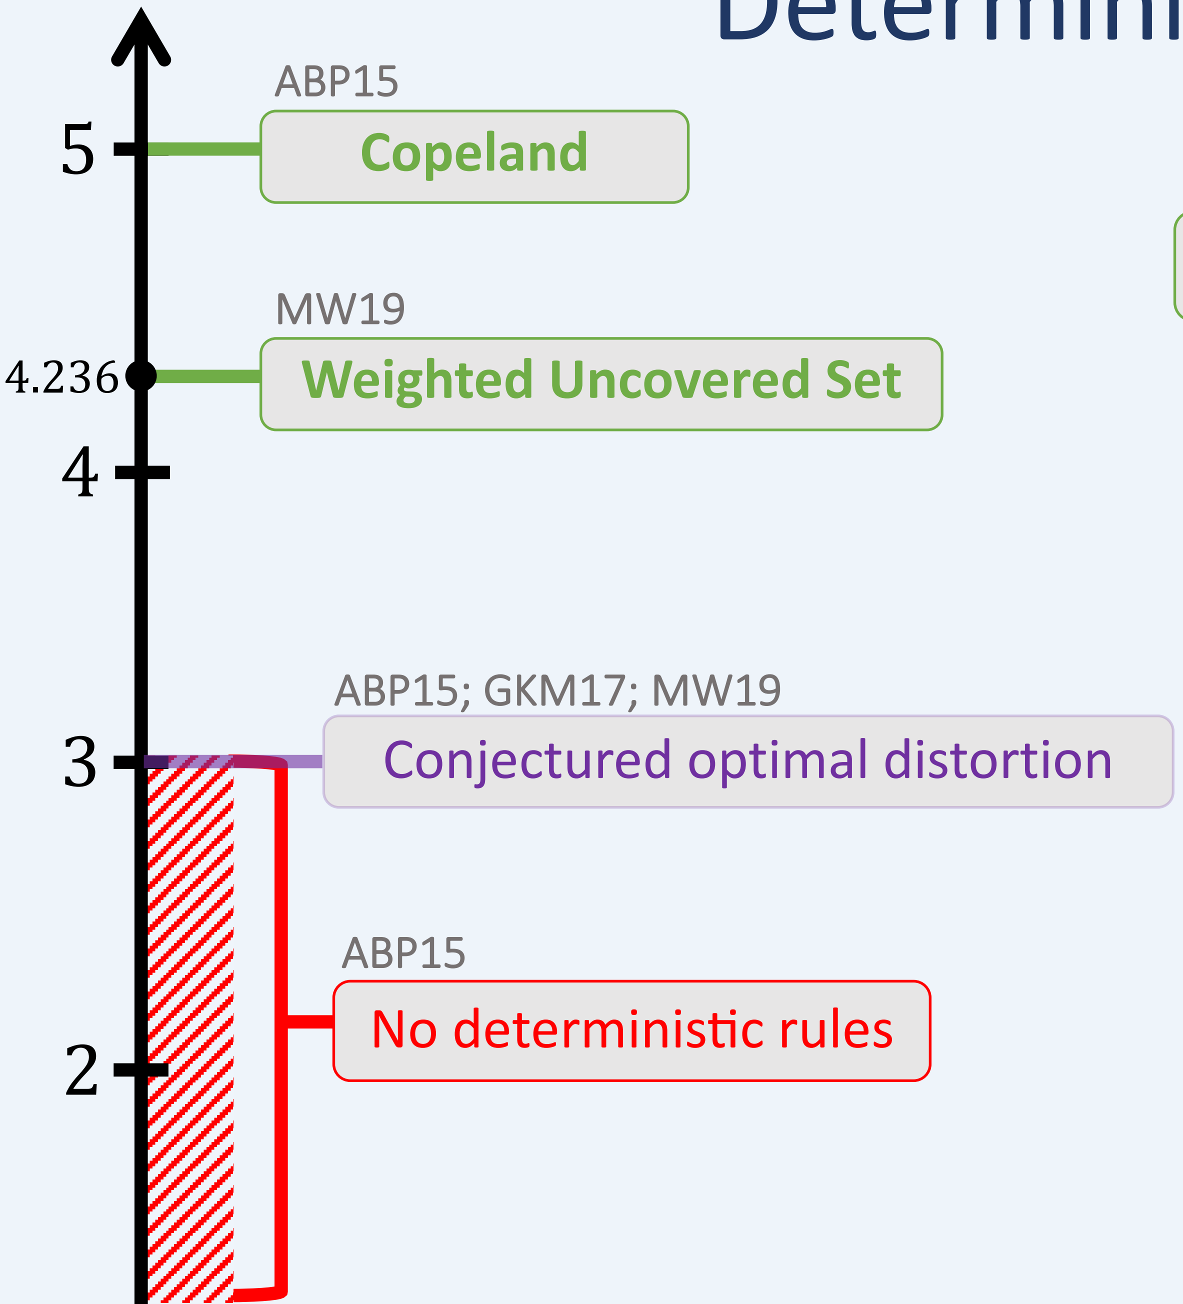

Gkatzelis–Halpern–Shah 2020

Proved MW19’s combinatorial conjecture 🎉

Deterministic Rules

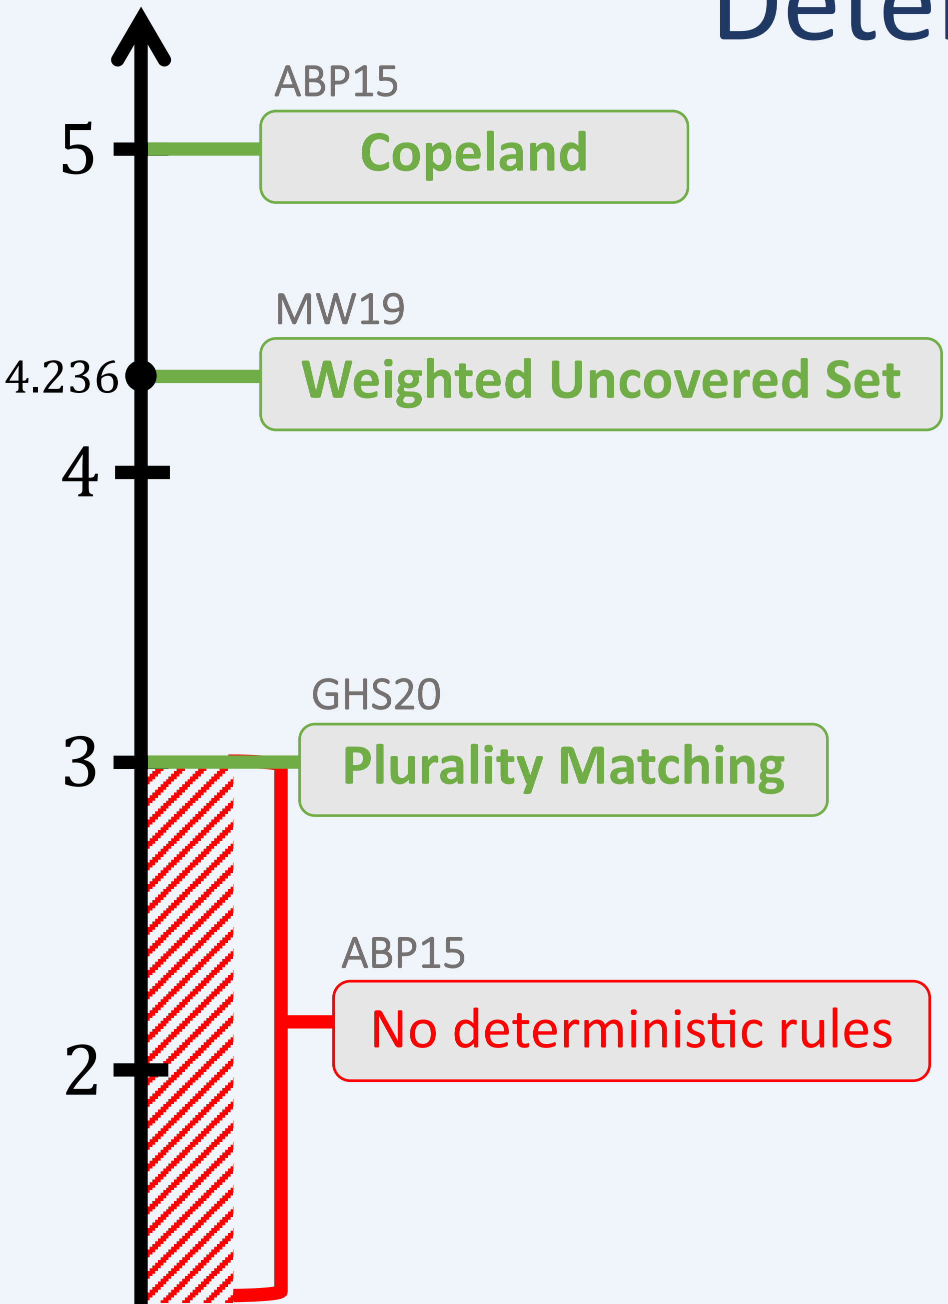

Gkatzelis–Halpern–Shah 2020

Proved MW19’s combinatorial conjecture 🎉

Deterministic Rules

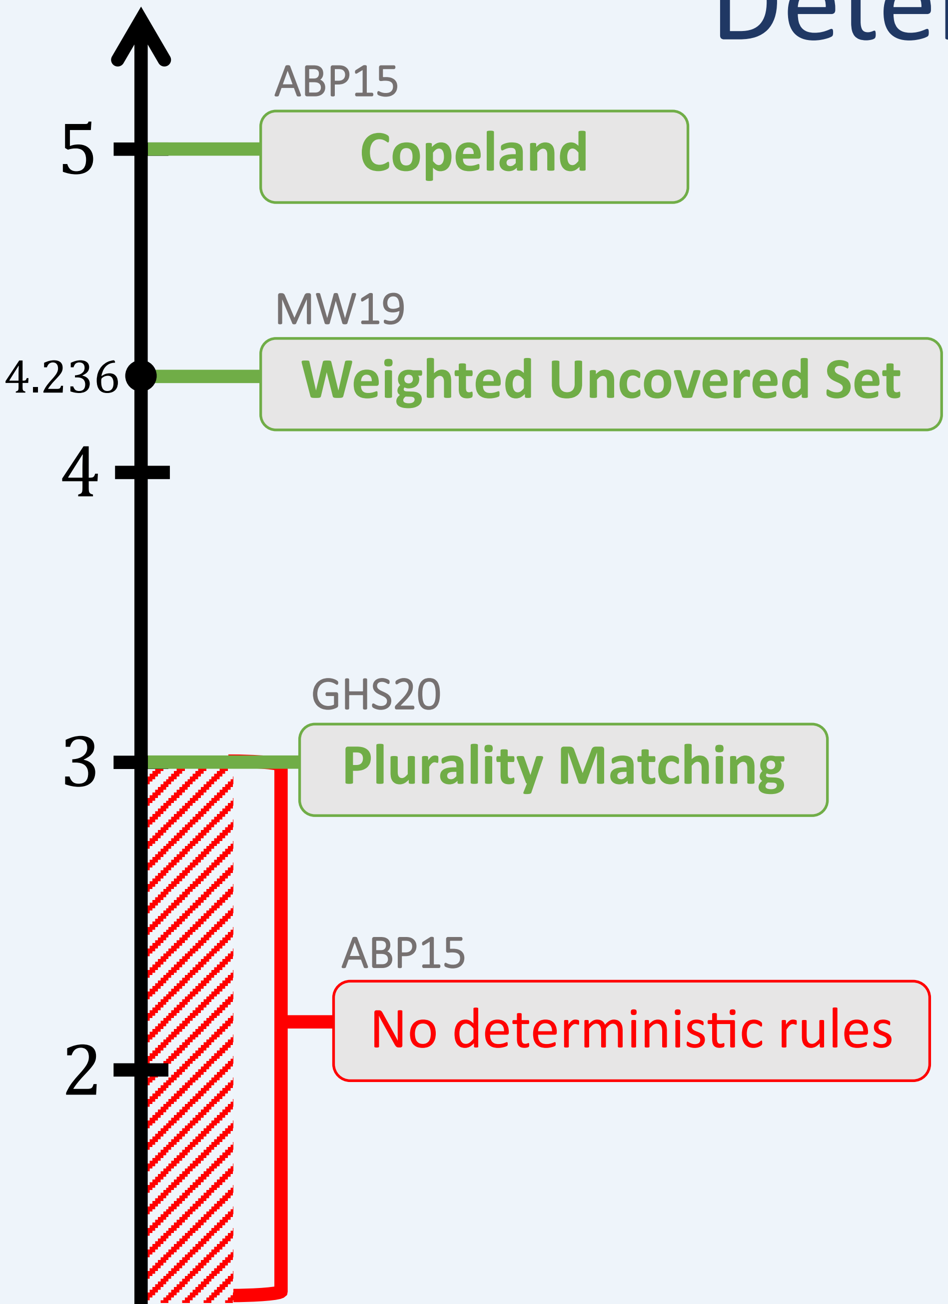

Kizilkaya–Kempe 2022

Plurality Veto: elegant **novel rule** with short proof of optimal distortion 🎉

Deterministic Rules

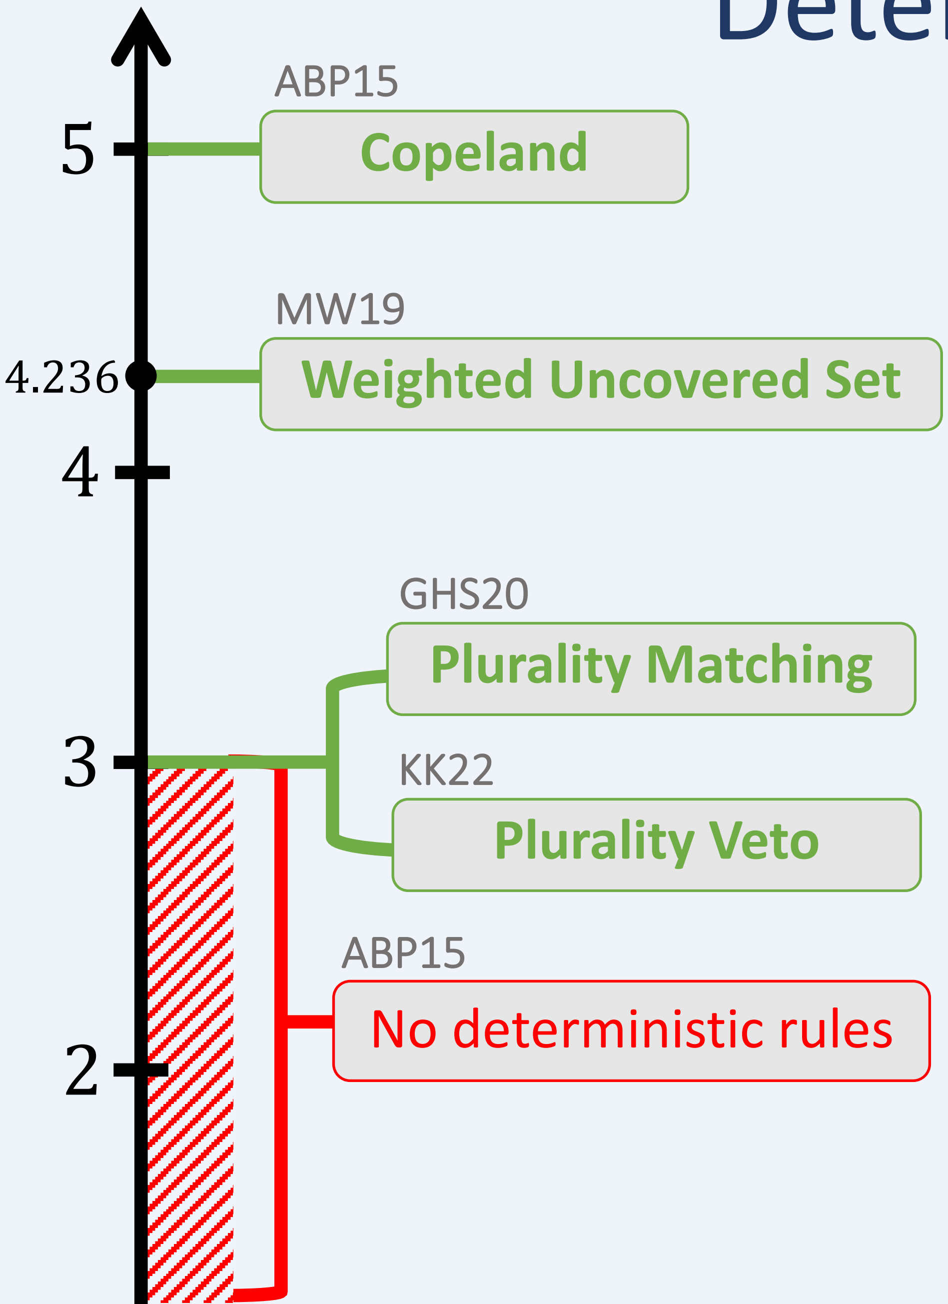

Kizilkaya–Kempe 2022

Plurality Veto: elegant **novel rule** with short proof of optimal distortion 🎉

Plurality Veto

Kizilkaya–Kempe 2022

Plurality Veto

Kizilkaya–Kempe 2022

Everyone: place token on favorite game

Plurality Veto

Kizilkaya–Kempe 2022

One by one: remove token from least favorite game

Plurality Veto

Kizilkaya–Kempe 2022

Winner: last game with tokens

Plurality Veto

Kizilkaya–Kempe 2022

Let j_v be the candidate vetoed by voter v, and let j^* be the final chosen candidate. Let P_j be the set of voters that rank candidate j first and let $\mathrm{plu}(j) = |P_j|$. Since j^* has positive score until the very end, it must be the case that for each $v \in V$, $j^* \succeq_v j_v$. Then we have that for any candidate i,

$$\begin{aligned}
\sum_{v\in V} d(j^*, v) &\leq \sum_{v\in V} d(j_v, v) && (j^* \succeq_v j_v) \\
&\leq \sum_{v\in V} (d(i,v) + d(i, j_v)) && (\text{triangle inequality}) \\
&= \sum_{v\in V} d(i,v) + \sum_{j\in C} \mathrm{plu}(j) d(i,j) && (j \text{ is vetoed } \mathrm{plu}(j) \text{ times}) \\
&= \sum_{v\in V} d(i,v) + \sum_{j\in C} \sum_{v\in P_j} d(i,j) && \\
&\leq \sum_{v\in V} d(i,v) + \sum_{j\in C} \sum_{v\in P_j} (d(i,v) + d(j,v)) && (\text{triangle inequality}) \\
&\leq \sum_{v\in V} d(i,v) + \sum_{j\in C} \sum_{v\in P_j} 2d(i,v) && (v \in P_j \text{ means } j \succeq_v i) \\
&= 3 \sum_{v\in V} d(i,v)
\end{aligned}$$

"Optimal Metric Distortion for Voting – A Proof from the Book"

Stanford Theory Dish Blog

Formal Description & Distortion Proof

Plurality Veto

- Initially, each candidate X,

$\text{score}(X) = \text{ \# first choice votes for X}$

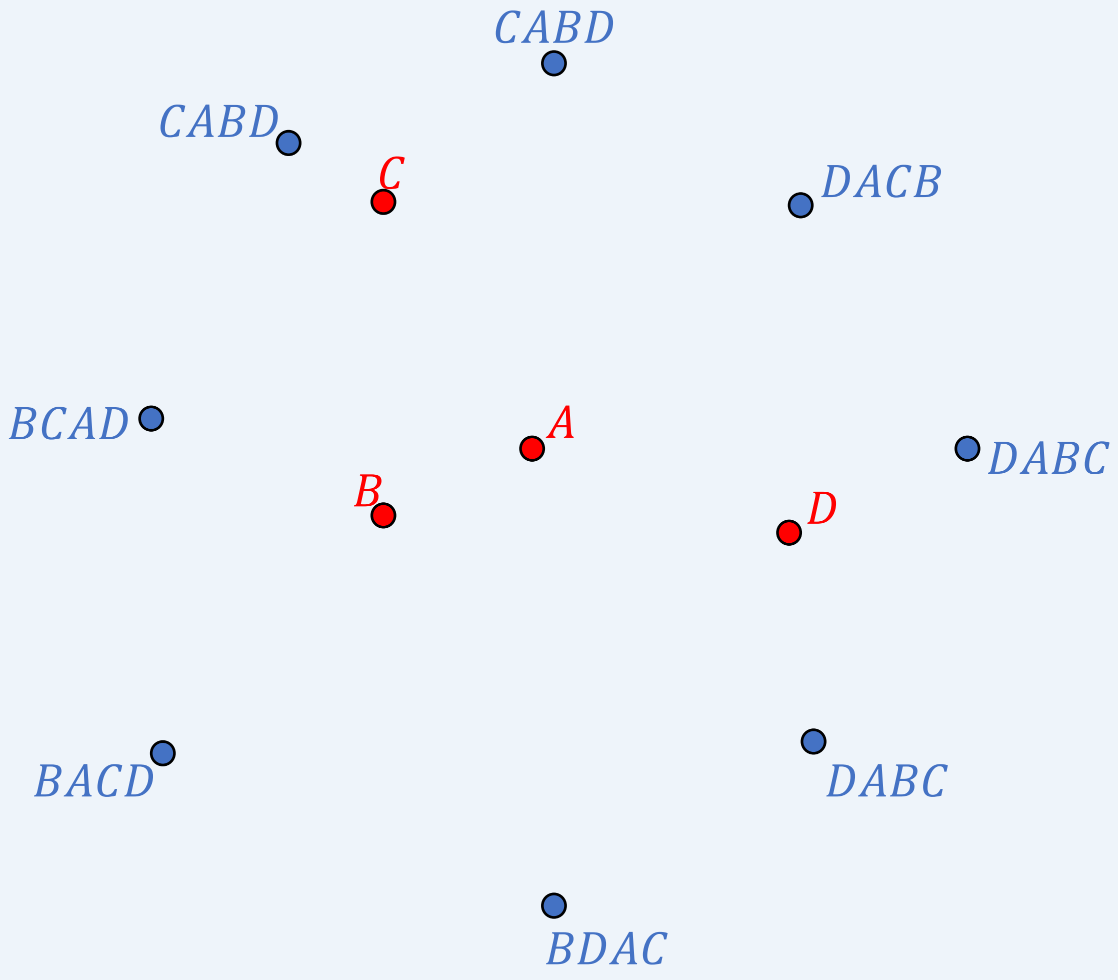

Plurality Veto

- Initially, each candidate X, $\text{score}(X) = \#\text{ first choice votes for X}$
- One by one, each voter decrements score (**veto**) of least favorite candidate with positive score
- Last candidate vetoed wins

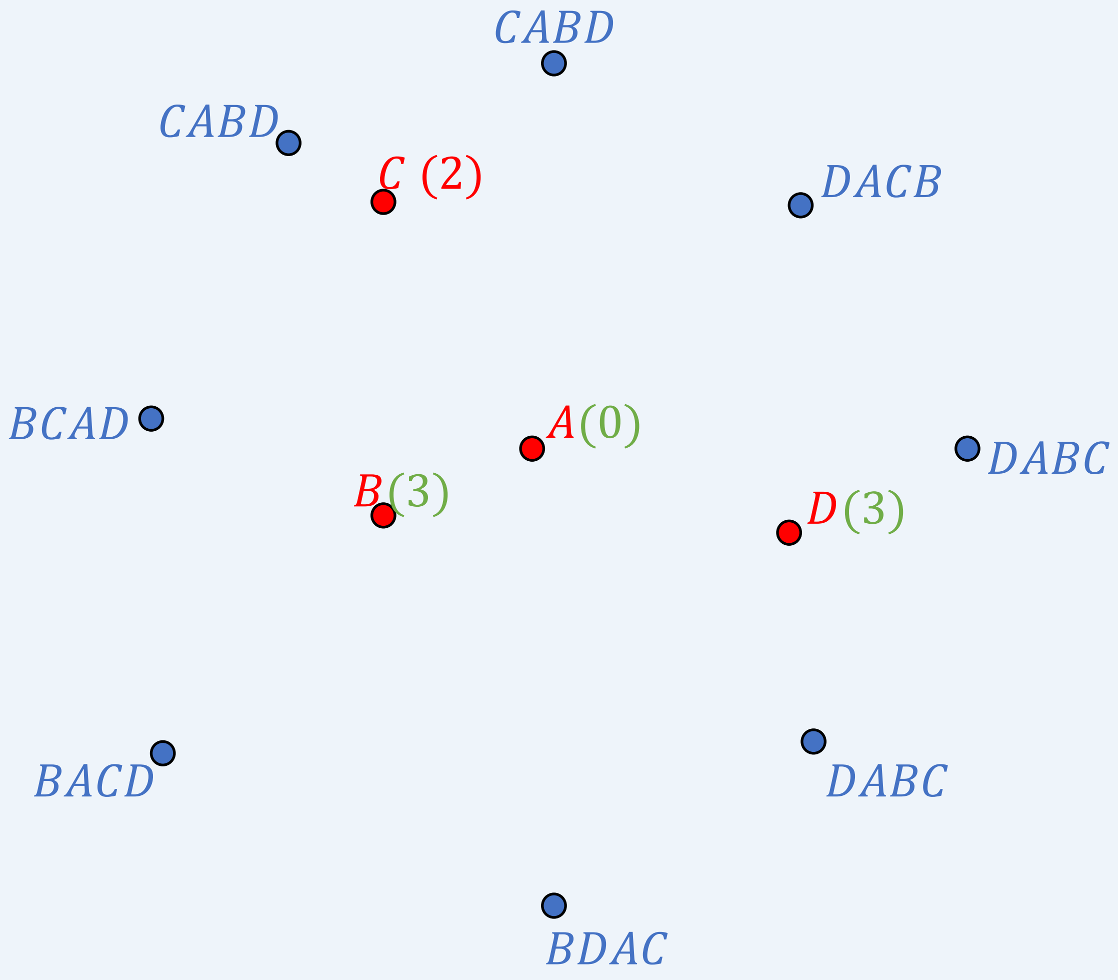

Plurality Veto

- Initially, each candidate X, $\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (**veto**) of least favorite candidate with positive score
- Last candidate vetoed wins

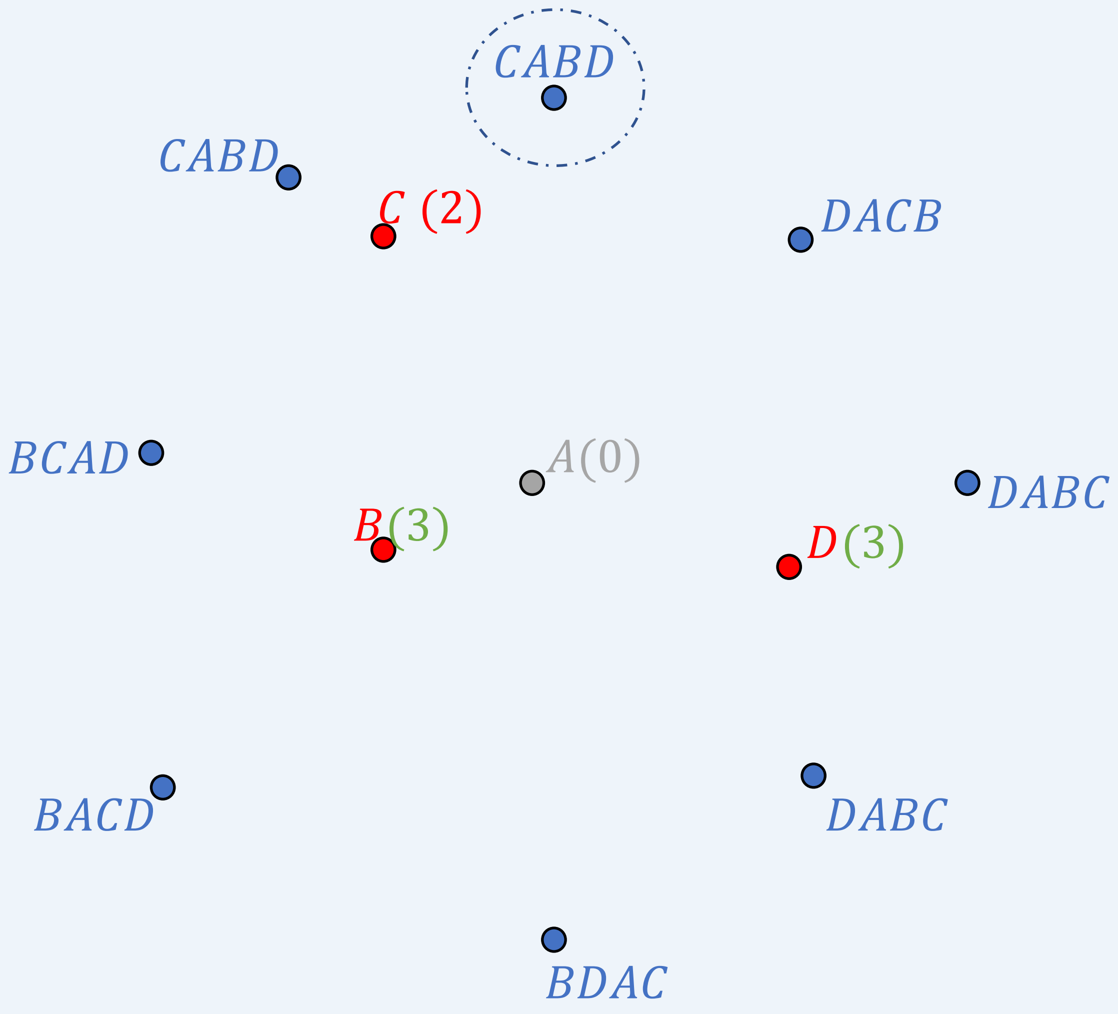

Plurality Veto

- Initially, each candidate X, $\text{score}(X) = \ \#\text{ first choice votes for X}$
- One by one, each voter decrements score (veto) of least favorite candidate with positive score
- Last candidate vetoed wins

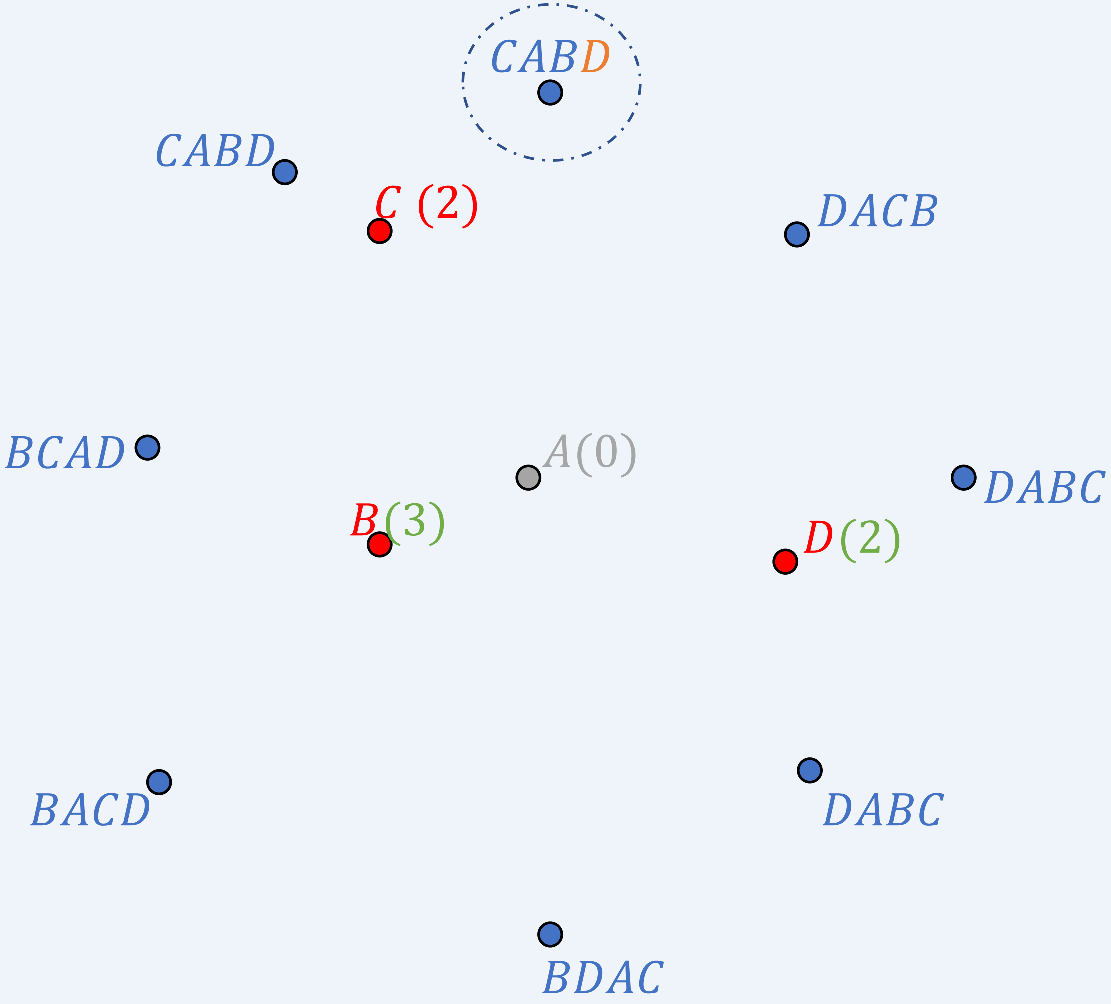

Plurality Veto

- Initially, each candidate X,
$\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (veto) of least favorite candidate with positive score
- Last candidate vetoed wins

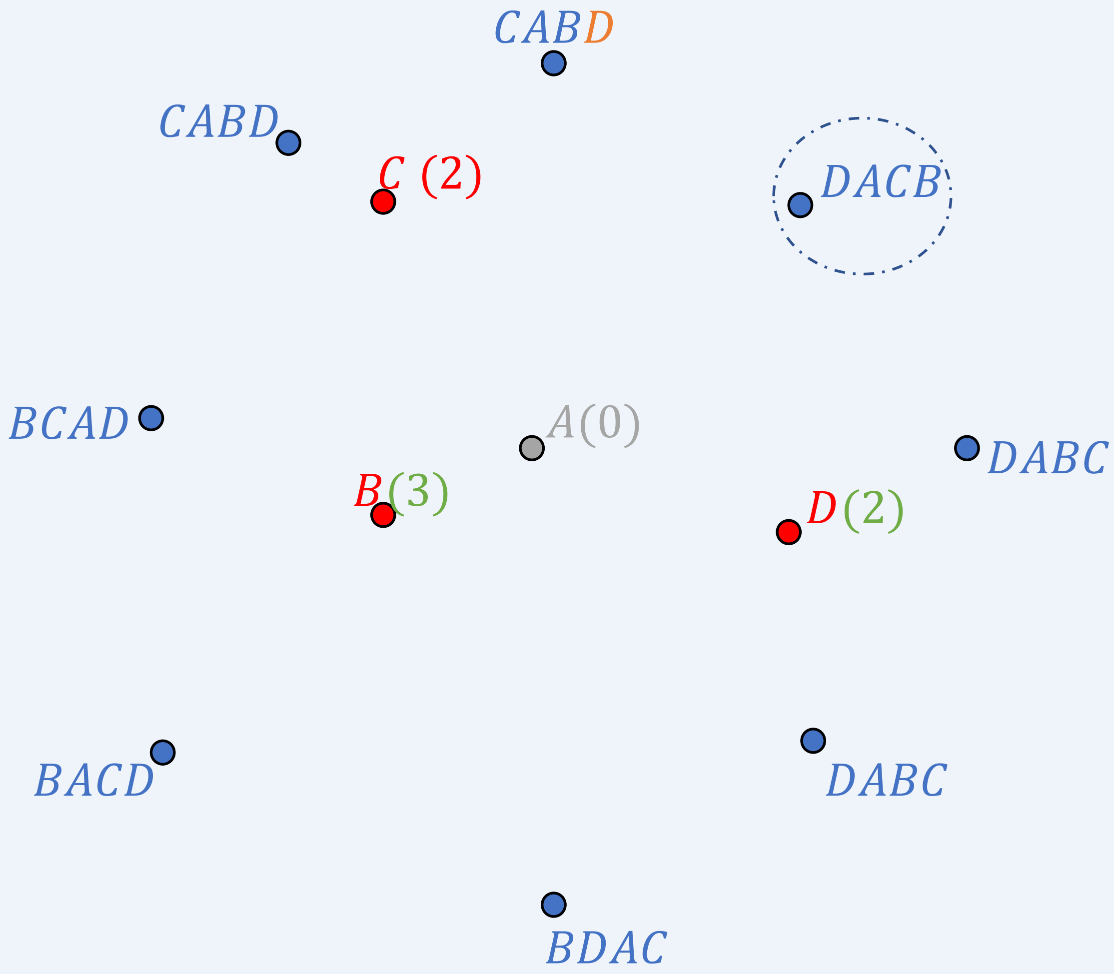

Plurality Veto

- Initially, each candidate X, $\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (veto) of least favorite candidate with positive score
- Last candidate vetoed wins

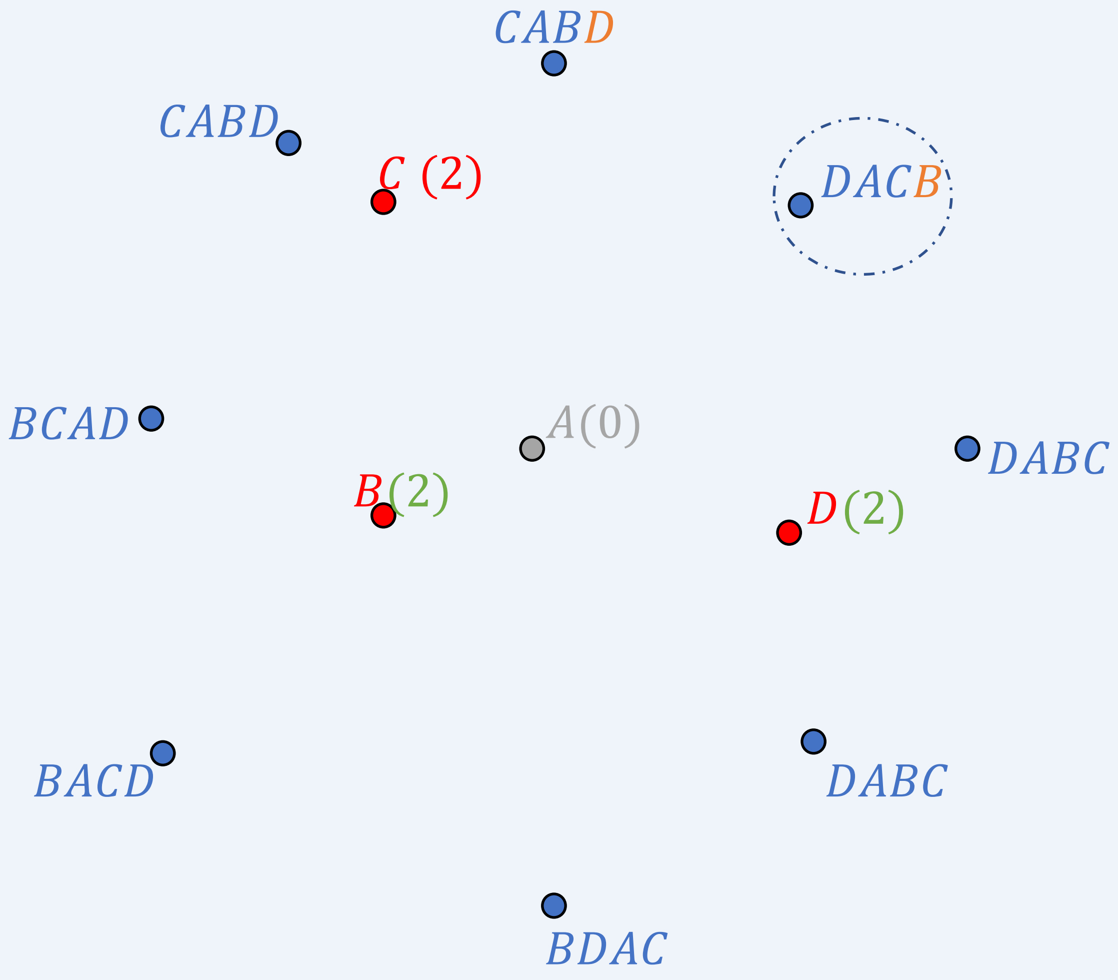

Plurality Veto

- Initially, each candidate X,
$\text{score}(X) = \#\ \text{first choice votes for X}$
- One by one, each voter decrements score (veto) of least favorite candidate with positive score
- Last candidate vetoed wins

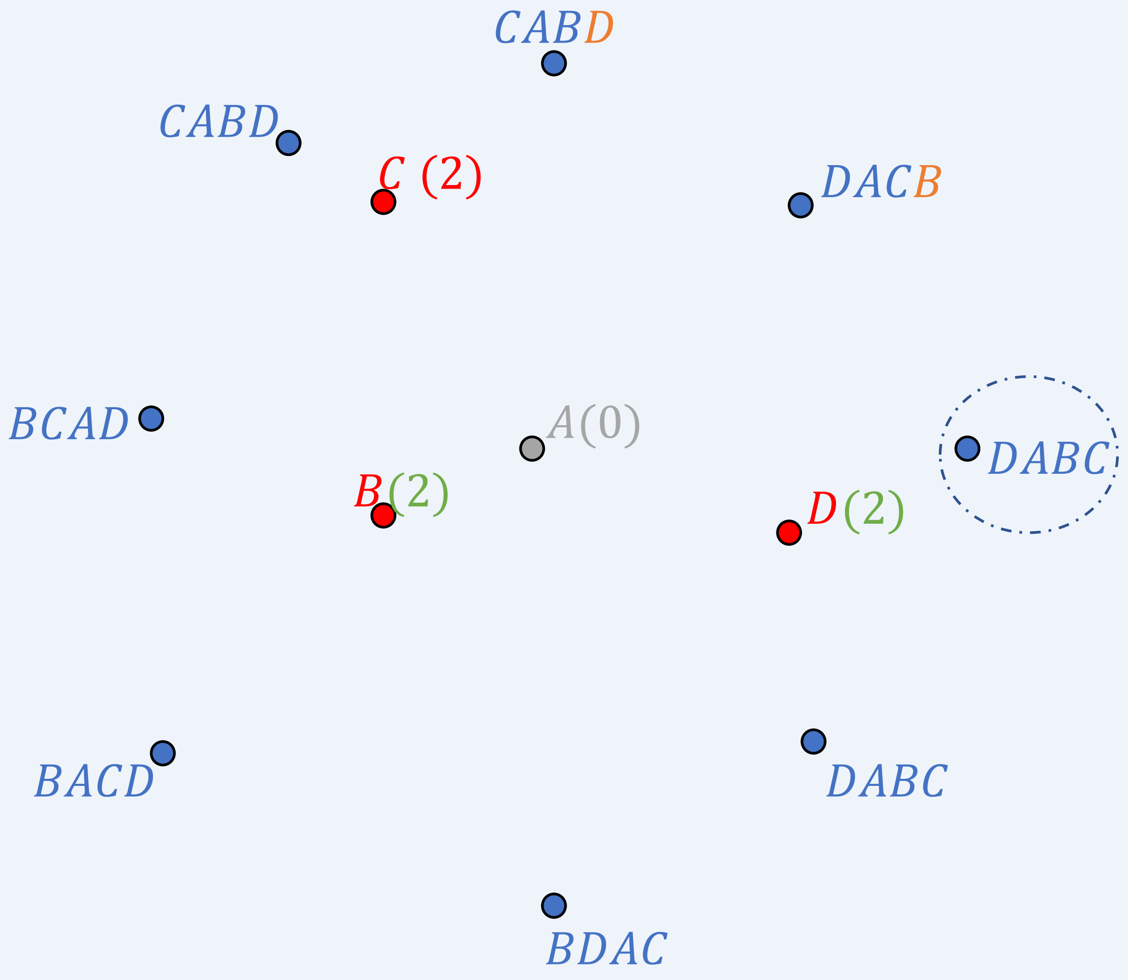

Plurality Veto

- Initially, each candidate X, $\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (veto) of least favorite candidate with positive score
- Last candidate vetoed wins

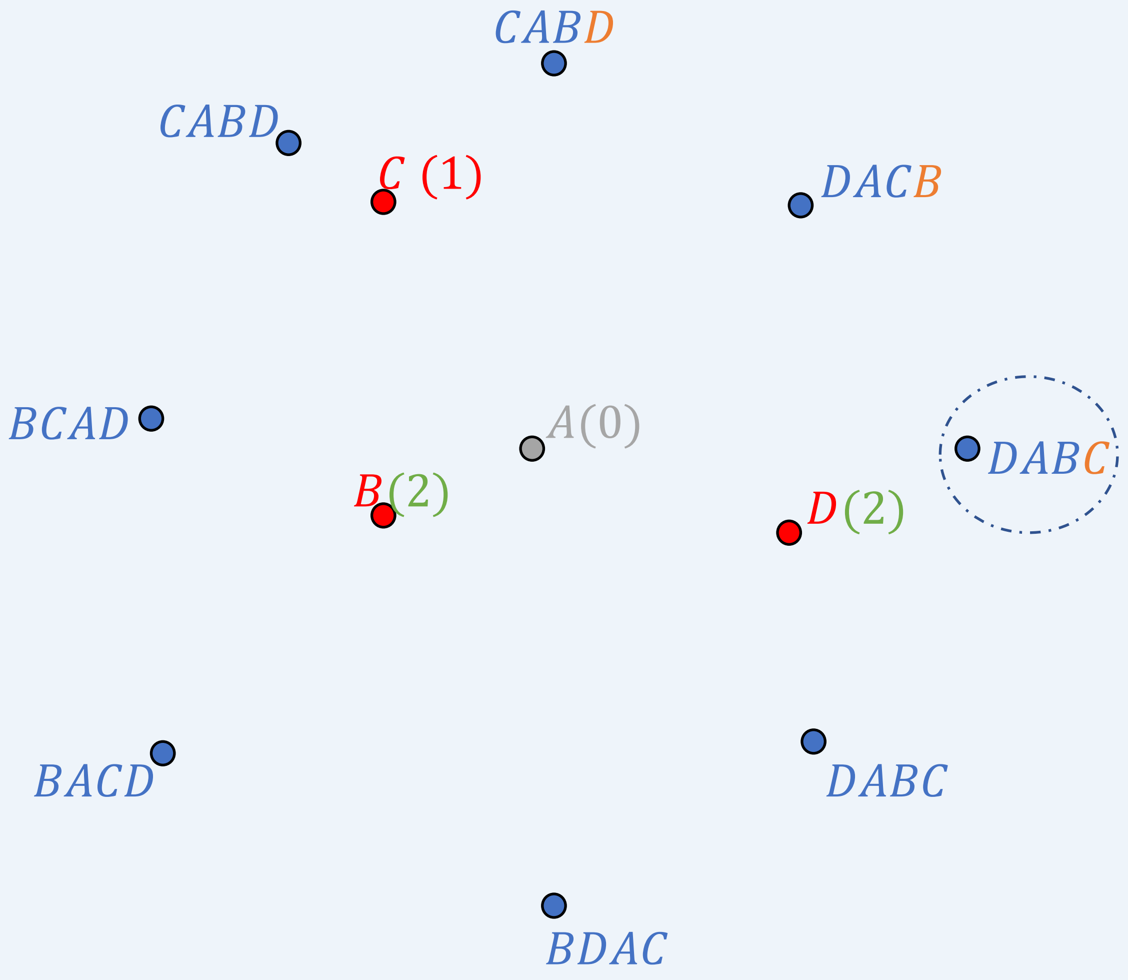

Plurality Veto

- Initially, each candidate X, $\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (veto) of least favorite candidate with positive score
- Last candidate vetoed wins

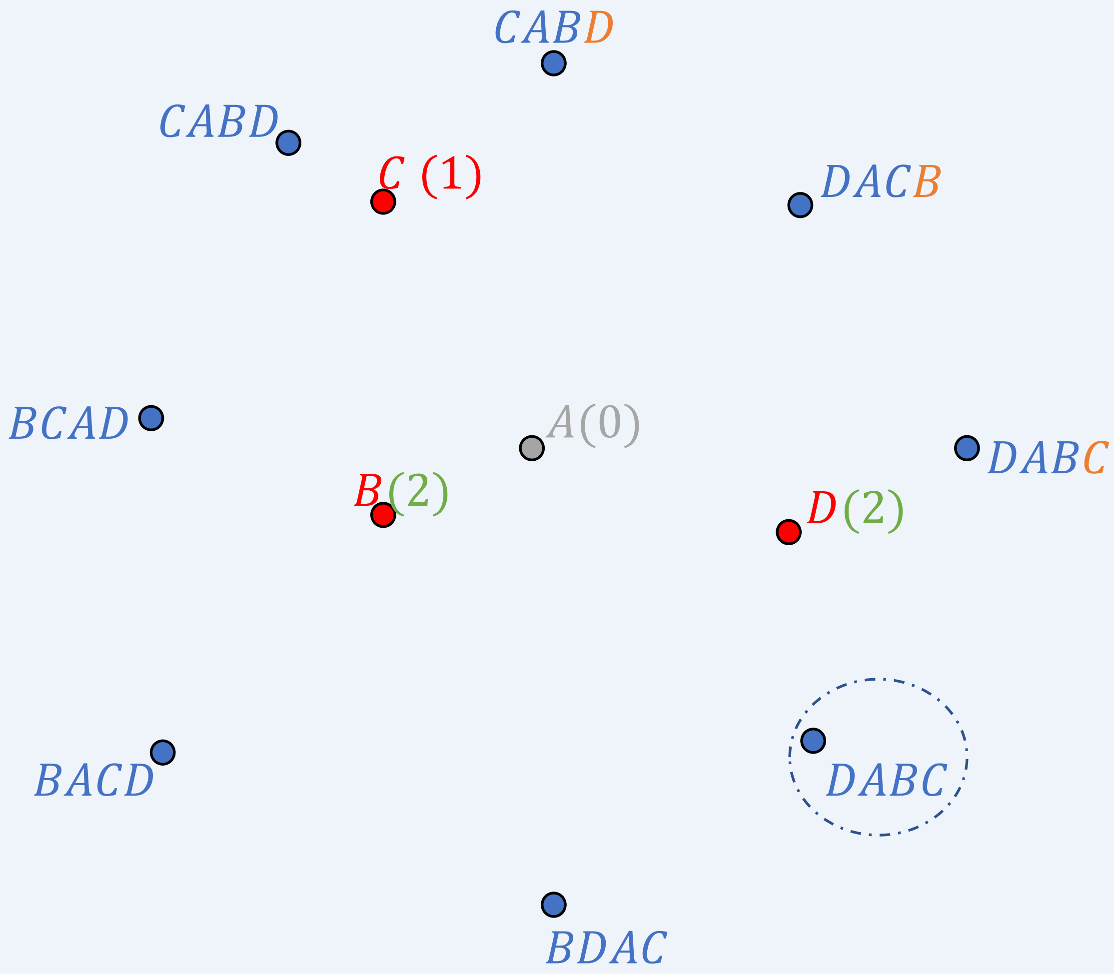

Plurality Veto

- Initially, each candidate X,
$\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (veto) of least favorite candidate with positive score
- Last candidate vetoed wins

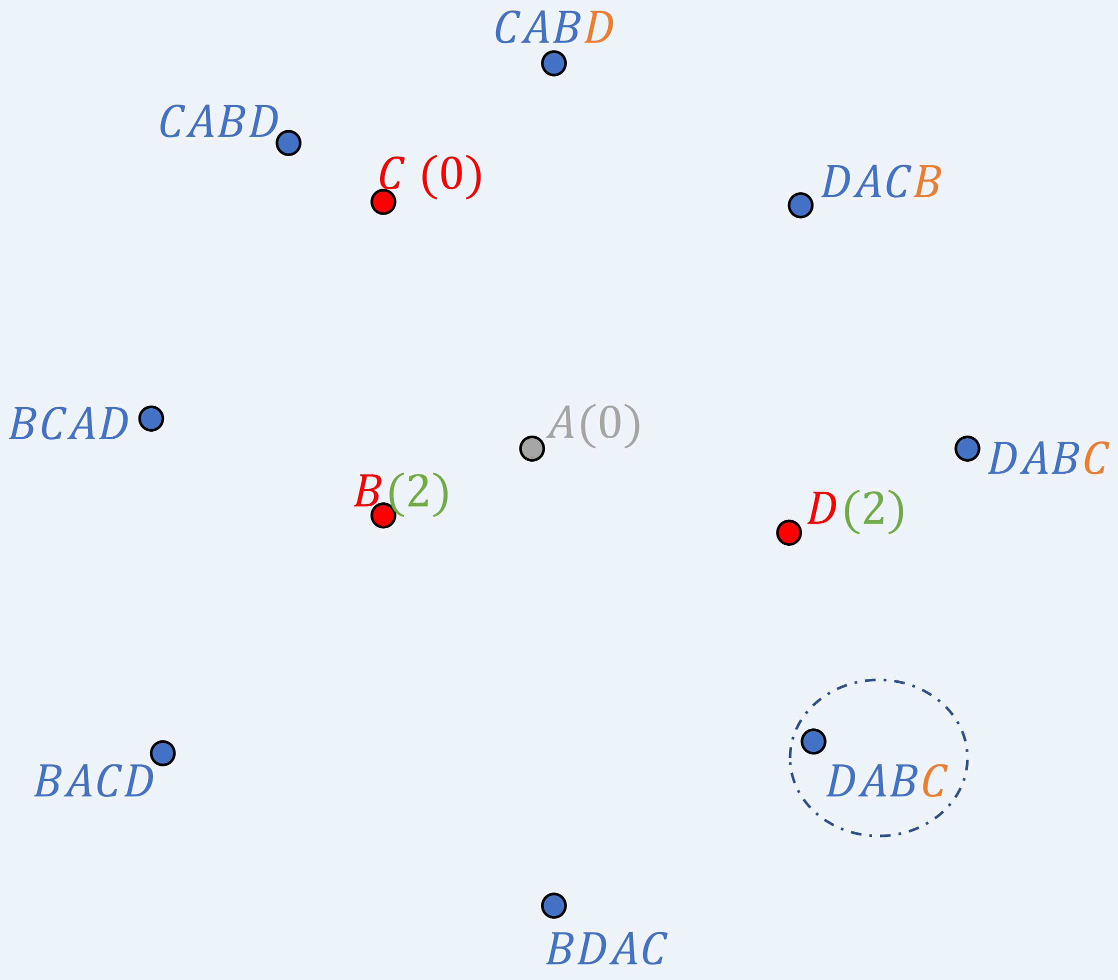

Plurality Veto

- Initially, each candidate X,
$\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (veto) of least favorite candidate with positive score
- Last candidate vetoed wins

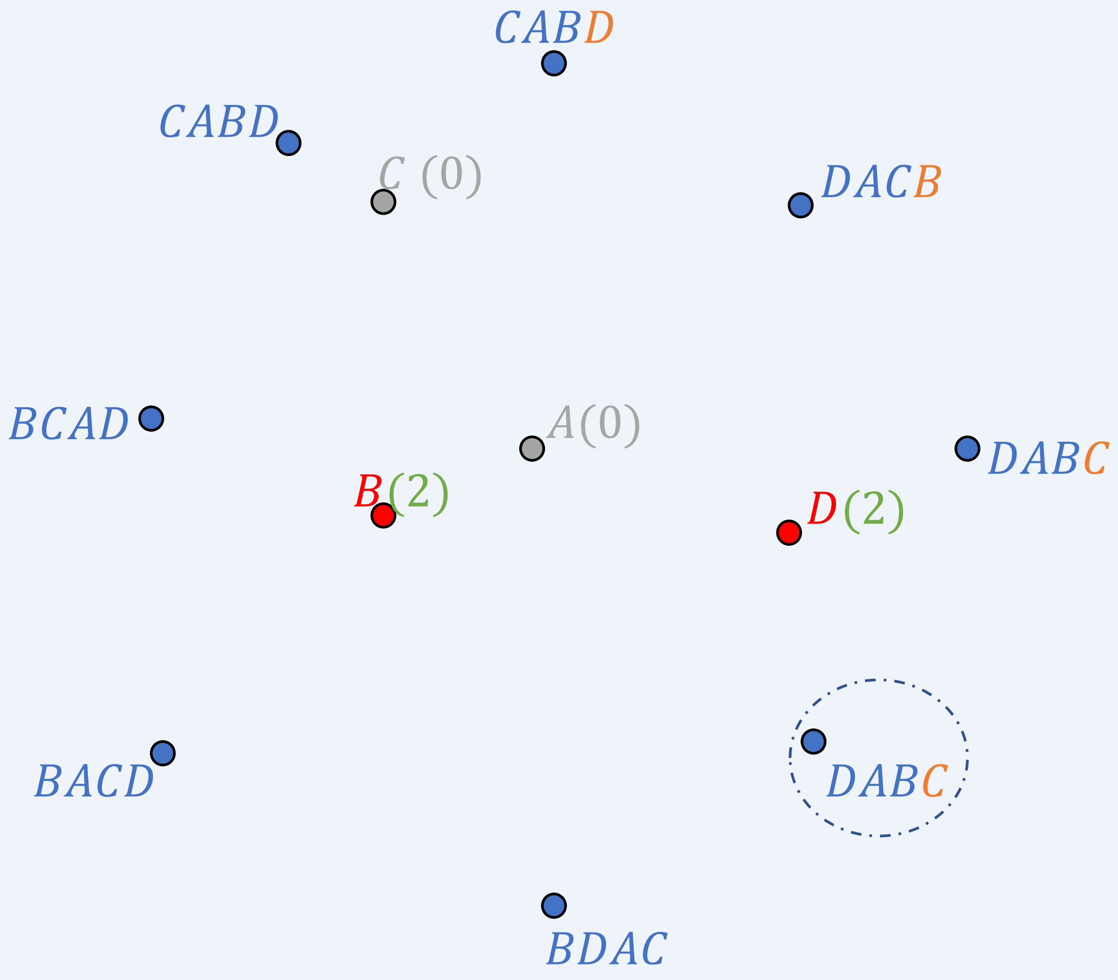

Plurality Veto

- Initially, each candidate X, $\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (veto) of least favorite candidate with positive score
- Last candidate vetoed wins

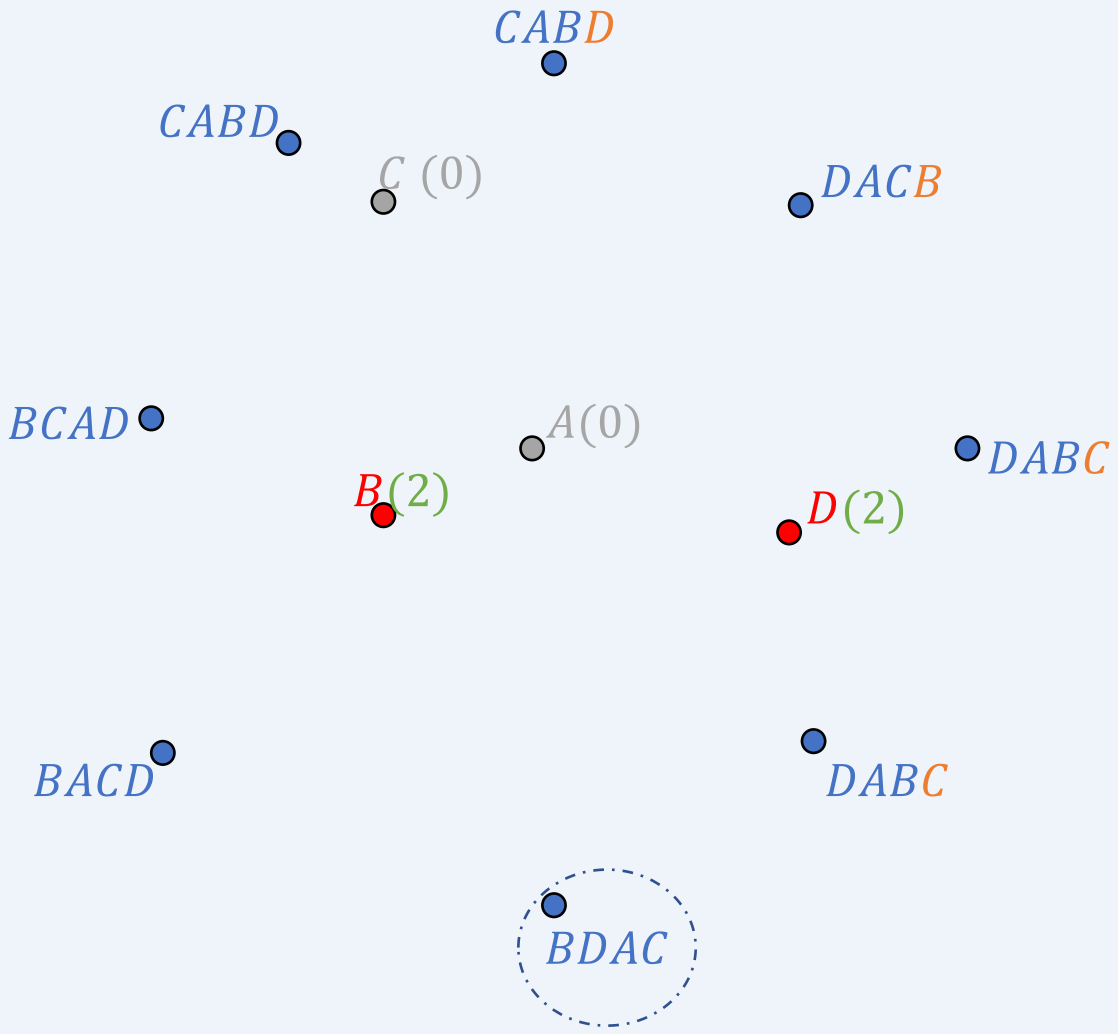

Plurality Veto

- Initially, each candidate X,
$\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (**veto**) of least favorite candidate with positive score
- Last candidate vetoed wins

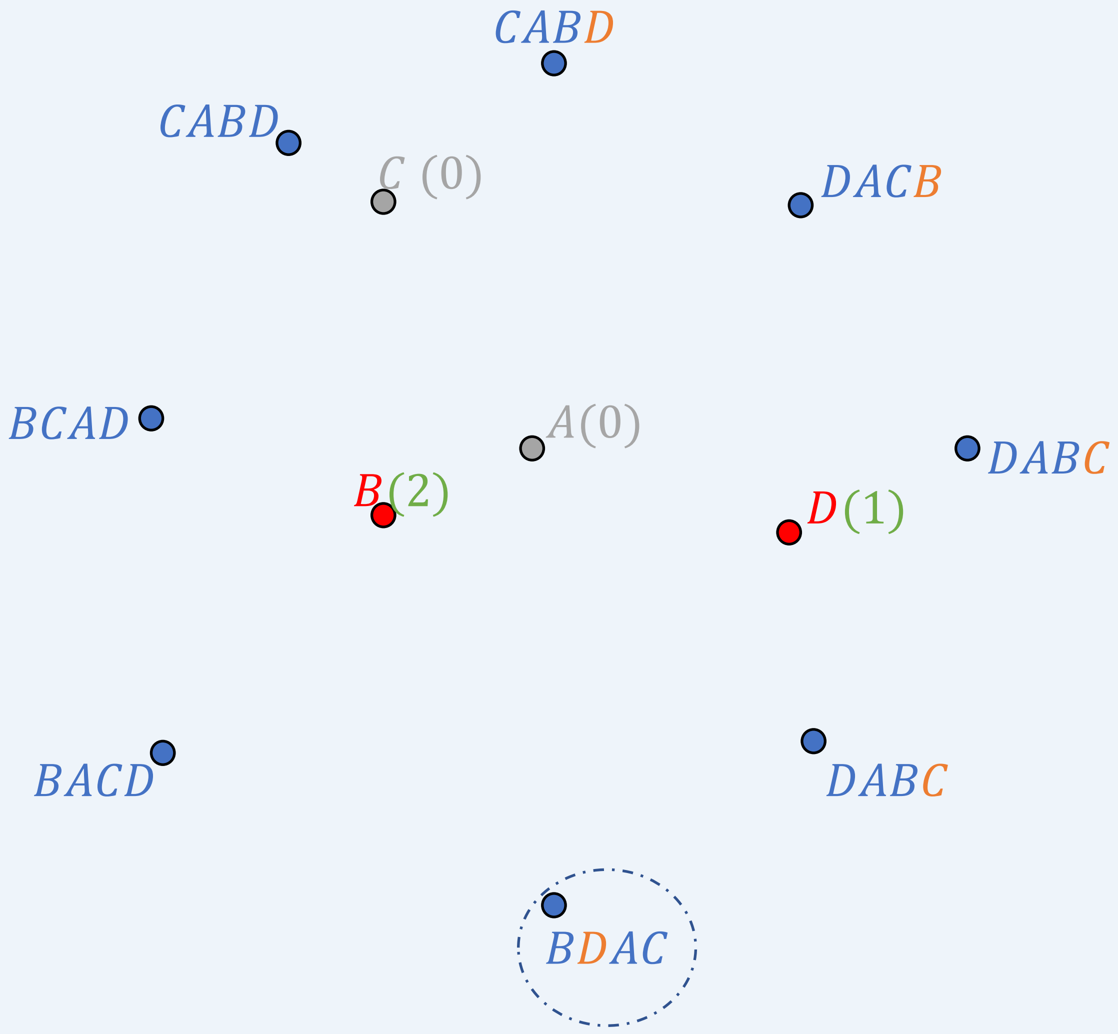

Plurality Veto

- Initially, each candidate X,
$\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (veto) of least favorite candidate with positive score
- Last candidate vetoed wins

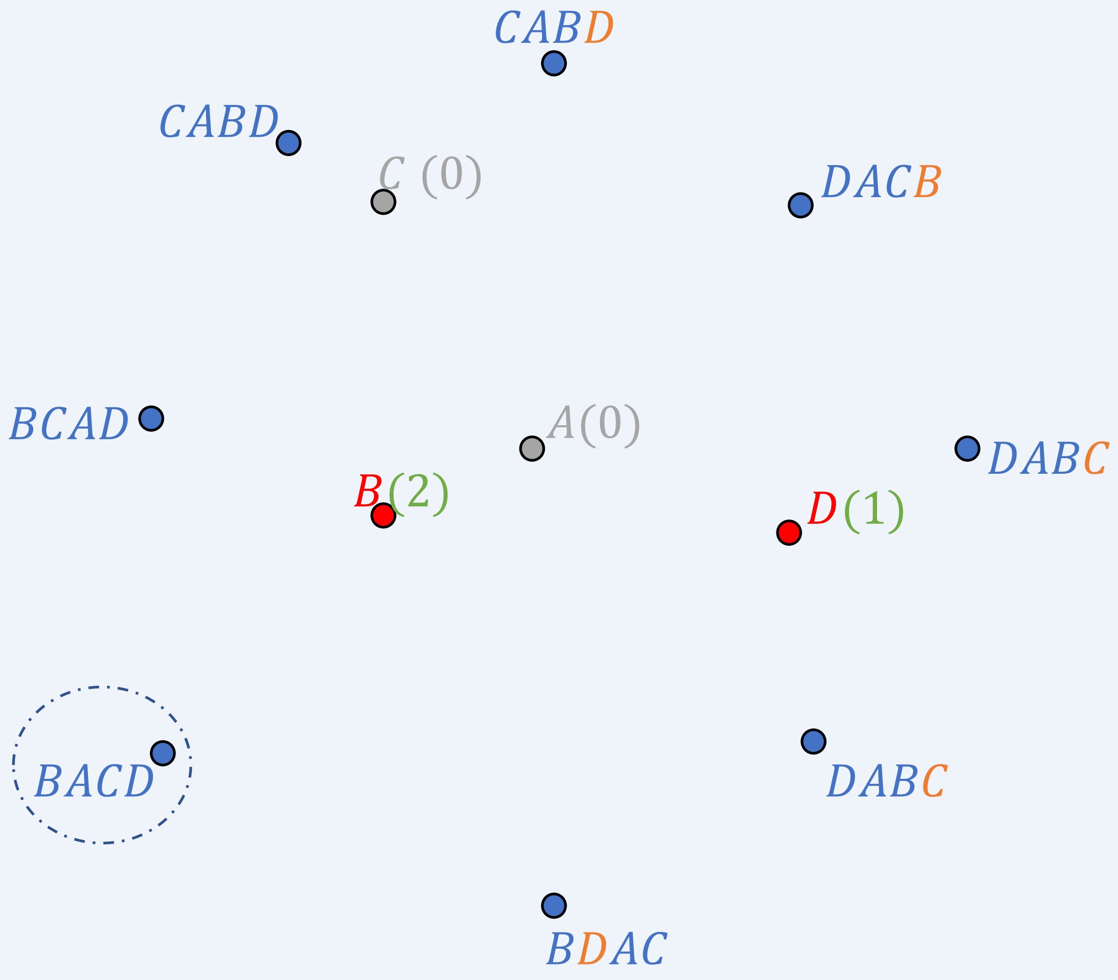

Plurality Veto

- Initially, each candidate X,
$\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (veto) of least favorite candidate with positive score
- Last candidate vetoed wins

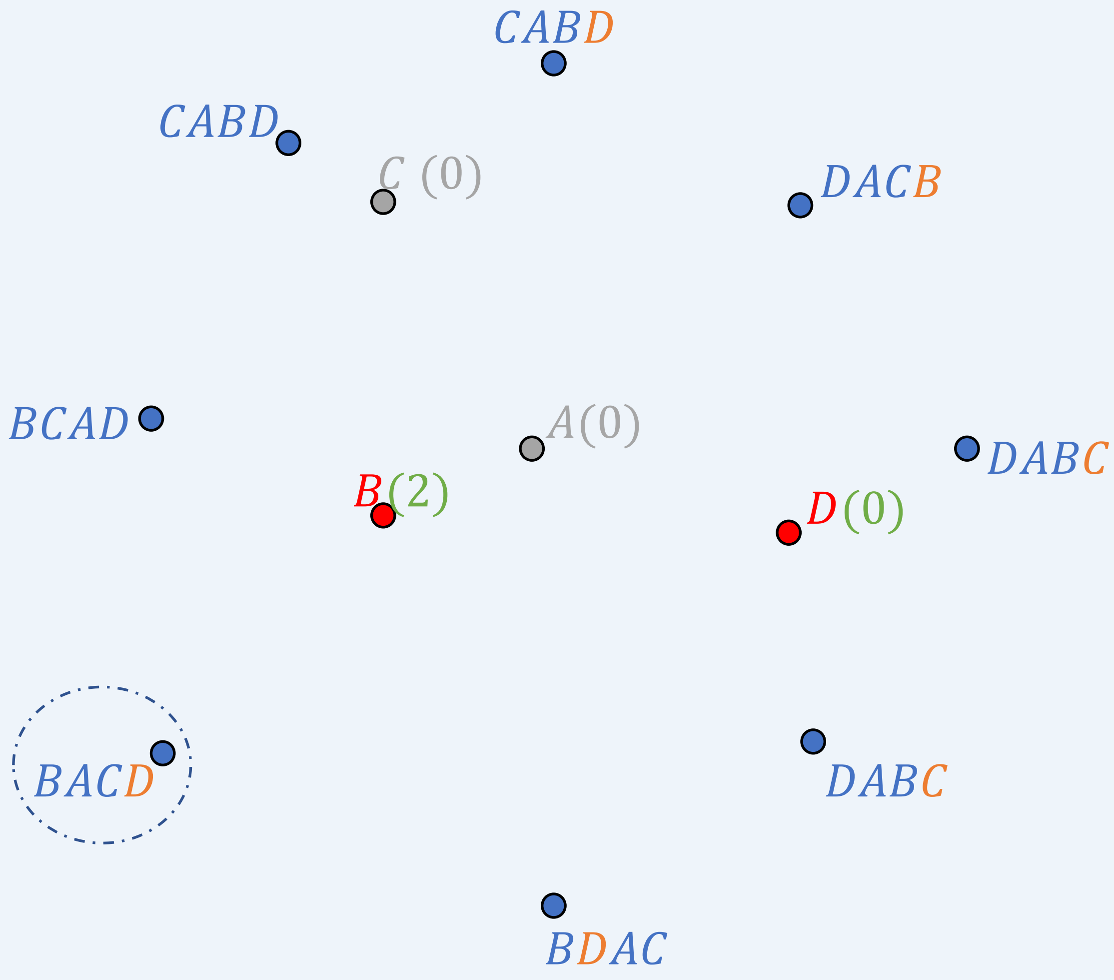

Plurality Veto

- Initially, each candidate X,
$\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (veto) of least favorite candidate with positive score
- Last candidate vetoed wins

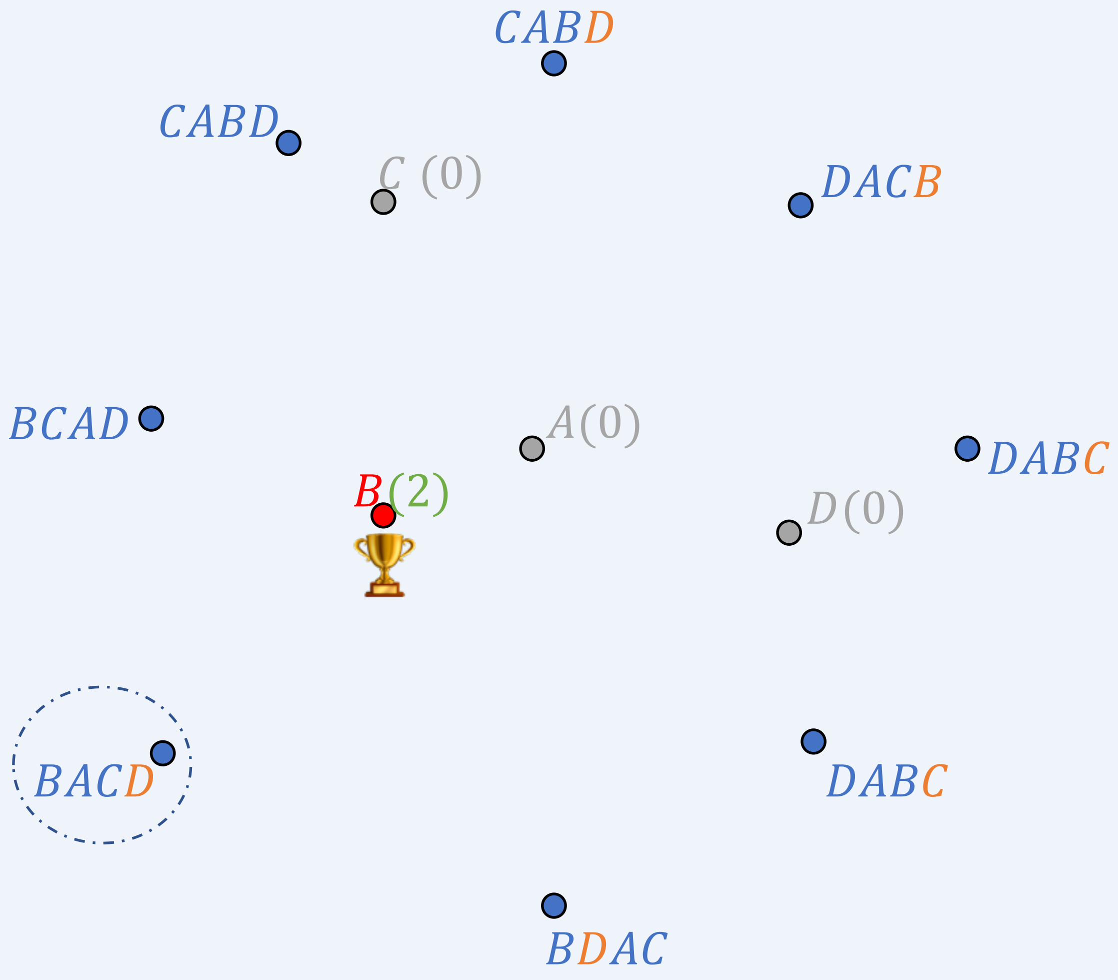

Plurality Veto

- Initially, each candidate X, $\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (**veto**) of least favorite candidate with positive score
- Last candidate vetoed wins!

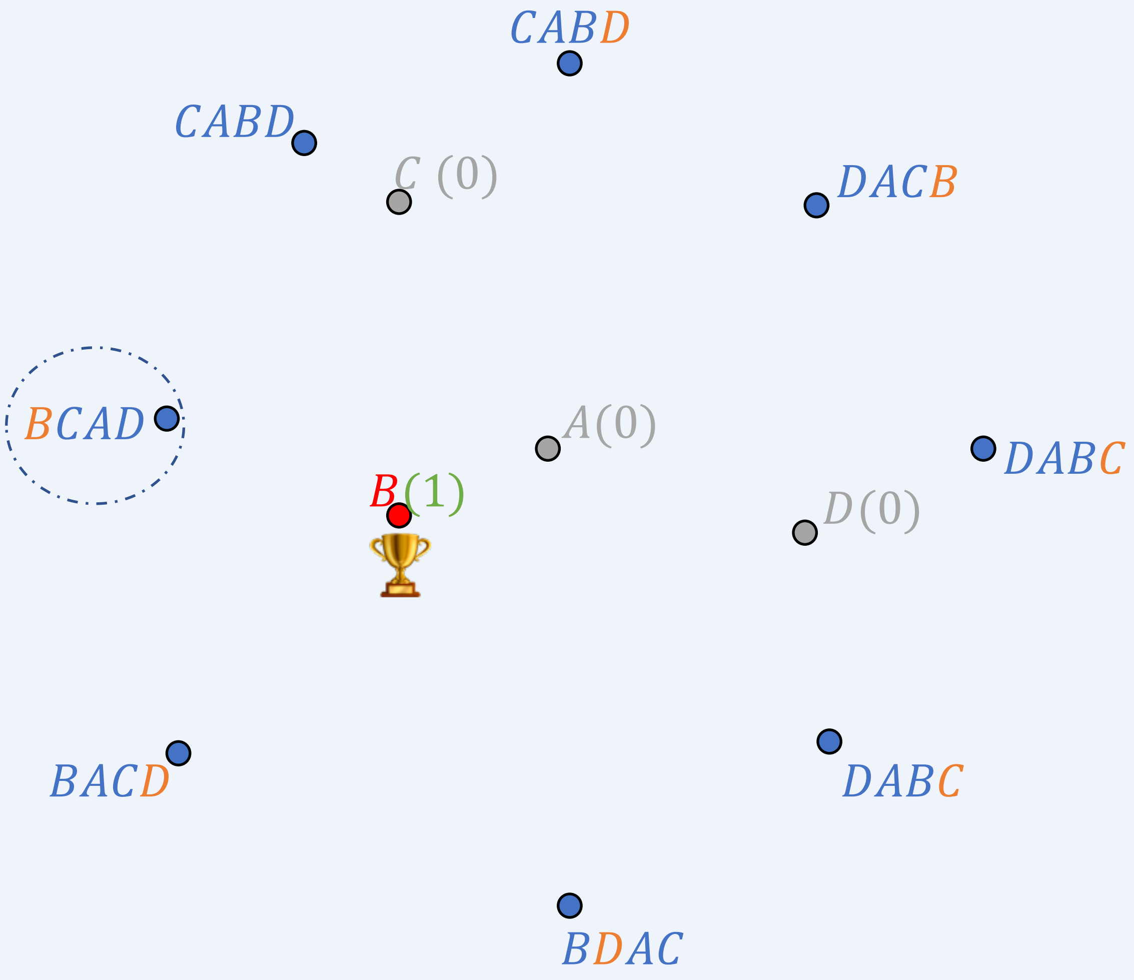

Plurality Veto

- Initially, each candidate X, $\text{score}(X) = \# \text{ first choice votes for X}$
- One by one, each voter decrements score (**veto**) of least favorite candidate with positive score
- Last candidate vetoed wins!

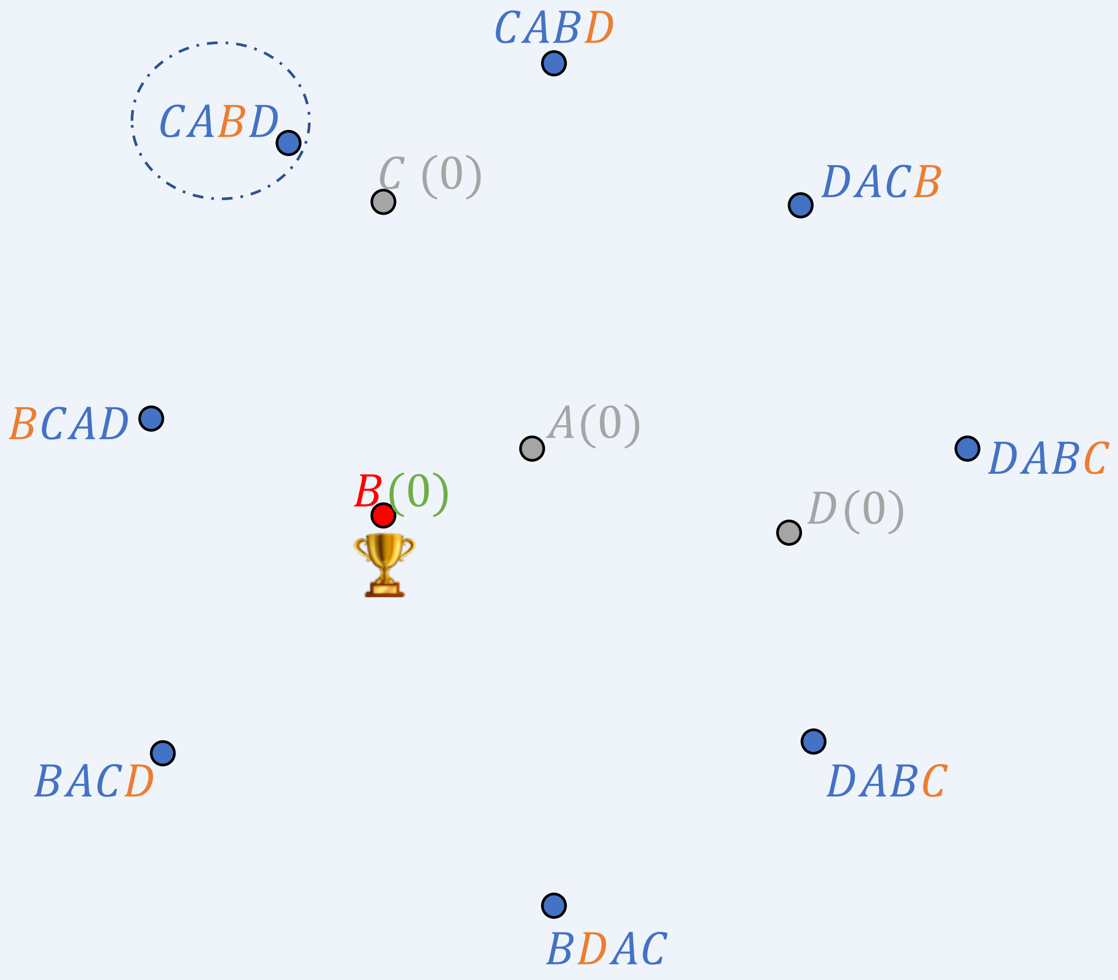

Plurality Veto

- Initially, each candidate X,
  $\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (**veto**) of least favorite candidate with positive score
- Last candidate vetoed wins!

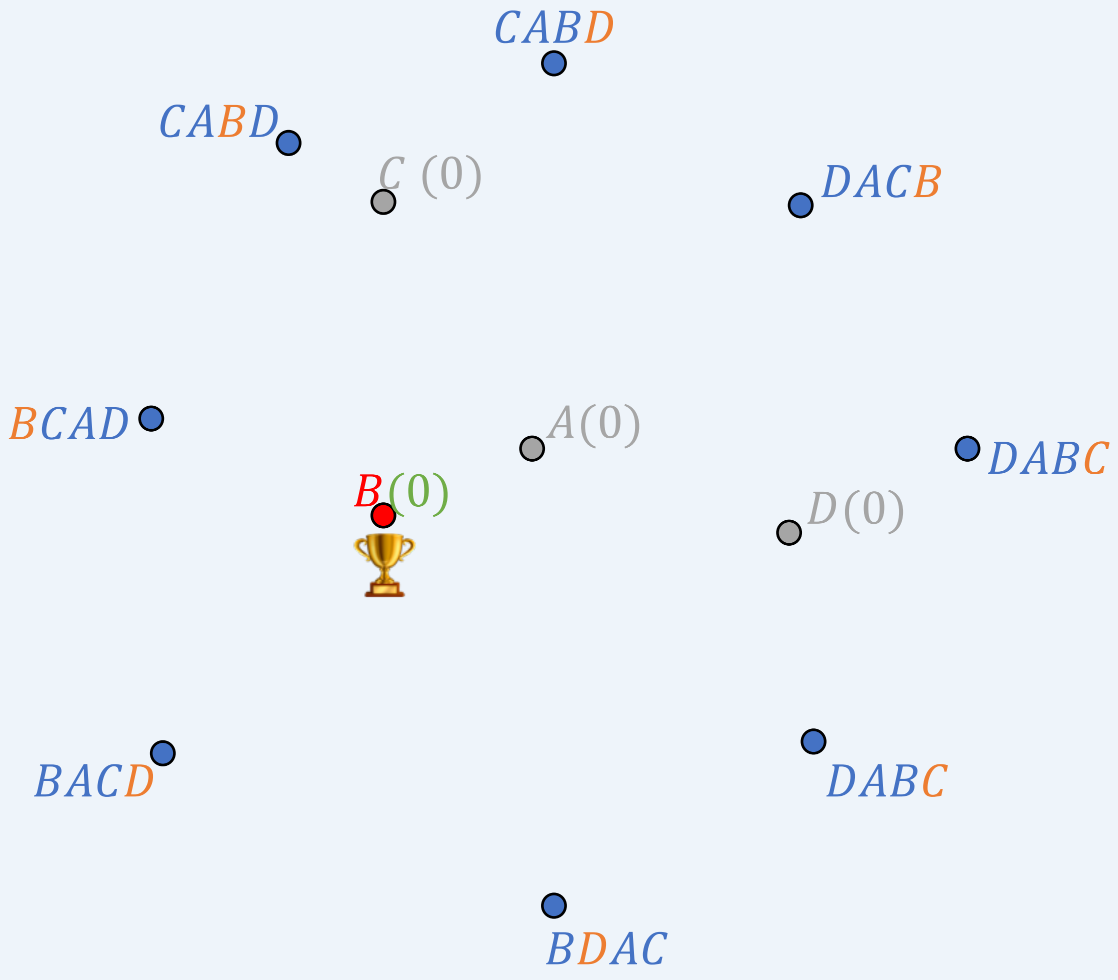

Proof of distortion 3

- Goal: $cost(B) \leq 3 \cdot cost(A)$
- Key observations:
  - Voters closer to B than veto
  - Candidates: # vetos = # first choice votes
- So far…

$$cost(B) \leq [\text{edges shown}]$$

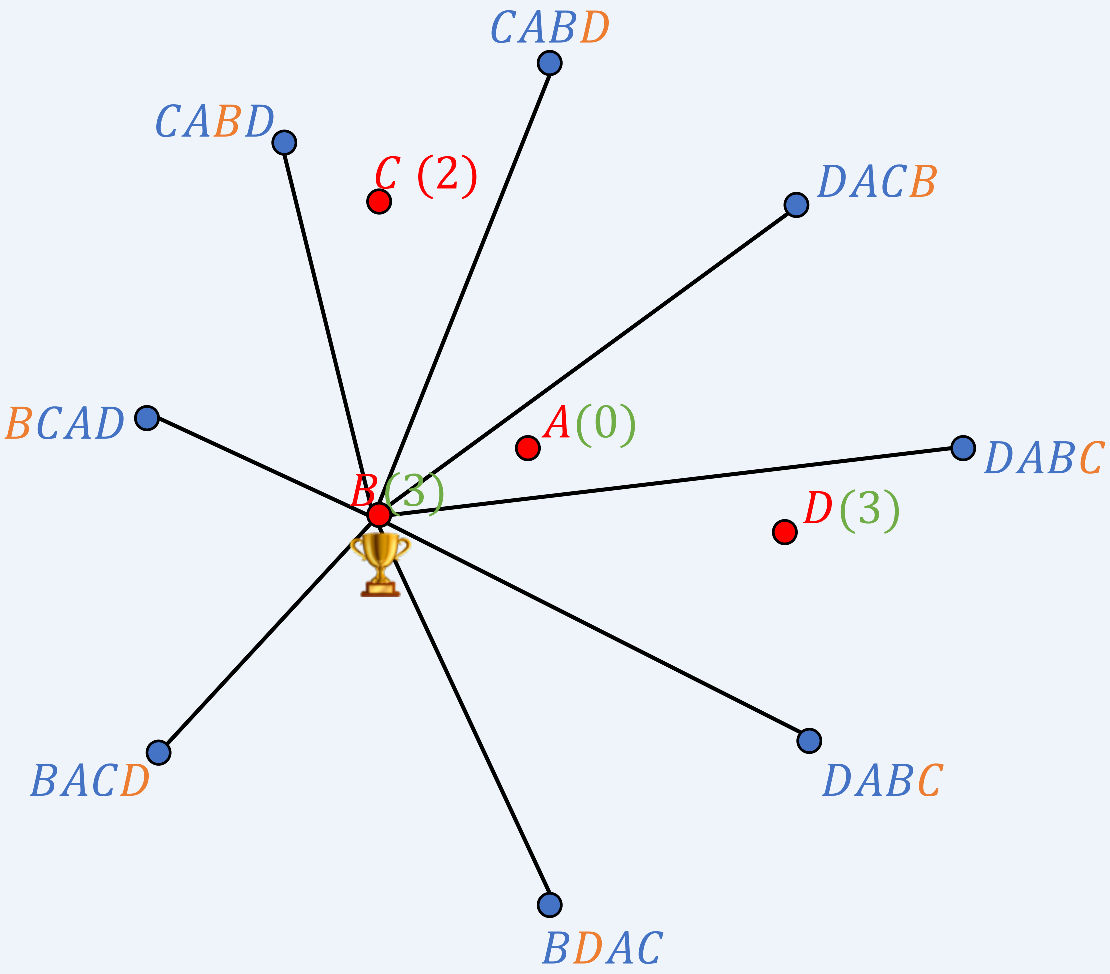

Plurality Veto

- Initially, each candidate X,
  $\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (veto) of least favorite candidate with positive score
- Last candidate vetoed wins!

Proof of distortion 3

- Goal: $cost(B) \leq 3 \cdot cost(A)$
- Key observations:
  - Voters closer to B than veto
  - Candidates: # vetos = # first choice votes
- So far…

$$cost(B) \leq [\text{edges shown}]$$

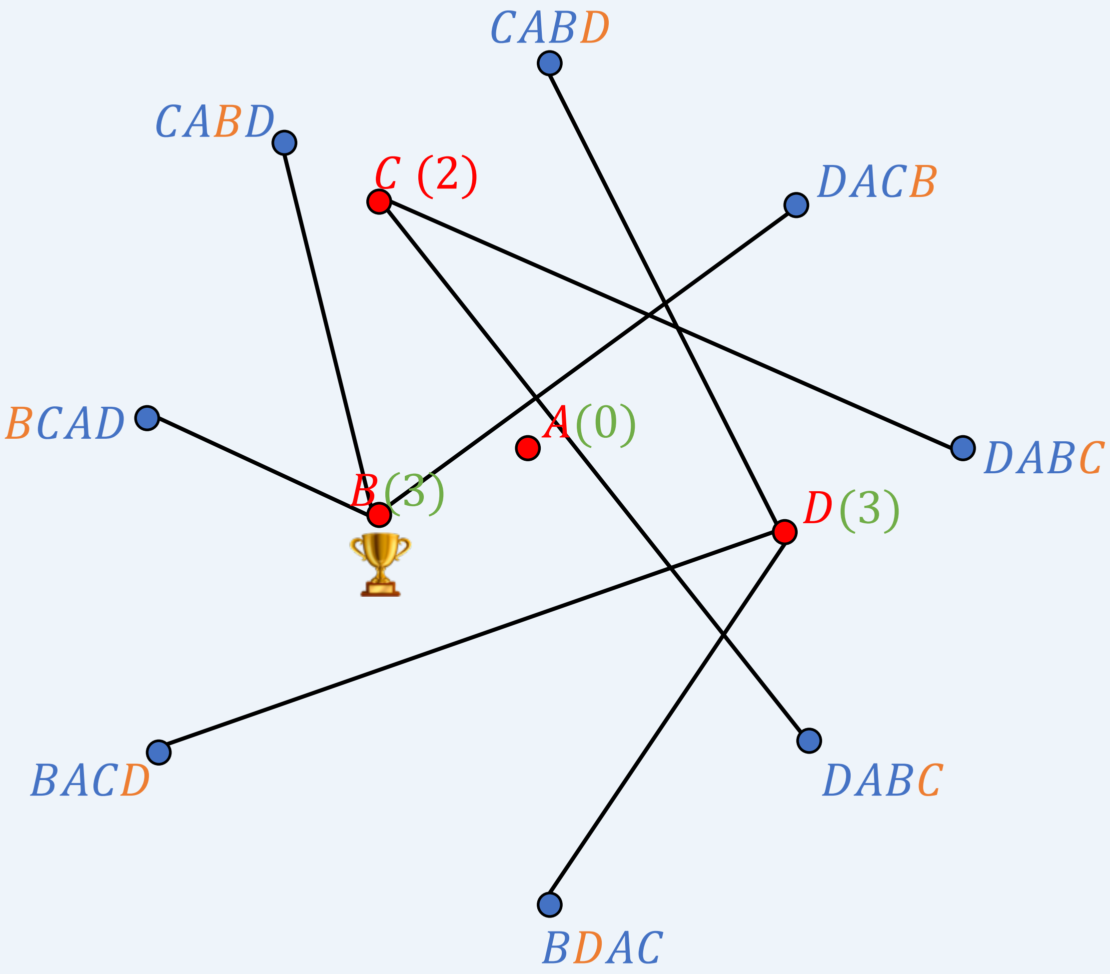

Plurality Veto

- Initially, each candidate X,
  $\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (veto) of least favorite candidate with positive score
- Last candidate vetoed wins!

Proof of distortion 3

- Goal: $cost(B) \leq 3 \cdot cost(A)$
- Key observations:
  - Voters closer to B than veto
  - Candidates: # vetos = # first choice votes
- So far…

$$cost(B) \leq [\text{edges shown}]$$

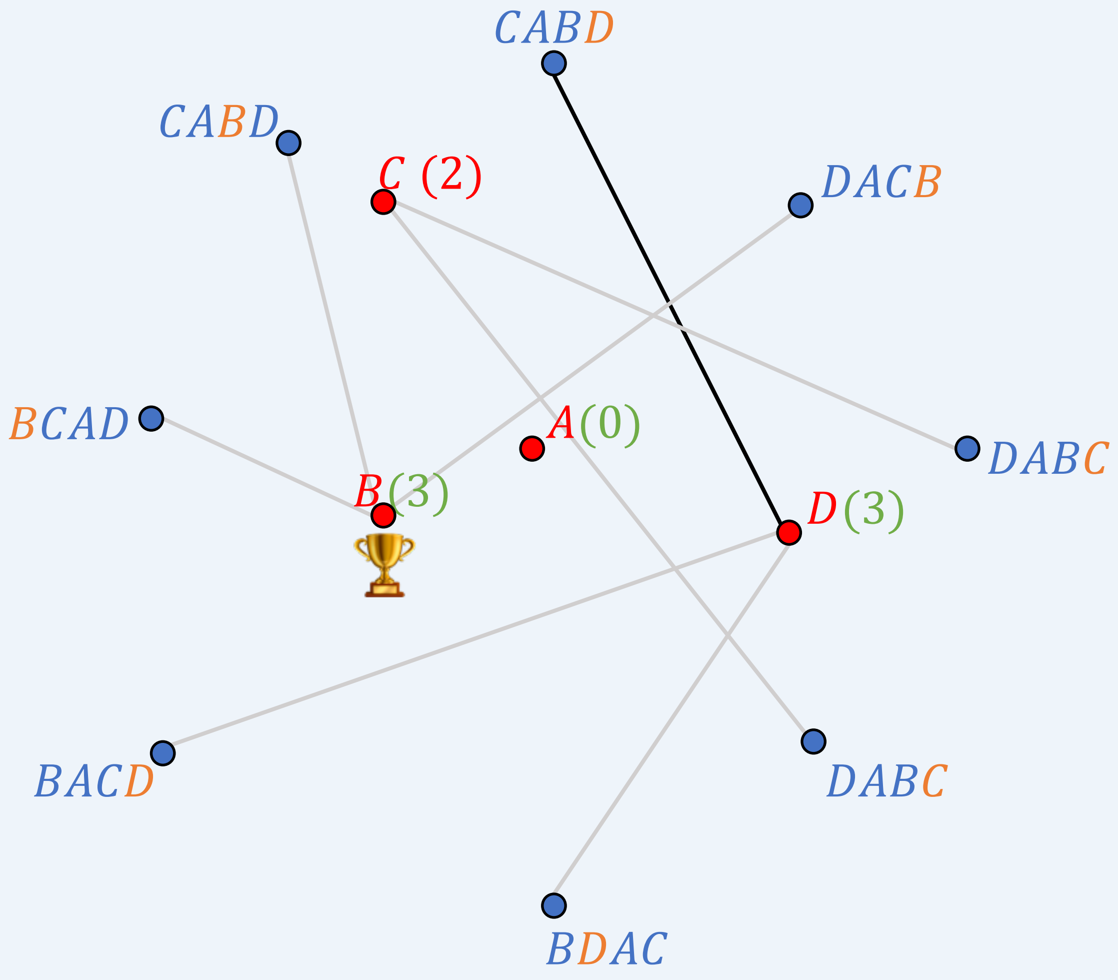

Plurality Veto

- Initially, each candidate X,
  $\text{score}(X) = \text{ \# first choice votes for X}$
- One by one, each voter decrements score (veto) of least favorite candidate with positive score
- Last candidate vetoed wins!

Proof of distortion 3

- Goal: $cost(B) \leq 3 \cdot cost(A)$
- Key observations:
  - Voters closer to B than veto
  - Candidates: # vetos = # first choice votes
- So far…

$$cost(B) \leq [\text{edges shown}]$$

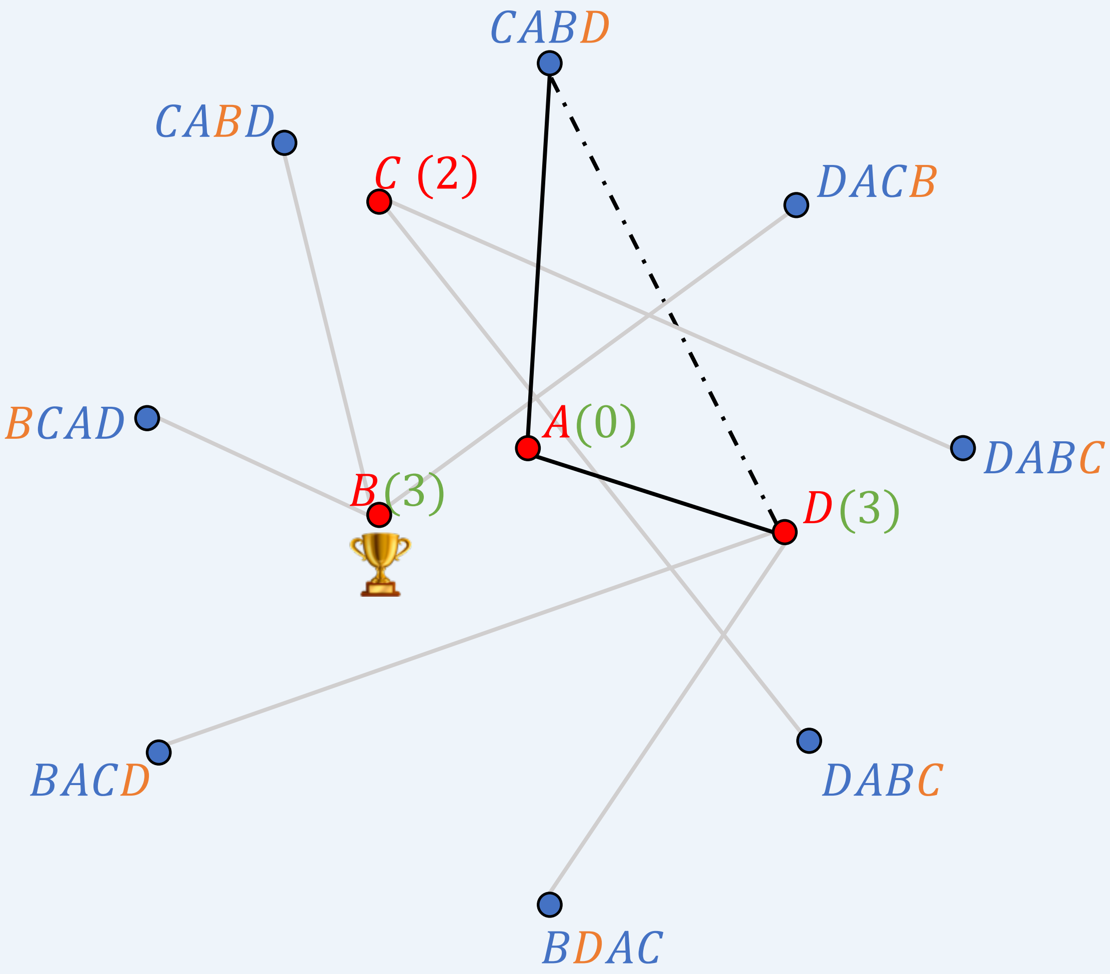

Plurality Veto

- Initially, each candidate X,
  $\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (veto) of least favorite candidate with positive score
- Last candidate vetoed wins!

Proof of distortion 3

- Goal: $cost(B) \leq 3 \cdot cost(A)$
- Key observations:
  - Voters closer to B than veto
  - Candidates: # vetos = # first choice votes
- So far…

$$cost(B) \leq [\text{edges shown}]$$

Plurality Veto

- Initially, each candidate X,
  $\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (**veto**) of least favorite candidate with positive score
- Last candidate vetoed wins!

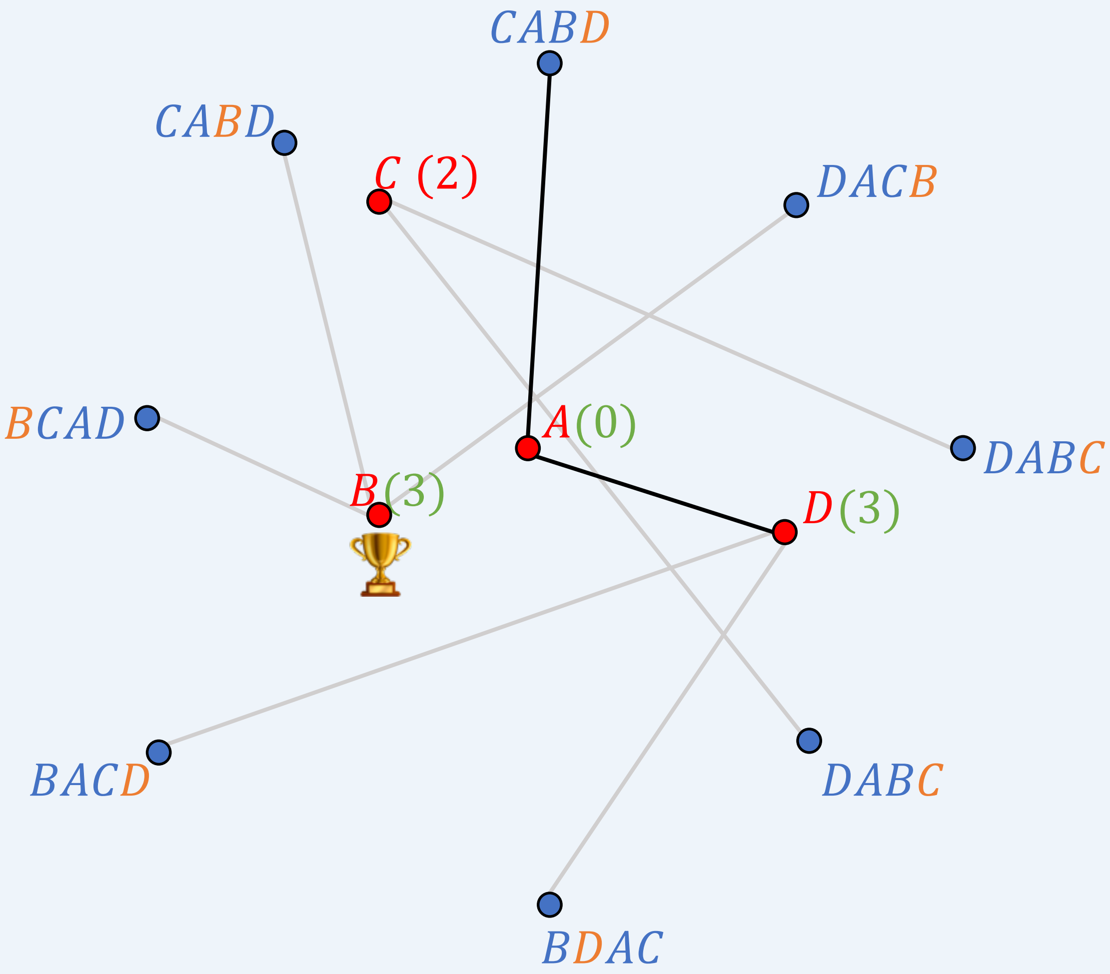

Proof of distortion 3

- Goal: $cost(B) \leq 3 \cdot cost(A)$
- Key observations:
  - Voters closer to B than veto
  - Candidates: # vetos = # first choice votes
- So far…

$$cost(B) \leq [\text{edges shown}]$$

Plurality Veto

- Initially, each candidate X,
  $\text{score}(X) = \text{ \# first choice votes for X}$
- One by one, each voter decrements score (**veto**) of least favorite candidate with positive score
- Last candidate vetoed wins!

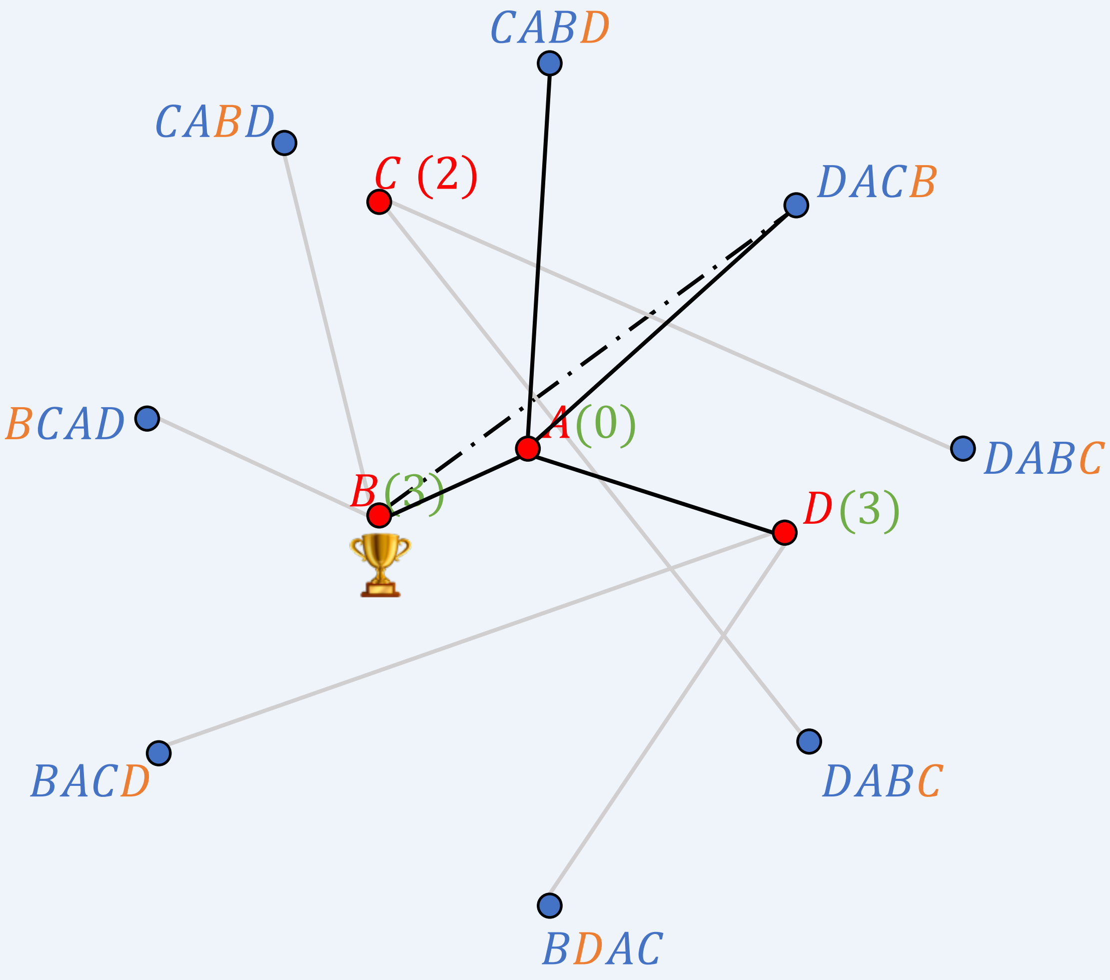

Proof of distortion 3

- Goal: $cost(B) \leq 3 \cdot cost(A)$
- Key observations:
  - Voters closer to B than veto
  - Candidates: # vetos = # first choice votes
- So far…

$$cost(B) \leq [\text{edges shown}]$$

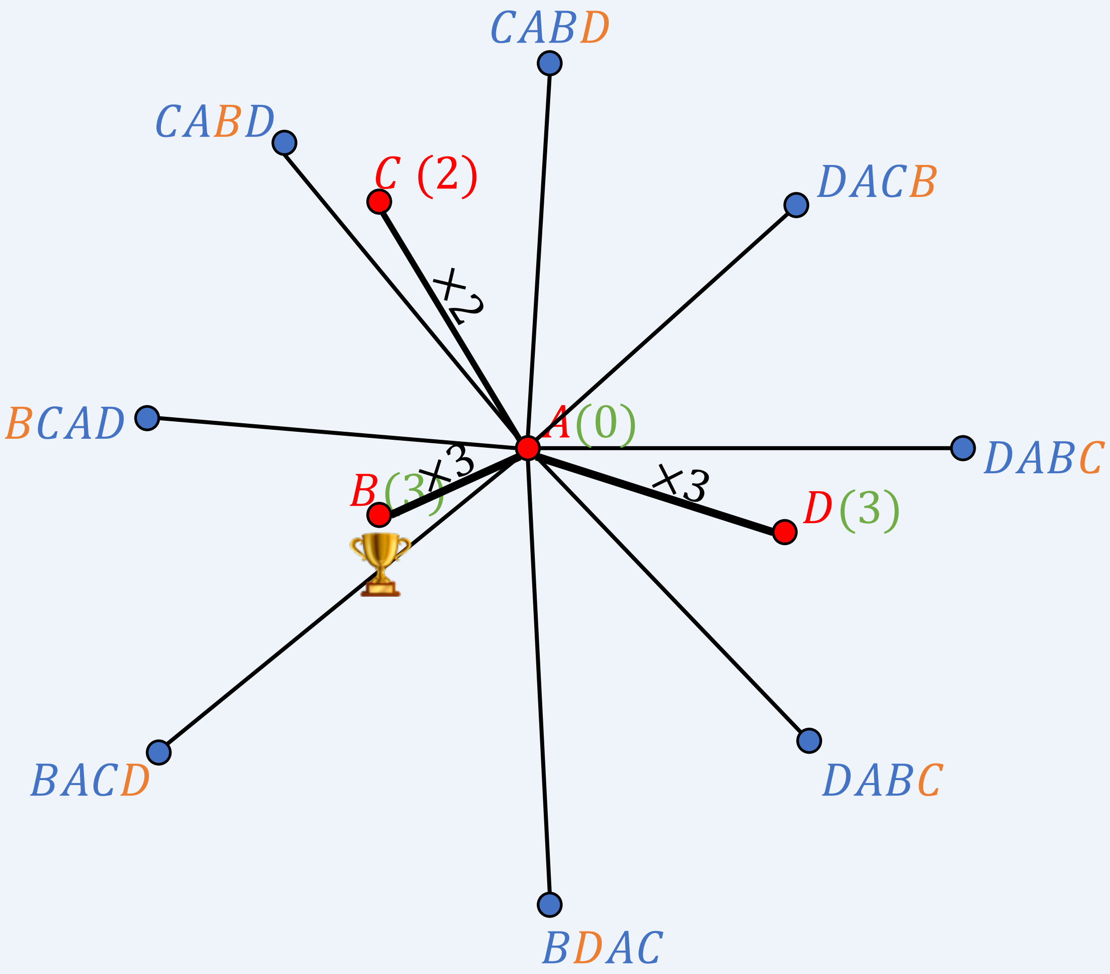

Plurality Veto

- Initially, each candidate X,
  $\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (**veto**) of least favorite candidate with positive score
- Last candidate vetoed wins!

Proof of distortion 3

- Goal: $cost(B) \leq 3 \cdot cost(A)$
- Key observations:
  - Voters closer to B than veto
  - Candidates: # vetos = # first choice votes
- So far…

$$cost(B) \leq [\text{edges shown}]$$

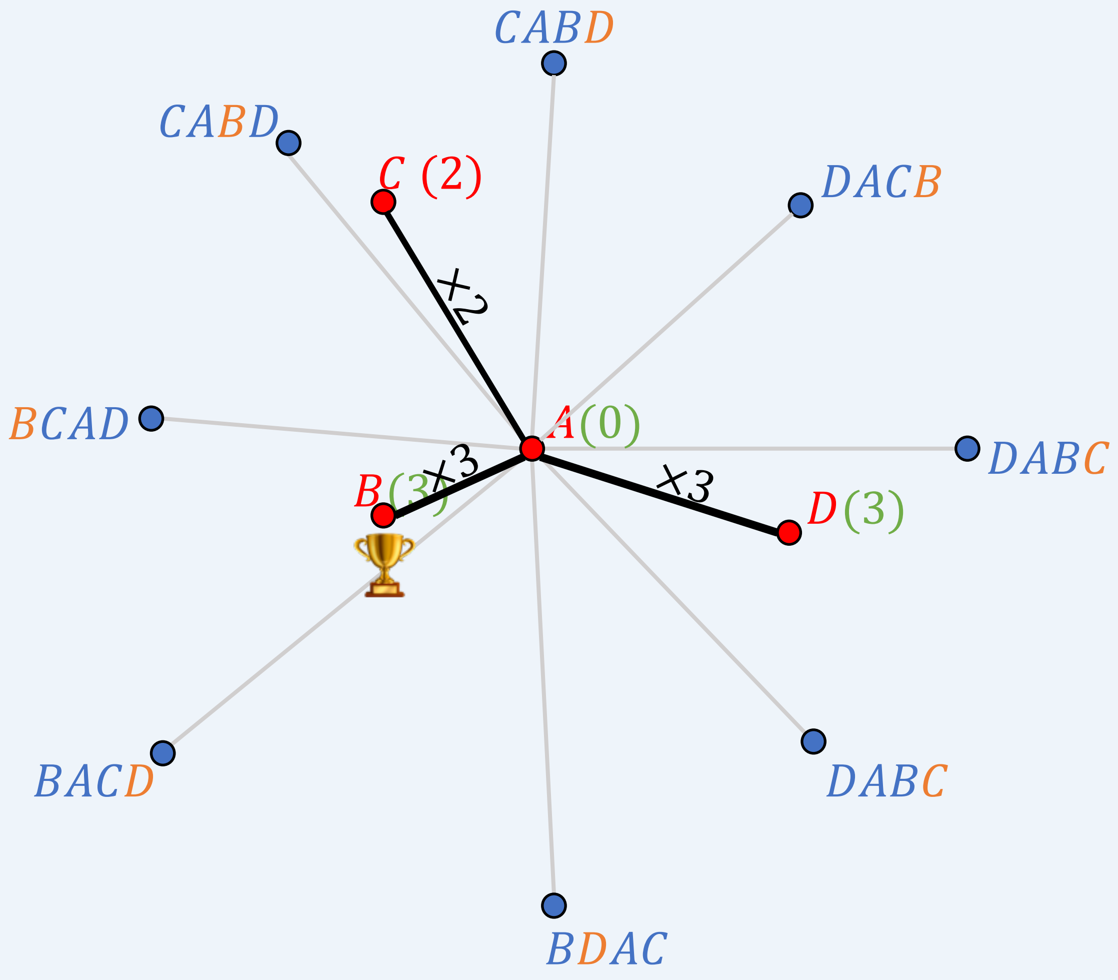

Plurality Veto

- Initially, each candidate X,
  $\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (veto) of least favorite candidate with positive score
- Last candidate vetoed wins!

Proof of distortion 3

- Goal: $cost(B) \leq 3 \cdot cost(A)$
- Key observations:
  - Voters closer to B than veto
  - Candidates: # vetos = # first choice votes
- So far…
  $cost(B) \leq cost(\mathrm{A}) + [\text{edges shown}]$

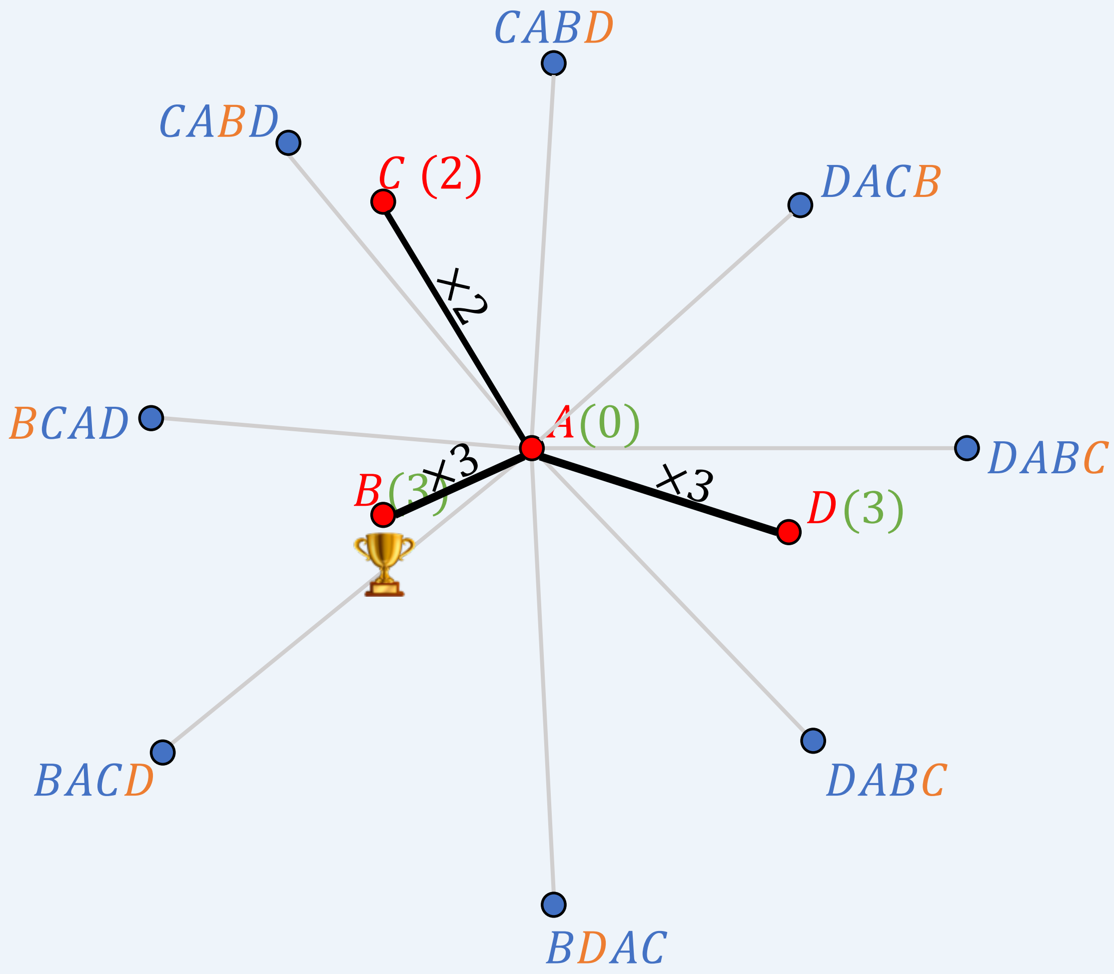

Plurality Veto

- Initially, each candidate X,
  $\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (veto) of least favorite candidate with positive score
- Last candidate vetoed wins!

Proof of distortion 3

- Goal: $cost(B) \leq 3 \cdot cost(A)$
- Key observations:
  - Voters closer to B than veto
  - Candidates: # vetos = # first choice votes
- So far…
  $cost(B) \leq cost(\text{A}) + [\text{edges shown}]$

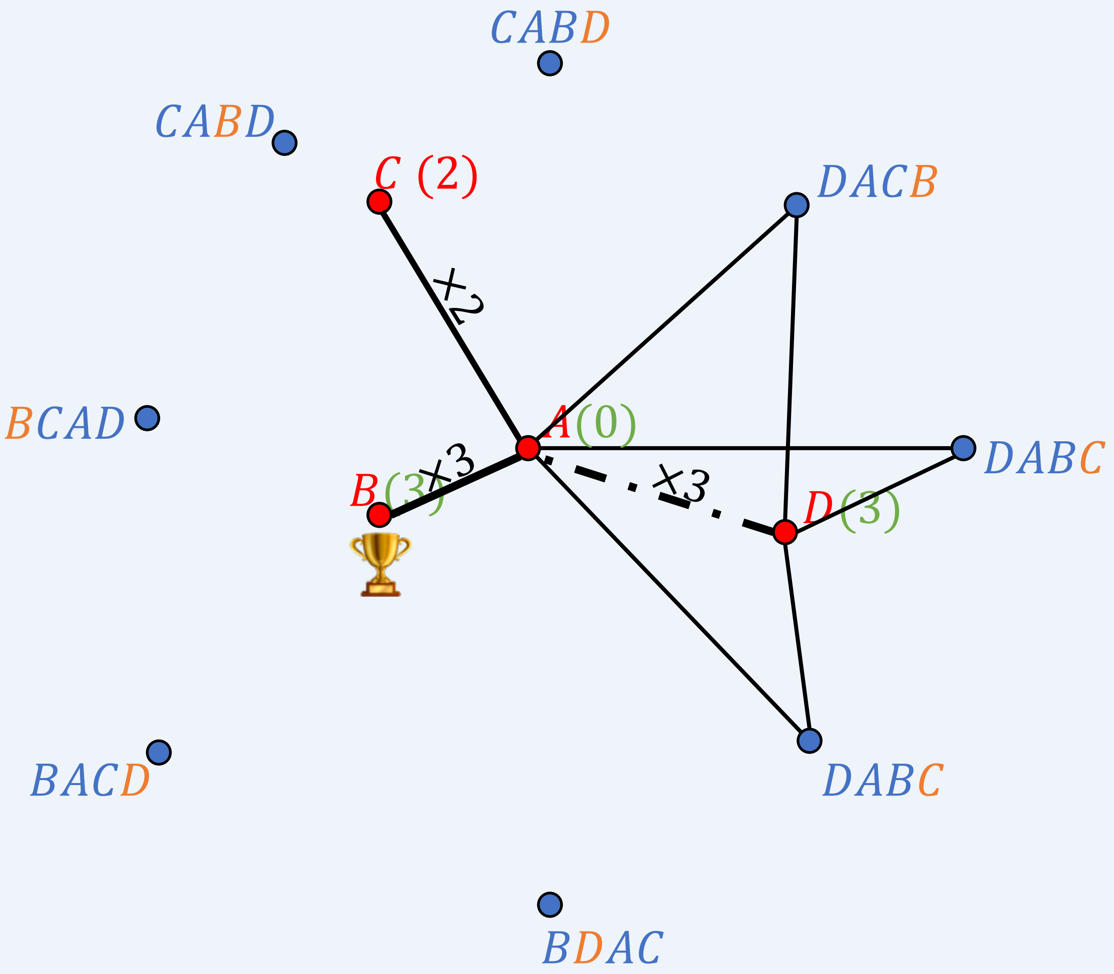

Plurality Veto

- Initially, each candidate X,
  $\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (veto) of least favorite candidate with positive score
- Last candidate vetoed wins!

Proof of distortion 3

- Goal: $cost(B) \leq 3 \cdot cost(A)$
- Key observations:
  - Voters closer to B than veto
  - Candidates: # vetos = # first choice votes
- So far…
  $cost(B) \leq cost(\mathrm{A}) + [\text{edges shown}]$

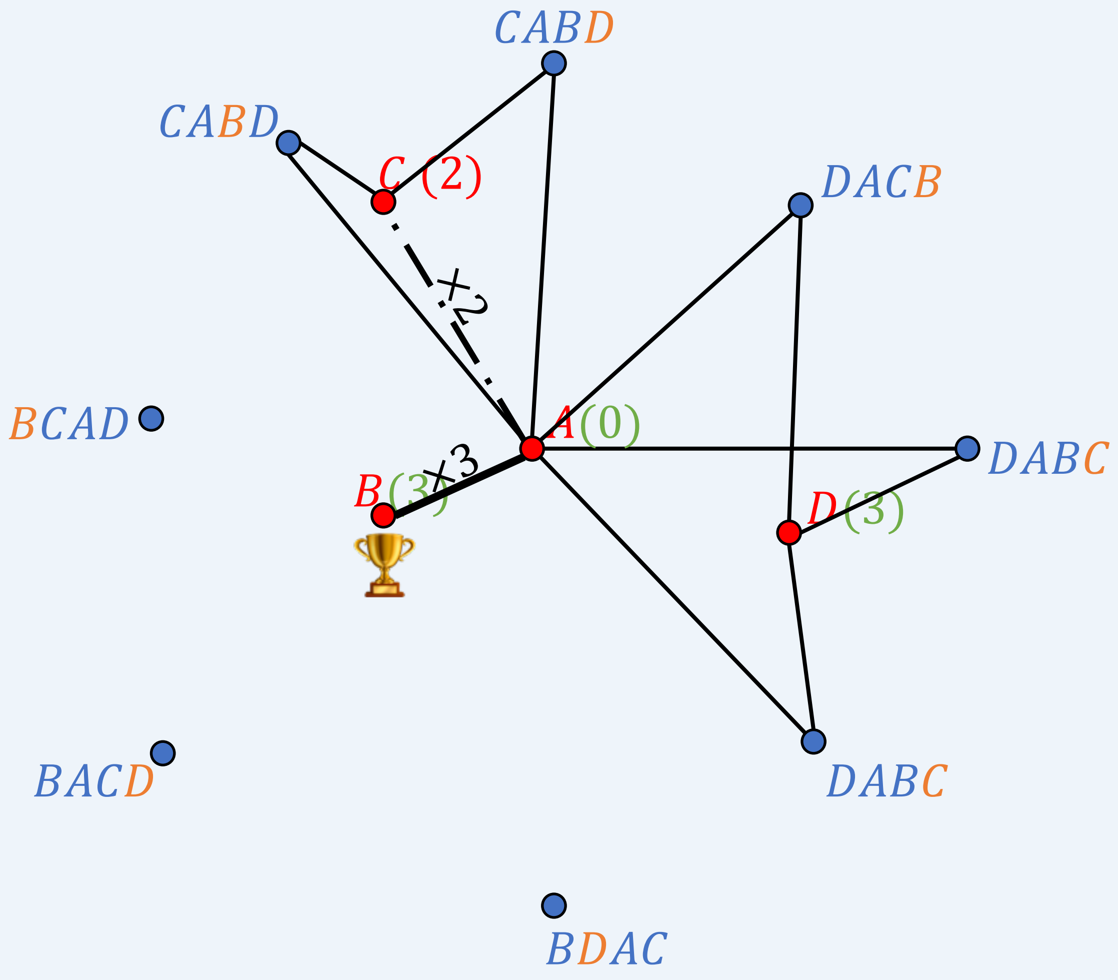

Plurality Veto

- Initially, each candidate X,
  $\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (veto) of least favorite candidate with positive score
- Last candidate vetoed wins!

Proof of distortion 3

- Goal: $cost(B) \leq 3 \cdot cost(A)$
- Key observations:
  - Voters closer to B than veto
  - Candidates: # vetos = # first choice votes
- So far…
  $cost(B) \leq cost(\mathrm{A}) + \text{[edges shown]}$

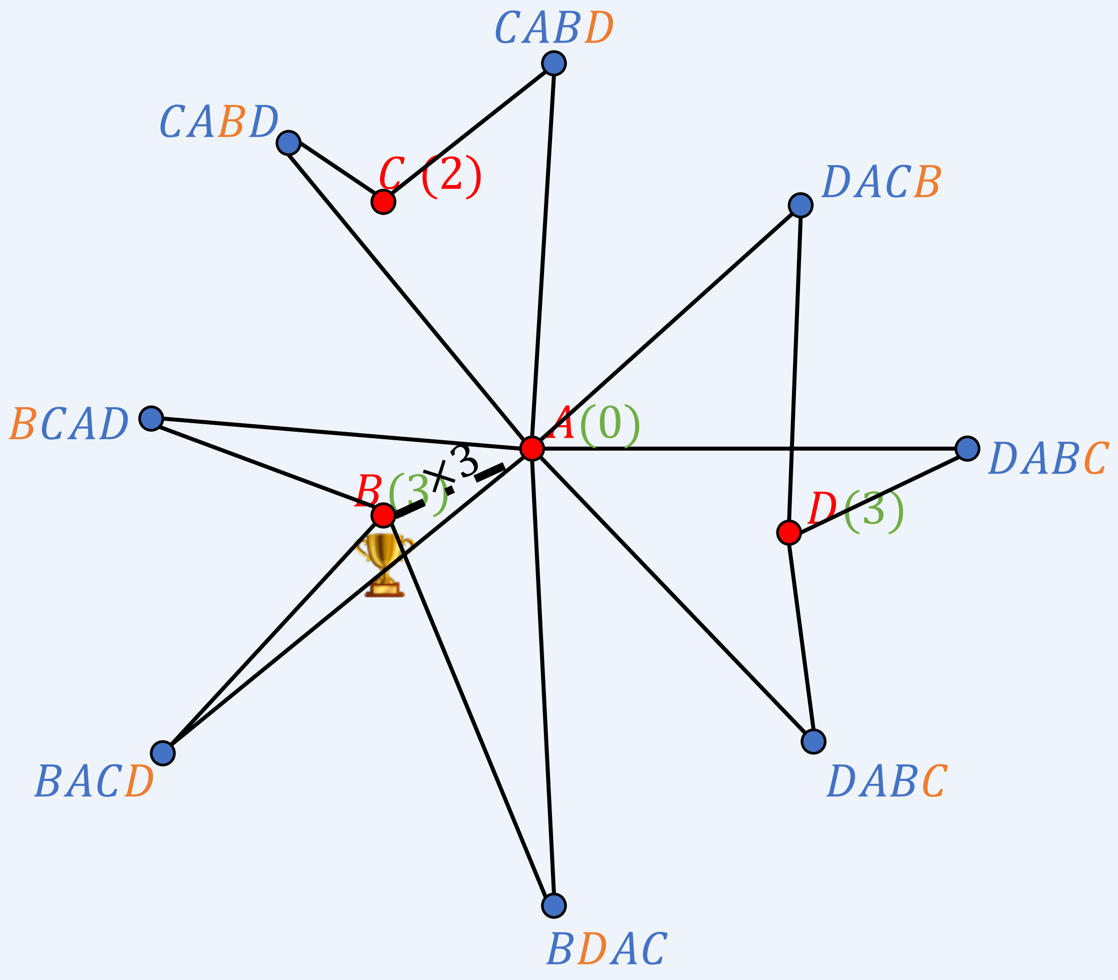

Plurality Veto

- Initially, each candidate X,
  $\mathrm{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (**veto**) of least favorite candidate with positive score
- Last candidate vetoed wins!

Proof of distortion 3

- Goal: $cost(B) \leq 3 \cdot cost(A)$
- Key observations:
  - Voters closer to B than veto
  - Candidates: # vetos = # first choice votes
- So far…
  $cost(B) \leq cost(\text{A}) + [\text{edges shown}]$

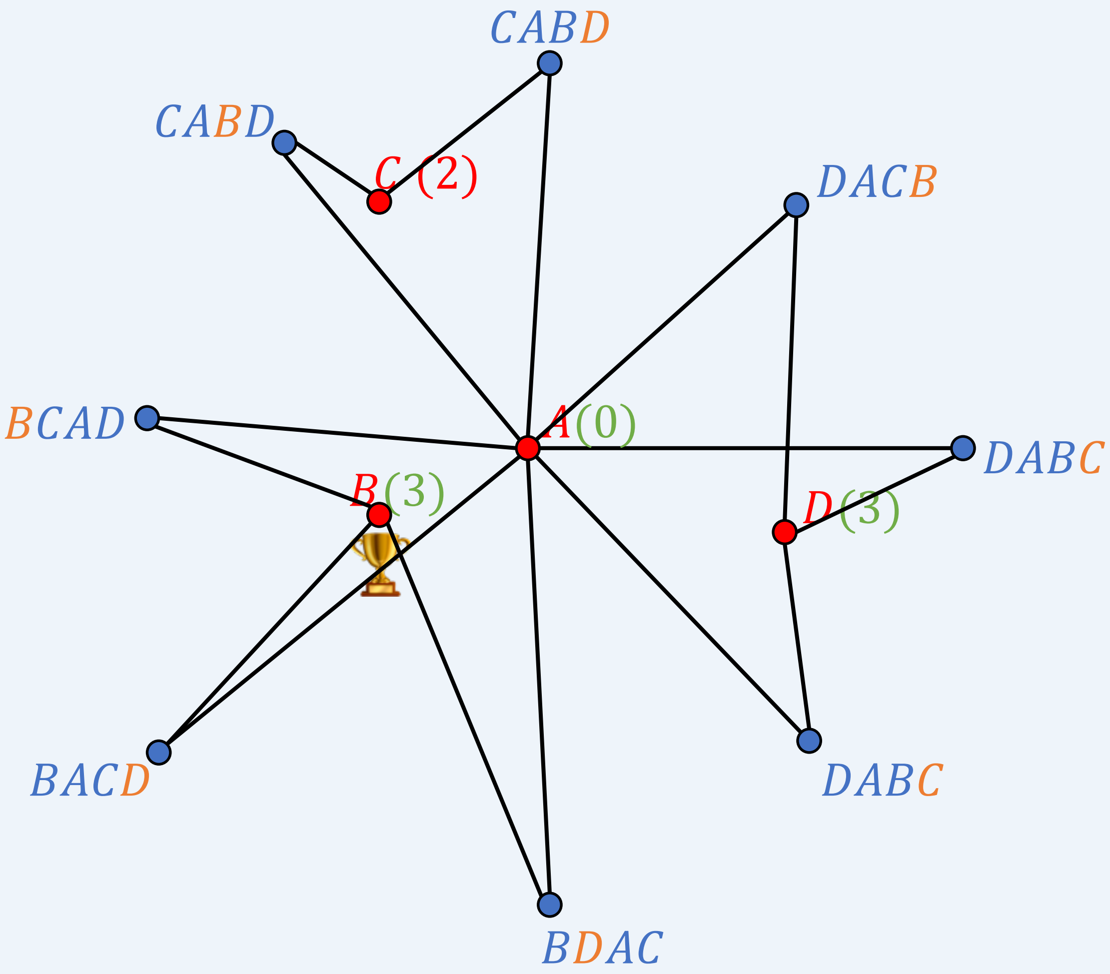

Plurality Veto

- Initially, each candidate X,
  $\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (veto) of least favorite candidate with positive score
- Last candidate vetoed wins!

Proof of distortion 3

- Goal: $cost(B) \leq 3 \cdot cost(A)$
- Key observations:
  - Voters closer to B than veto
  - Candidates: # vetos = # first choice votes
- So far…
$cost(B) \leq 2 \cdot cost(\mathrm{A}) + [\text{edges shown}]$

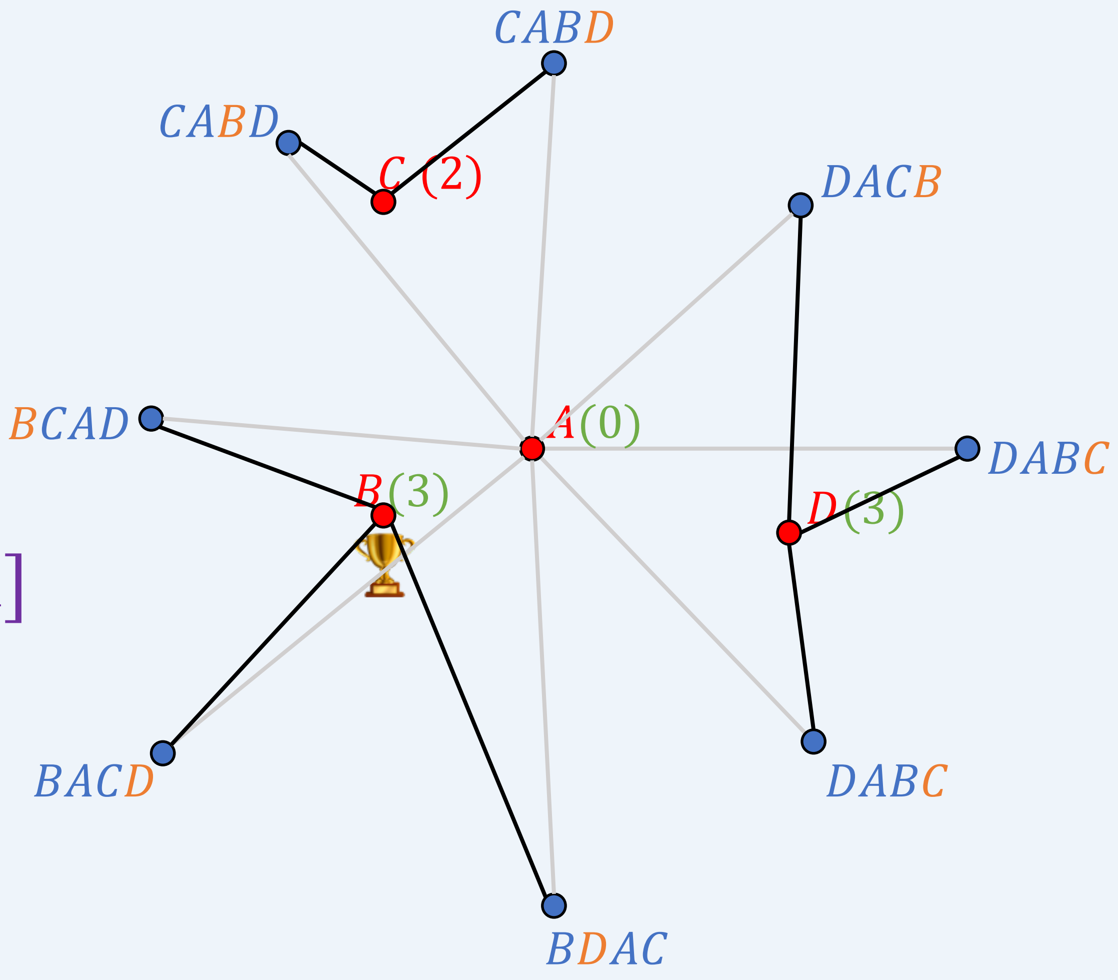

Plurality Veto

- Initially, each candidate X,
  $\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (veto) of least favorite candidate with positive score
- Last candidate vetoed wins!

Proof of distortion 3

- Goal: $cost(B) \leq 3 \cdot cost(A)$
- Key observations:
  - Voters closer to B than veto
  - Candidates: # vetos = # first choice votes
- So far…

$cost(B) \leq 2 \cdot cost(\mathrm{A}) + [\text{edges shown}]$

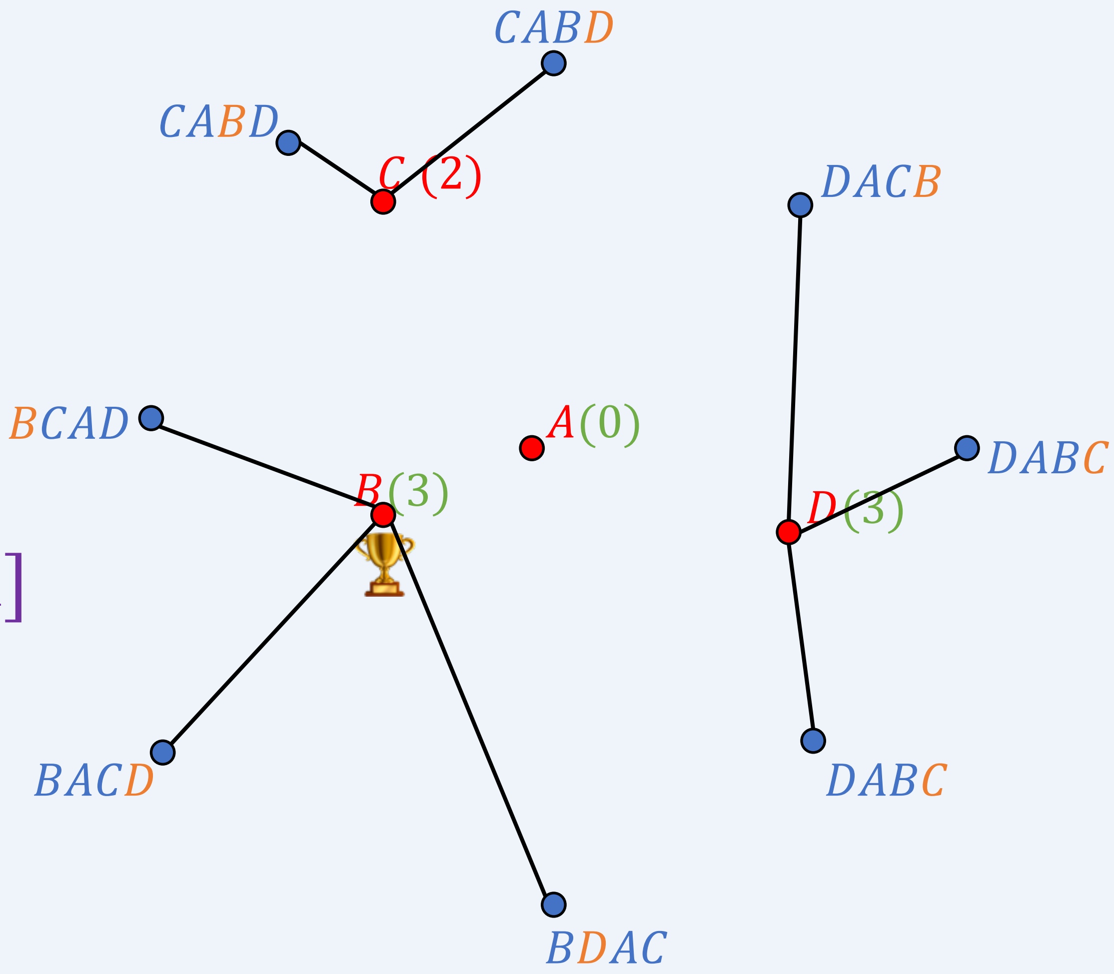

Plurality Veto

- Initially, each candidate X,
  $\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (**veto**) of least favorite candidate with positive score
- Last candidate vetoed wins!

Proof of distortion 3

- Goal: $cost(B) \leq 3 \cdot cost(A)$
- Key observations:
  - Voters closer to B than veto
  - Candidates: # vetos = # first choice votes
- So far…

$cost(B) \leq 2 \cdot cost(\text{A}) + [\text{edges shown}]$

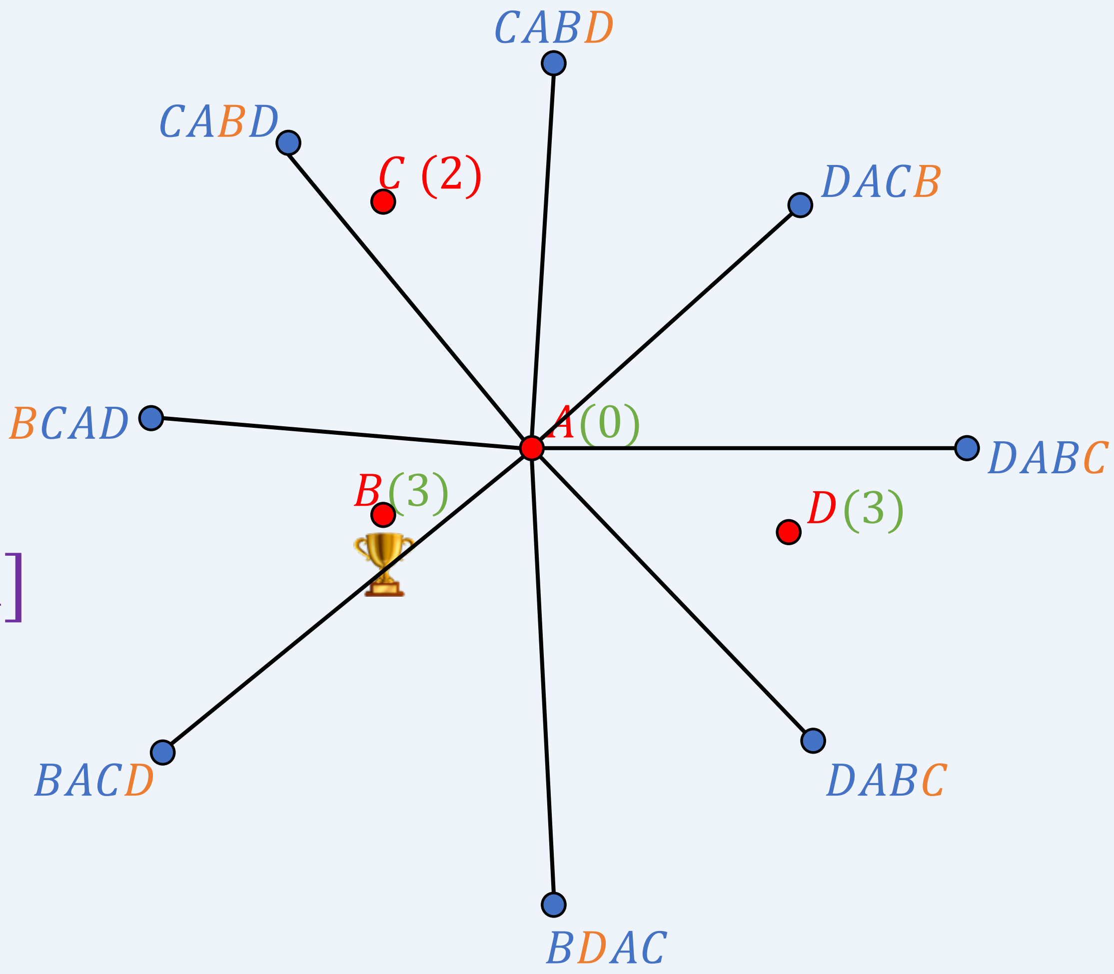

Plurality Veto

- Initially, each candidate X,
  $\text{score}(X) = \text{\# first choice votes for X}$
- One by one, each voter decrements score (**veto**) of least favorite candidate with positive score
- Last candidate vetoed wins!

Randomized Rules

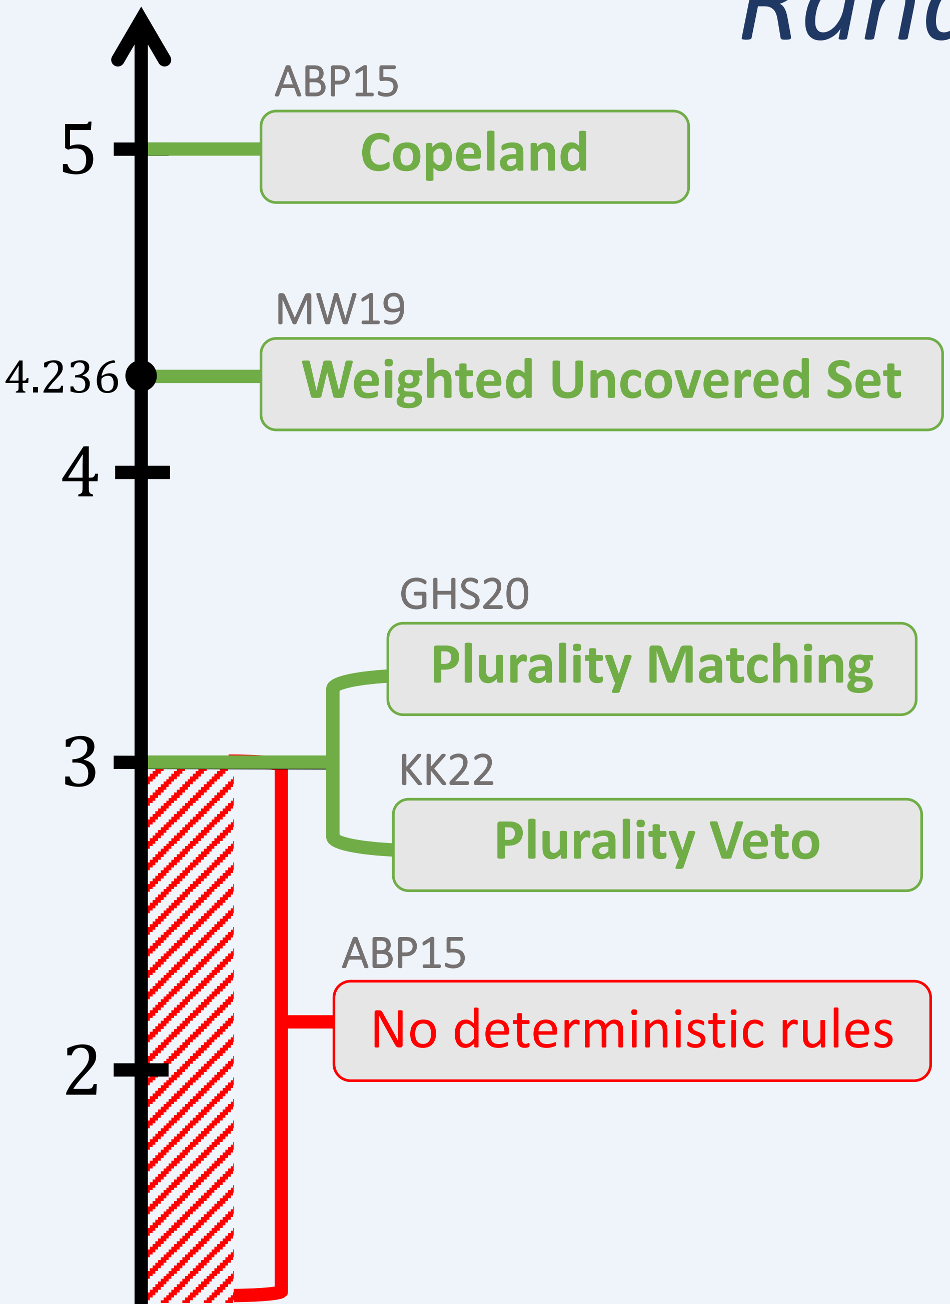

Randomized Rules

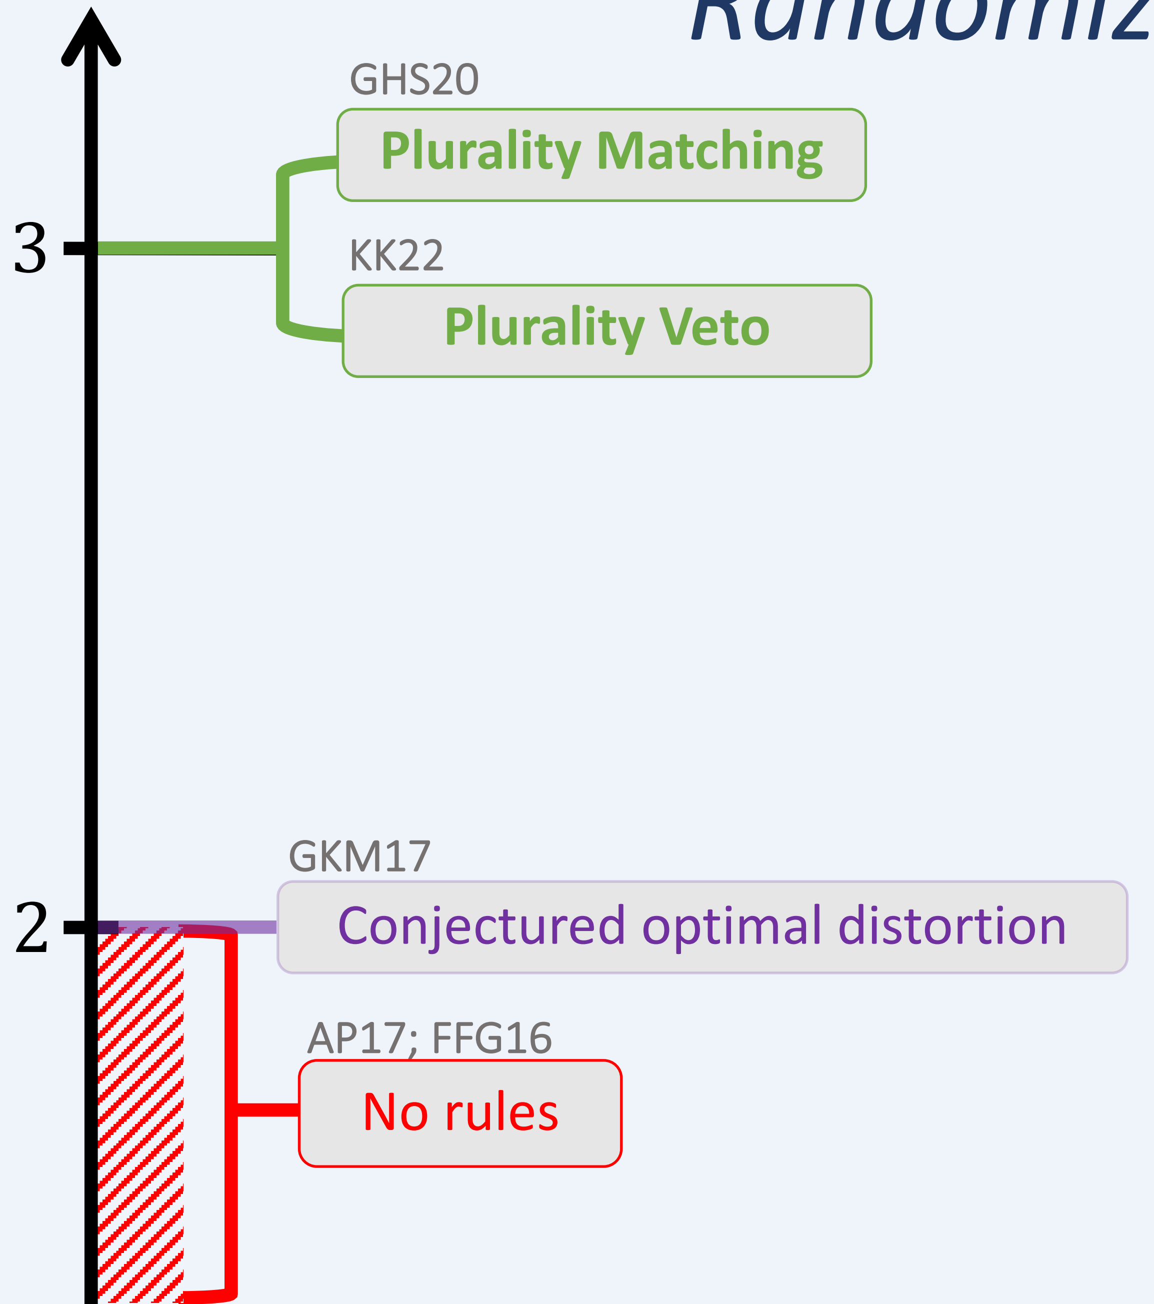

Anshelevich–Postl 2017;

Feldman–Fiat–Golomb 2016

Optimal randomized distortion is ≥ 2

Random Dictator has distortion 3

Randomized Rules

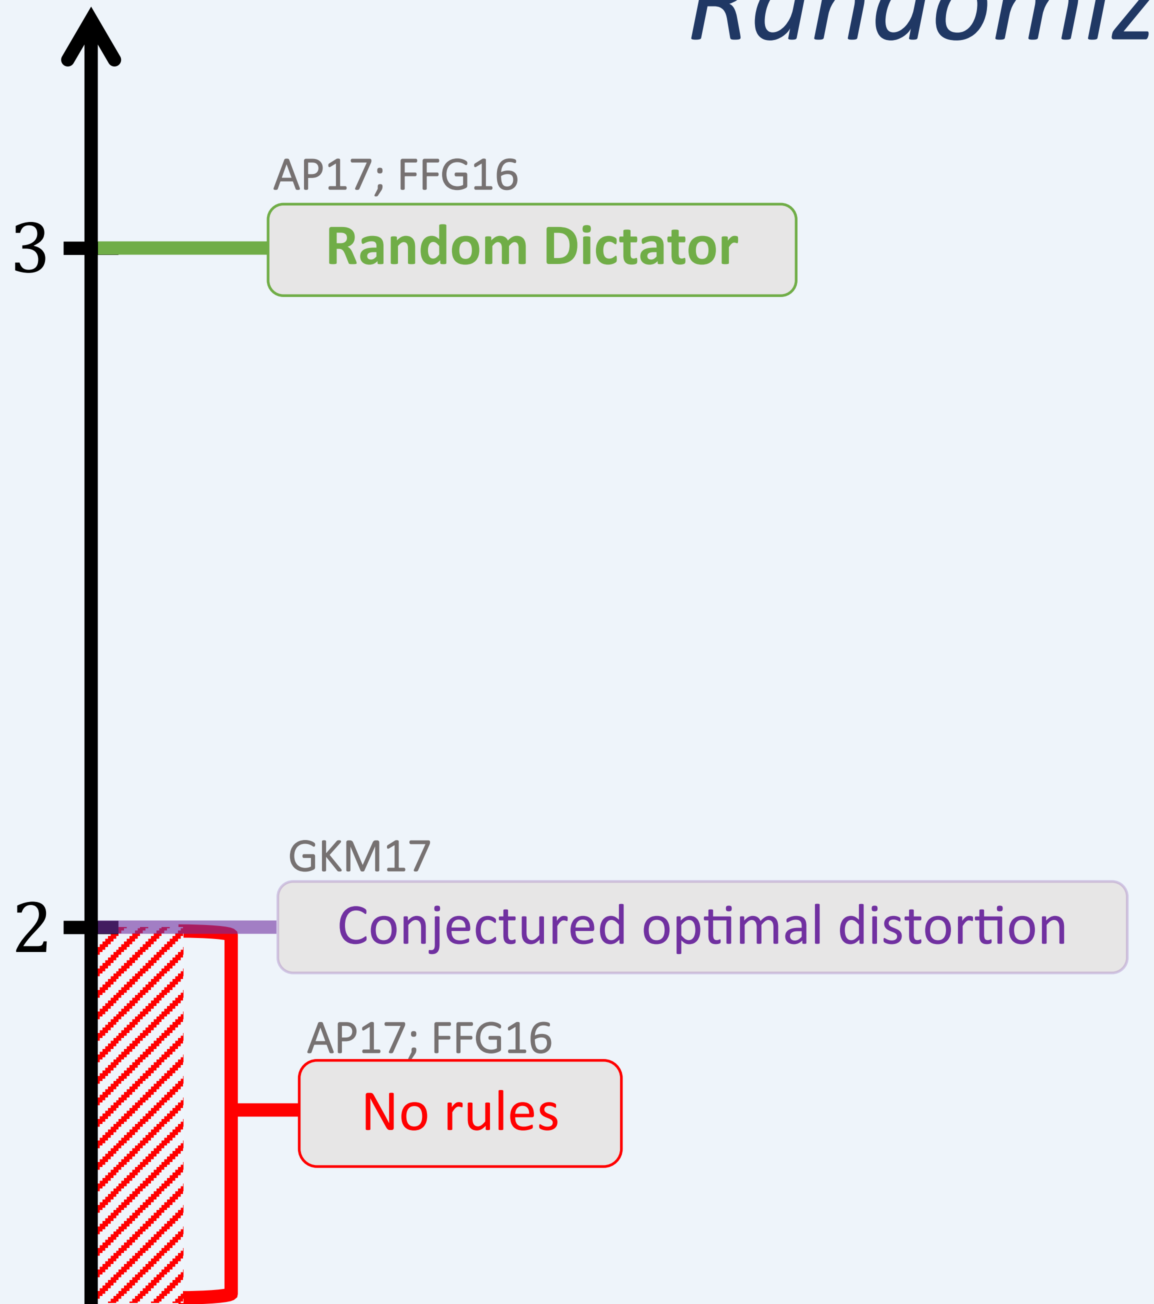

Anshelevich–Postl 2017;

Feldman–Fiat–Golomb 2016

Optimal randomized distortion is ≥ 2

Random Dictator has distortion 3

- Random voter's favorite candidate

Randomized Rules

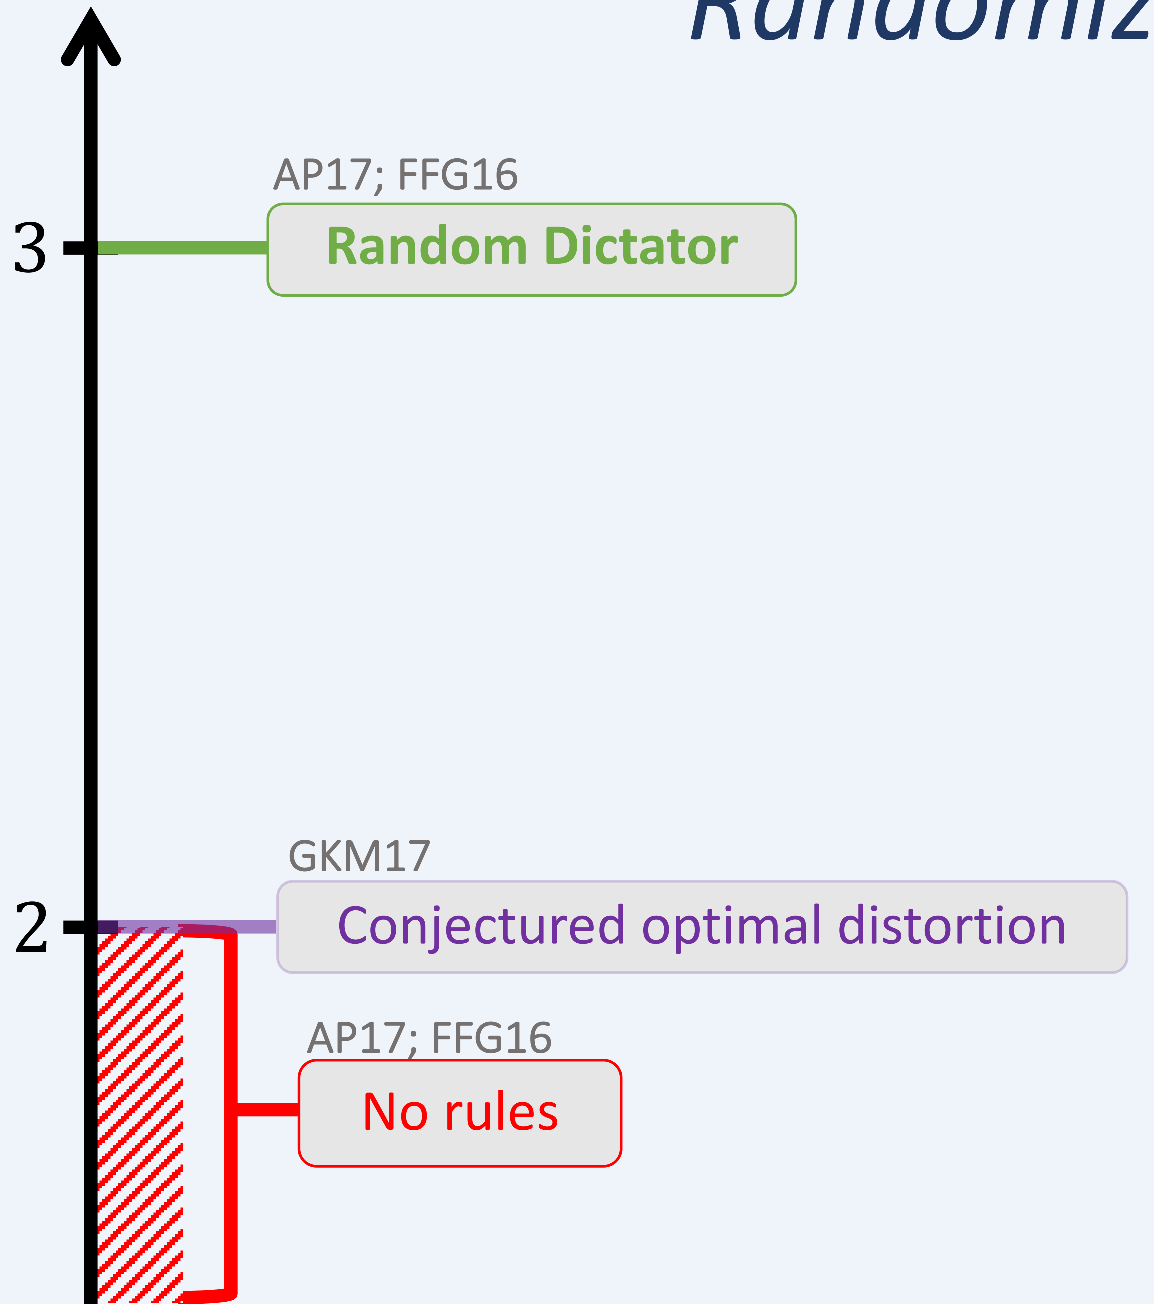

Surprisingly difficult to improve!

- Natural classes of rules fail
  - Tournament rules (GKM17), e.g., Copeland, Ranked Pairs, Schulze, Maximal Lotteries, (Weighted) Uncovered set…
  - Top O(1) choices of voters (GAX17), e.g., Plurality, Plurality Veto, Single Transferable Vote, Random Dictator, Smart Dictator, Proportional to Squares…
- Improvements only in restricted settings
  - Few voters or candidates (AP17, FGMP19, Kem20, GHS20)
  - Restricted metrics/more information (FFG16, FGMS17, BFGT23)

What's the right answer?

Randomized Rules

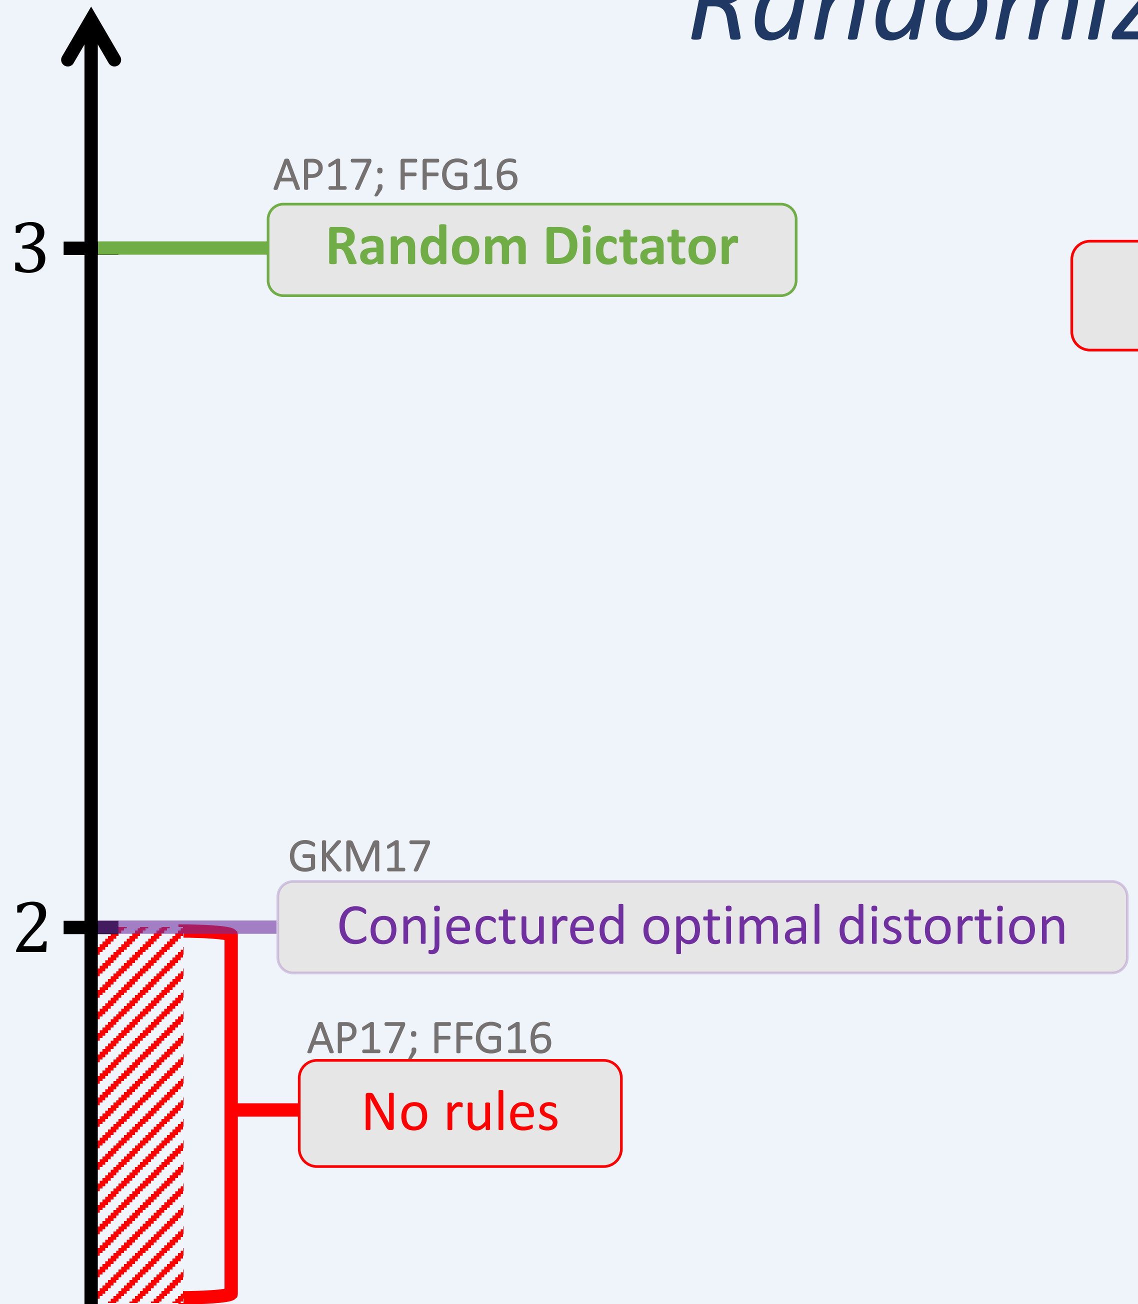

Charikar–R. 2022

Optimal randomized distortion is $\geq \mathbf{2.112}$

Randomized Rules

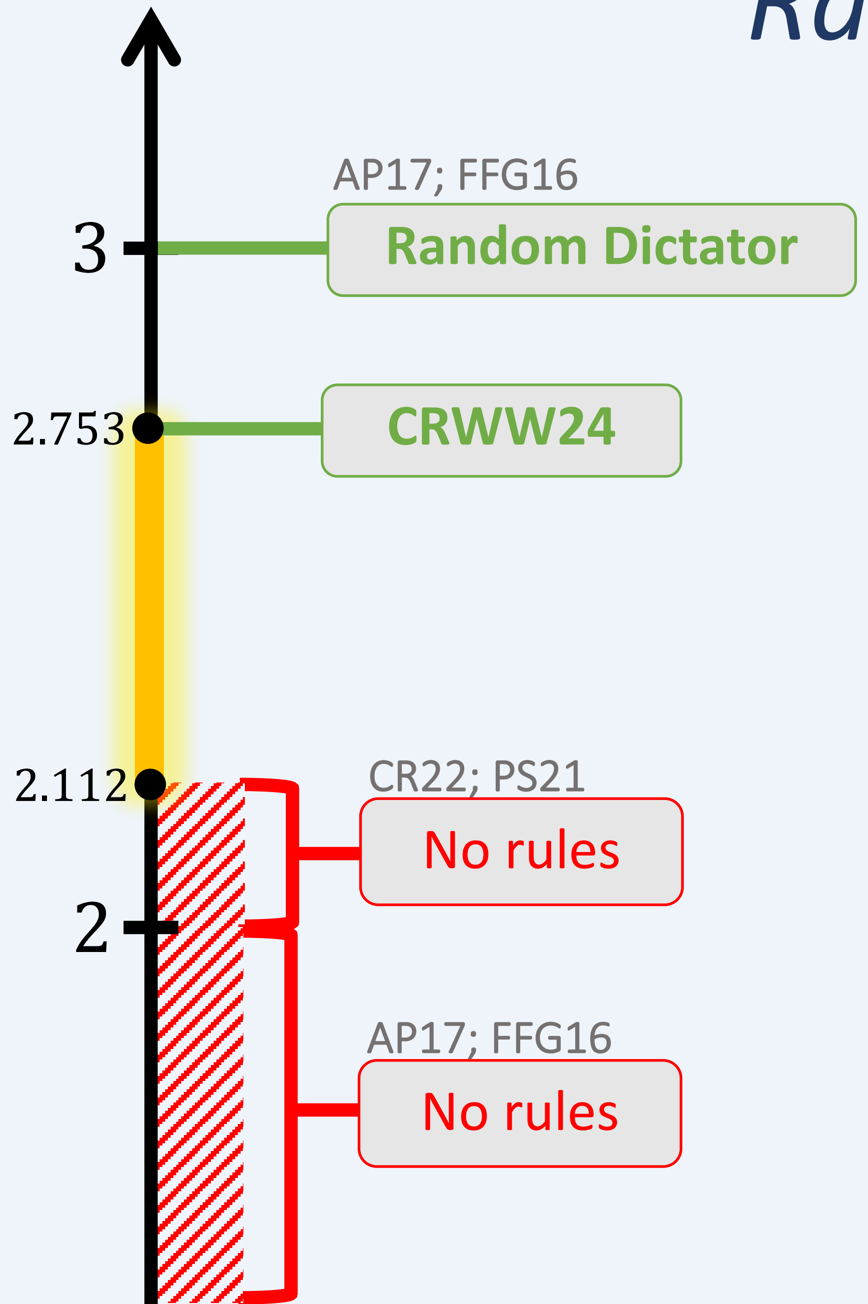

Charikar–R. 2022

Optimal randomized distortion is $\geq \mathbf{2.112}$

- Independently, Pulyassary–Swamy got 2.063

Charikar–R.–Wang–Wu 2024

New randomized rule with distortion $\leq \mathbf{2.753}$

- Powerful techniques to determine distortion precisely
- Uses simple voting rules!

Beating distortion 3

Part 1: Maximal Lotteries

Alice and Bob observe an election and play a game:

- Each picks a candidate
- Random voter is chosen
- Winner is whose candidate the voter prefers

What is the best strategy?

- Symmetric zero sum game
- Exists mixed-strategy Nash equilibrium
- Voting rule: use equilibrium distribution

Voter	Ranking
1	$CABD$
2	$DACB$
3	$DABC$
4	$DABC$
5	$BDAC$
6	$BACD$
7	$BCAD$
8	$CBAD$

A

Alice

B

Bob

Part 1: Maximal Lotteries

"Maximal lotteries were first considered by Kreweras (1965) and rediscovered and studied in detail by Fishburn (1984a). … rediscovered again by economists (Laffond et al., 1993), mathematicians (Fisher and Ryan, 1995), political scientists (Felsenthal and Machover, 1992), and computer scientists (Rivest and Shen, 2010)"

Felix Brandt (2017), *Recent Results in Probabilistic Social Choice*

Part 1: Maximal Lotteries

Nicholas de Condorcet

Condorcet's Paradox

Not always a candidate that beats all others

Maximal Lotteries

Always a ***distribution* over candidates** that beats all others

Part 1: Maximal Lotteries

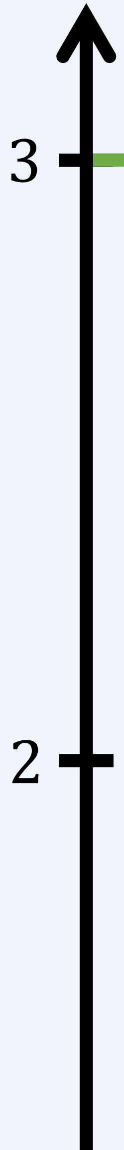

This work

Maximal Lotteries has distortion **3**

- Optimal for tournament rules!
- Key feature of analysis:
  - When ML has high distortion, get precise structure on metric

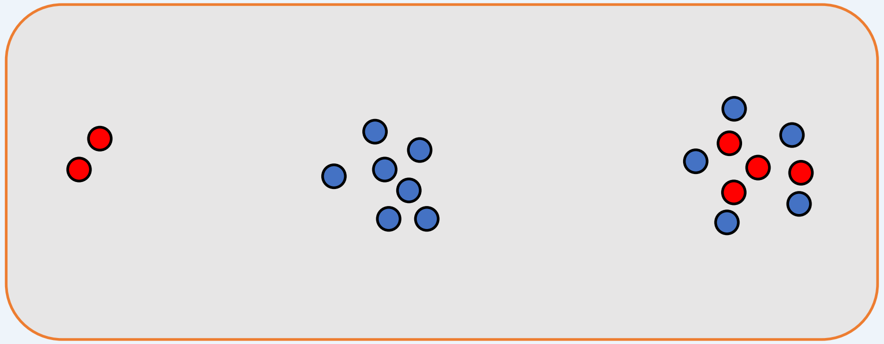

Distortion $\approx 3 \Longrightarrow$ structured metric

Part 2: Random Consensus Builder

Idea: Random Dictator, but strong consensus can overrule

Threshold for overruling: β (say 2/3)

Random voter's ranking:

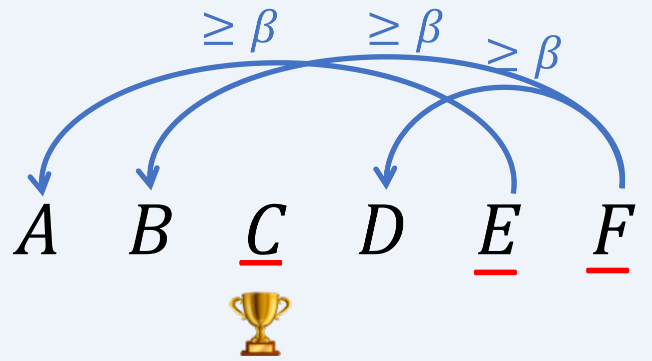

$\beta = 1$ is Random Dictator, $\beta = \frac{1}{2}$ is like Copeland

Distortion interpolates between 3 and 5

ML mixed with RCB

With probability p:

- Run Maximal Lotteries

With probability $1 - p$:

- Randomly pick threshold $\beta \sim [\frac{1}{2}, B]$
- Run Random Consensus Builder

With $p = \frac{1}{\sqrt{2}}, B = \sqrt{2} - \frac{1}{2}$, gets distortion $\mathbf{2\sqrt{2} \approx 2.82}$

To get 2.753, need something a little different...

Part 3: RaDiUS

- Idea: Use consensus to identify "good" shortlist of candidates
- Random voter picks from this set
- "Good set": **weighted uncovered set** from Munagala–Wang 2019

$\beta \in [\frac{1}{2}, 1]$. $\boldsymbol{A}$ such that for all $\boldsymbol{B}$

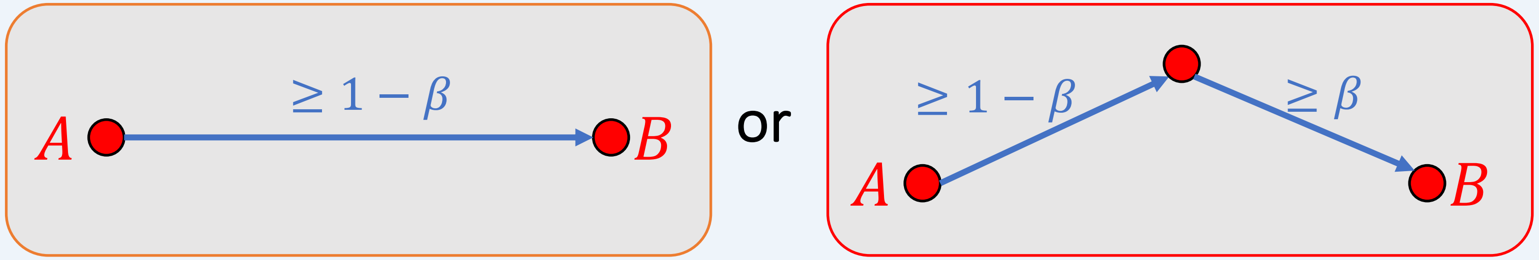

Random Dictator in the (Weighted) Uncovered Set

ML mixed with RaDiUS

With probability p:

- Run Maximal Lotteries

With probability $1-p$:

- Randomly pick threshold $\beta \sim [\frac{1}{2}, B]$ **according to pdf $\boldsymbol{\rho}(\cdot)$**
- Run **RaDiUS**

With $\rho(\beta) = \frac{p}{(1-p)(1-\beta^2)}$, $p = \frac{1}{1+\int_{\frac{1}{2}}^{B} \frac{d\beta}{1-\beta^2}}$, $B = 0.876$, gets distortion $\mathbf{2.753}$

Key Analysis Ideas

Idea 1: Hard Metrics

Voter	Ranking
1	AB
2	BA

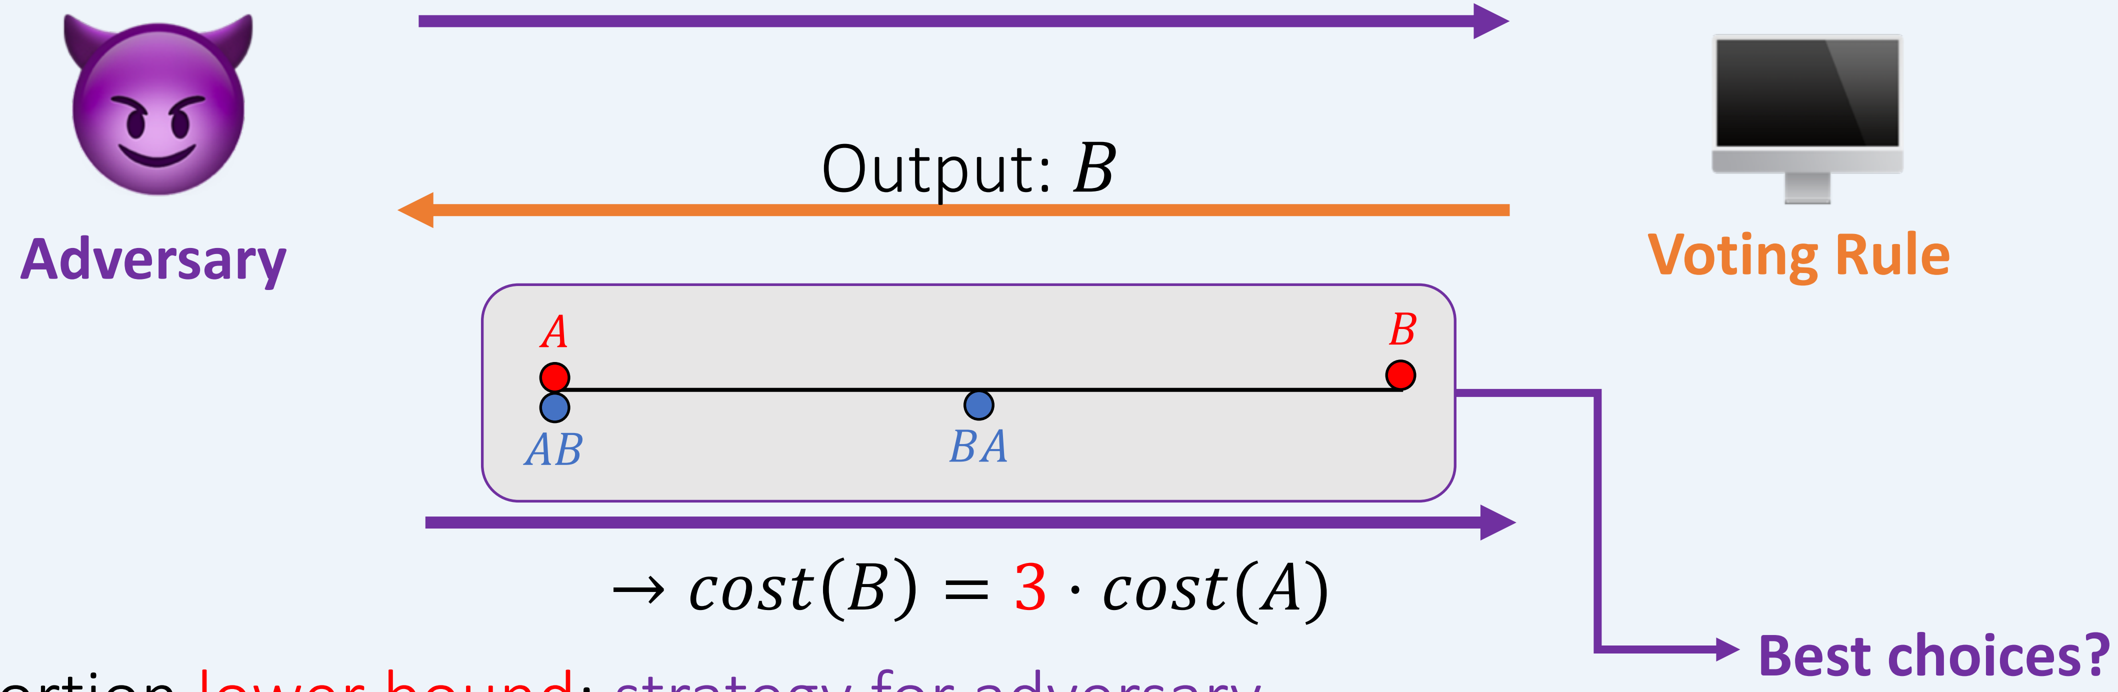

- Distortion lower bound: strategy for adversary
- Distortion upper bound: strategy for voting rule

Idea 1: Hard Metrics

- Make metric worse for rule?
- Keep only dists from candidates to $\boldsymbol{A}$
- Make voter dists to A minimal
- Make voter dists to others maximal

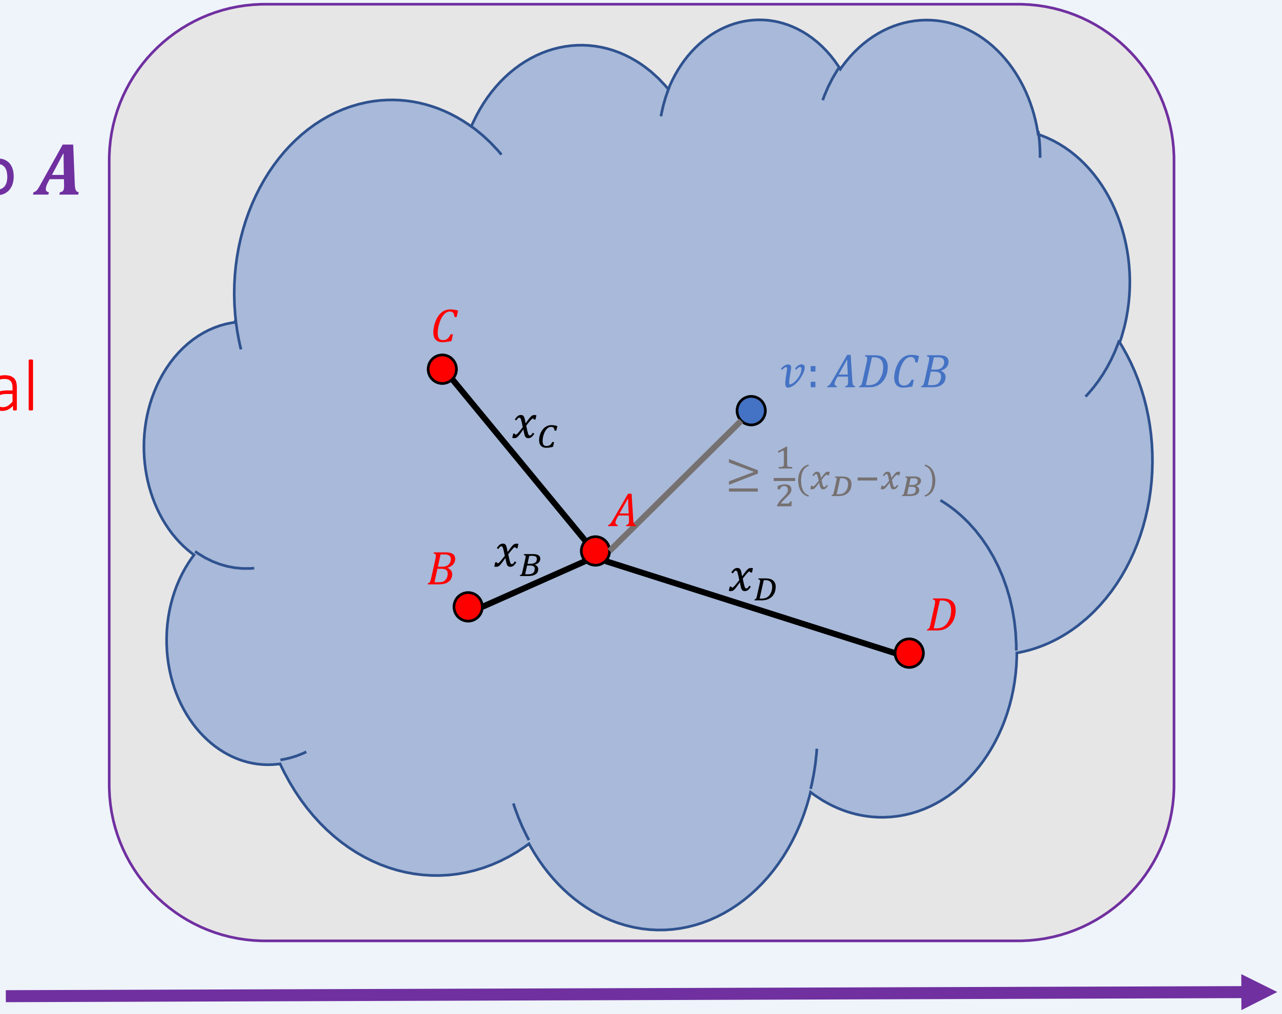

True opt: A

Idea 1: Hard Metrics

- Make metric worse for rule?
- Keep only dists from candidates to $\boldsymbol{A}$
- Make voter dists to A minimal
- Make voter dists to others maximal

Key observations:

- Changes increase distortion!
- Distances define a metric!
- $\Longrightarrow$ Can assume metric looks like this

Notes: factor of 2, difference between costs

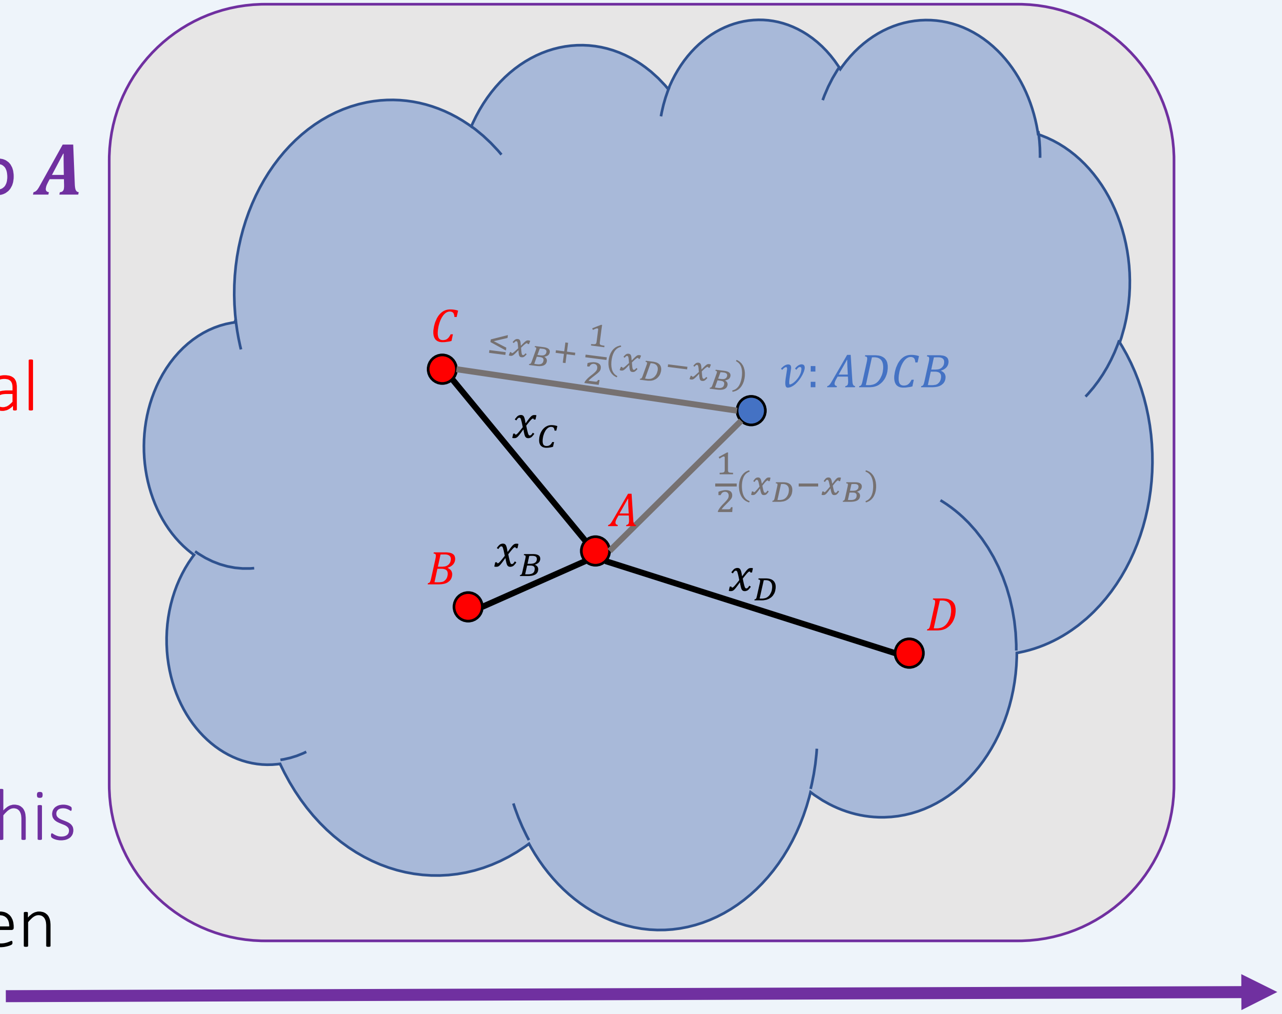

True opt: A

Idea 2: Continuous View of Costs

$$2 \cdot cost(A) = \frac{1}{n} \sum_{v \in V} 2d(v, A)$$

$$= \mathbb{E}_{v \sim V}\left[2d(v, A)\right]$$

$$= \int_0^\infty \Pr_{v \sim V}\left[2d(v, A) > t\right] \mathrm{d}t$$

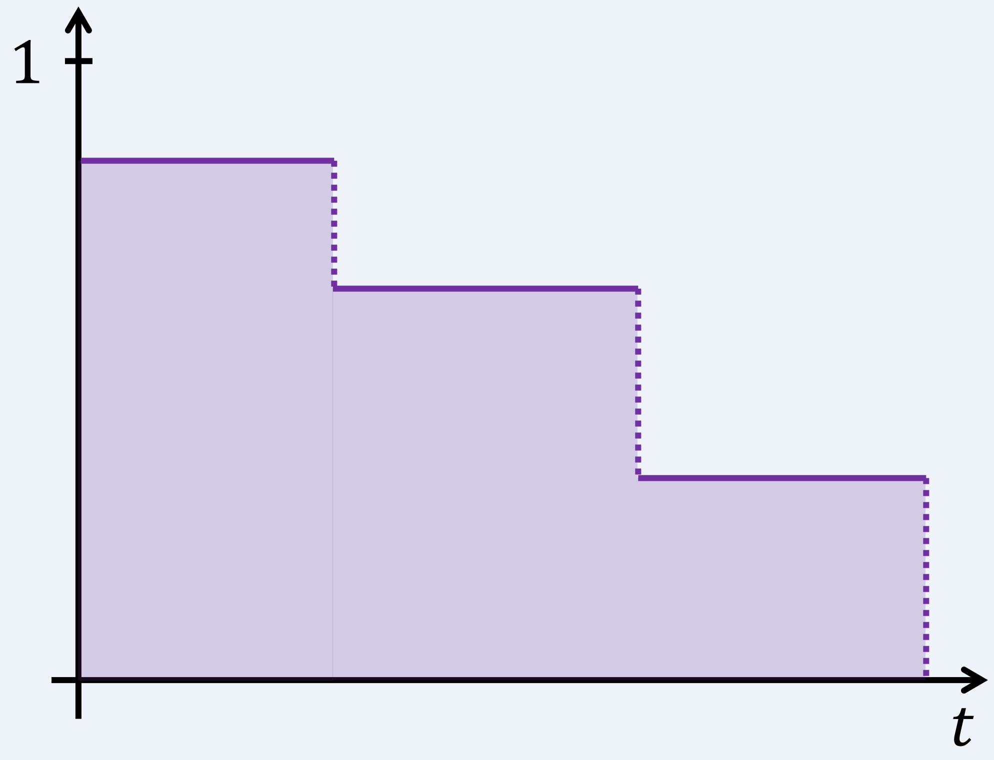

Idea 2: Continuous View of Costs

$$2 \cdot cost(A) = \int_0^\infty \Pr_{v\sim V}[2d(v,A) > t]\,\mathrm{d}t$$

$$cost(B) - cost(A) = \mathbb{E}_{v\sim V}[d(v,B) - d(v,A)]$$

$$= \int_0^\infty \Pr_{v\sim V}[d(v,B) - d(v,A) > t]\,\mathrm{d}t$$

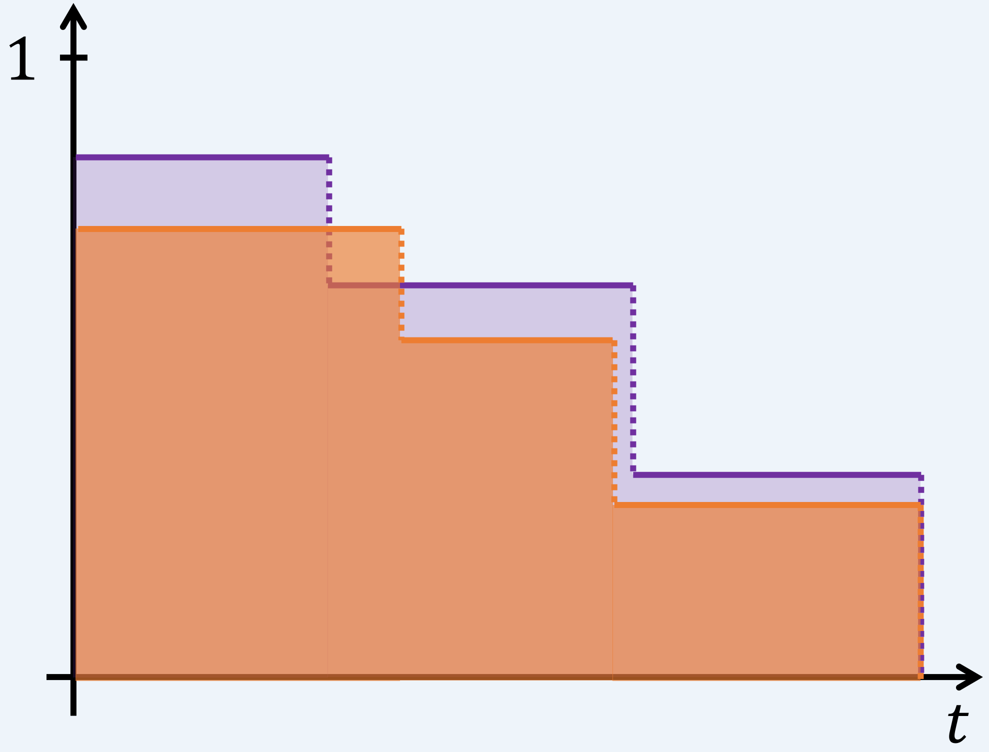

Distortion 3: orange < purple

Idea 2: Continuous View of Costs

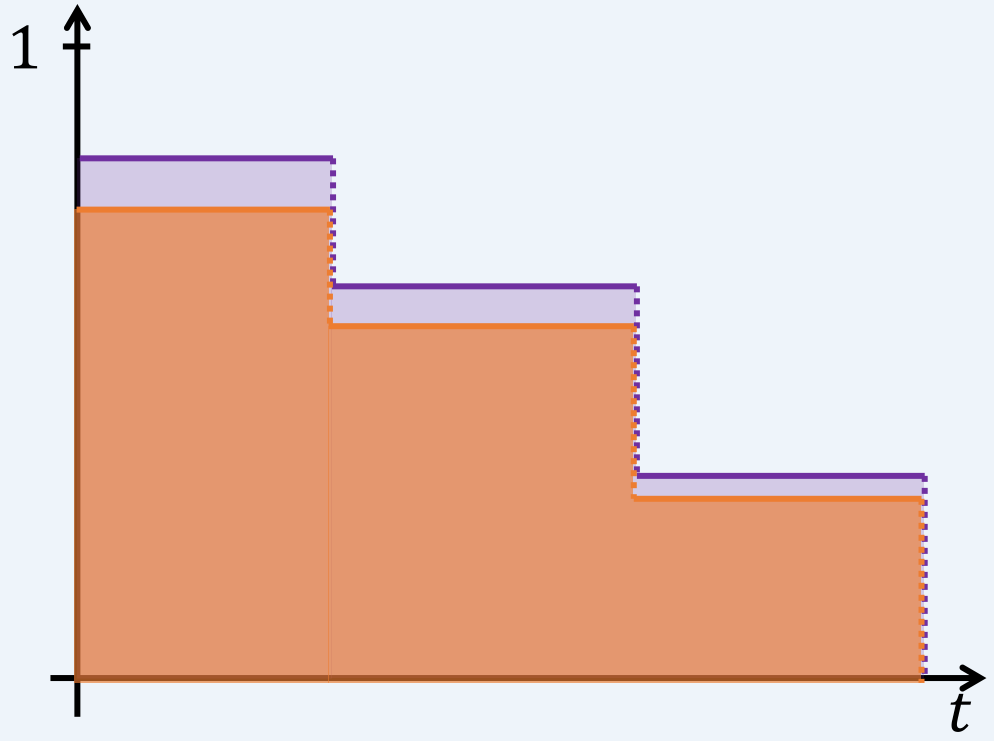

Proofs for Random Dictator, Plurality Matching, Plurality Veto

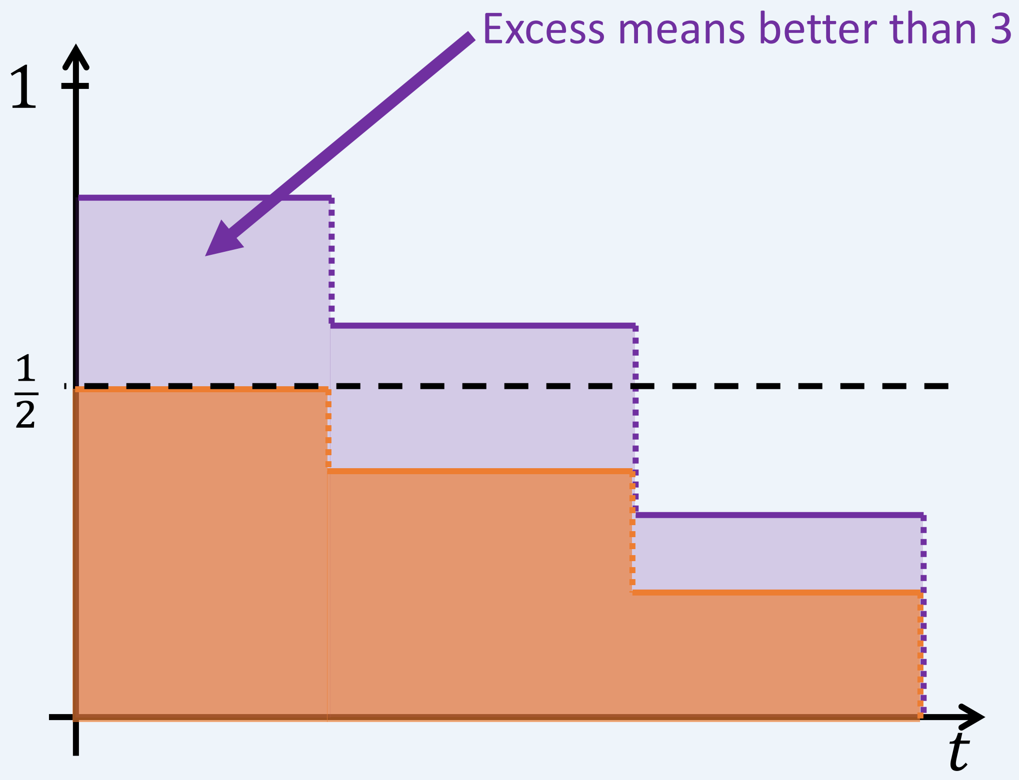

Proof for **Maximal Lotteries**

Idea 2: Continuous View of Costs

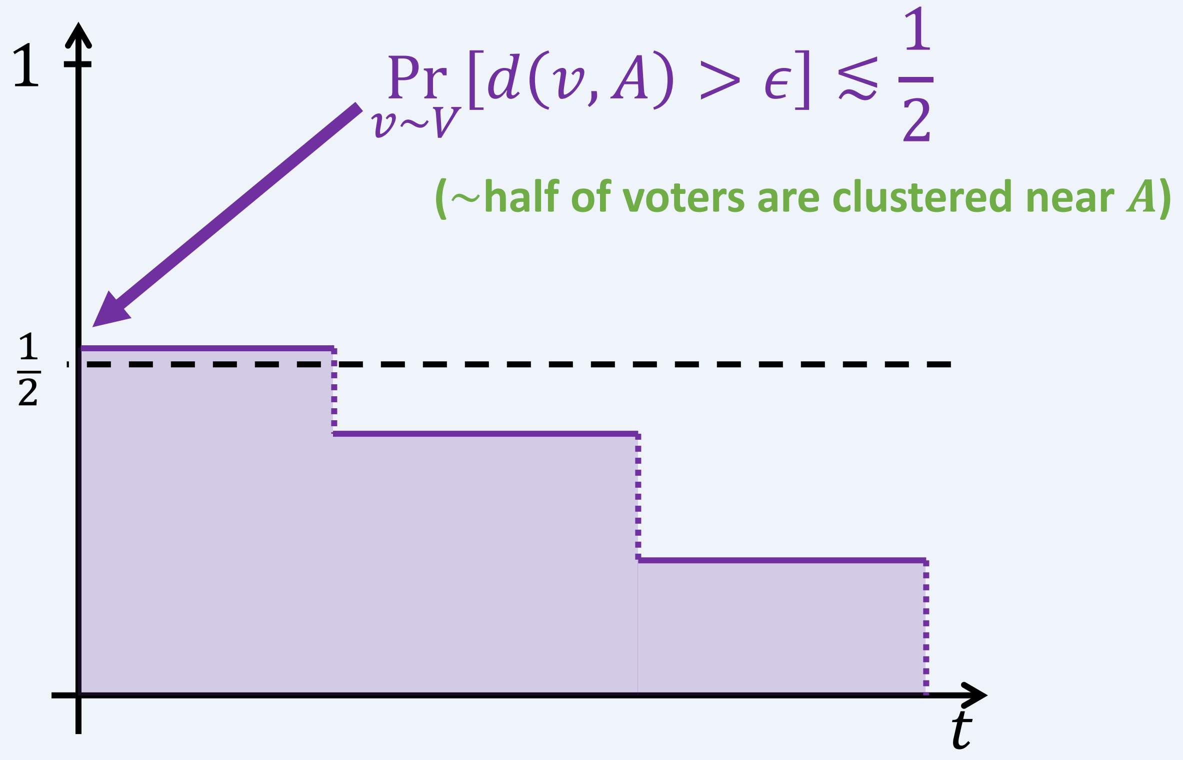

Hard instances for **Maximal Lotteries**

Open Problems

State of the Art (Voting)

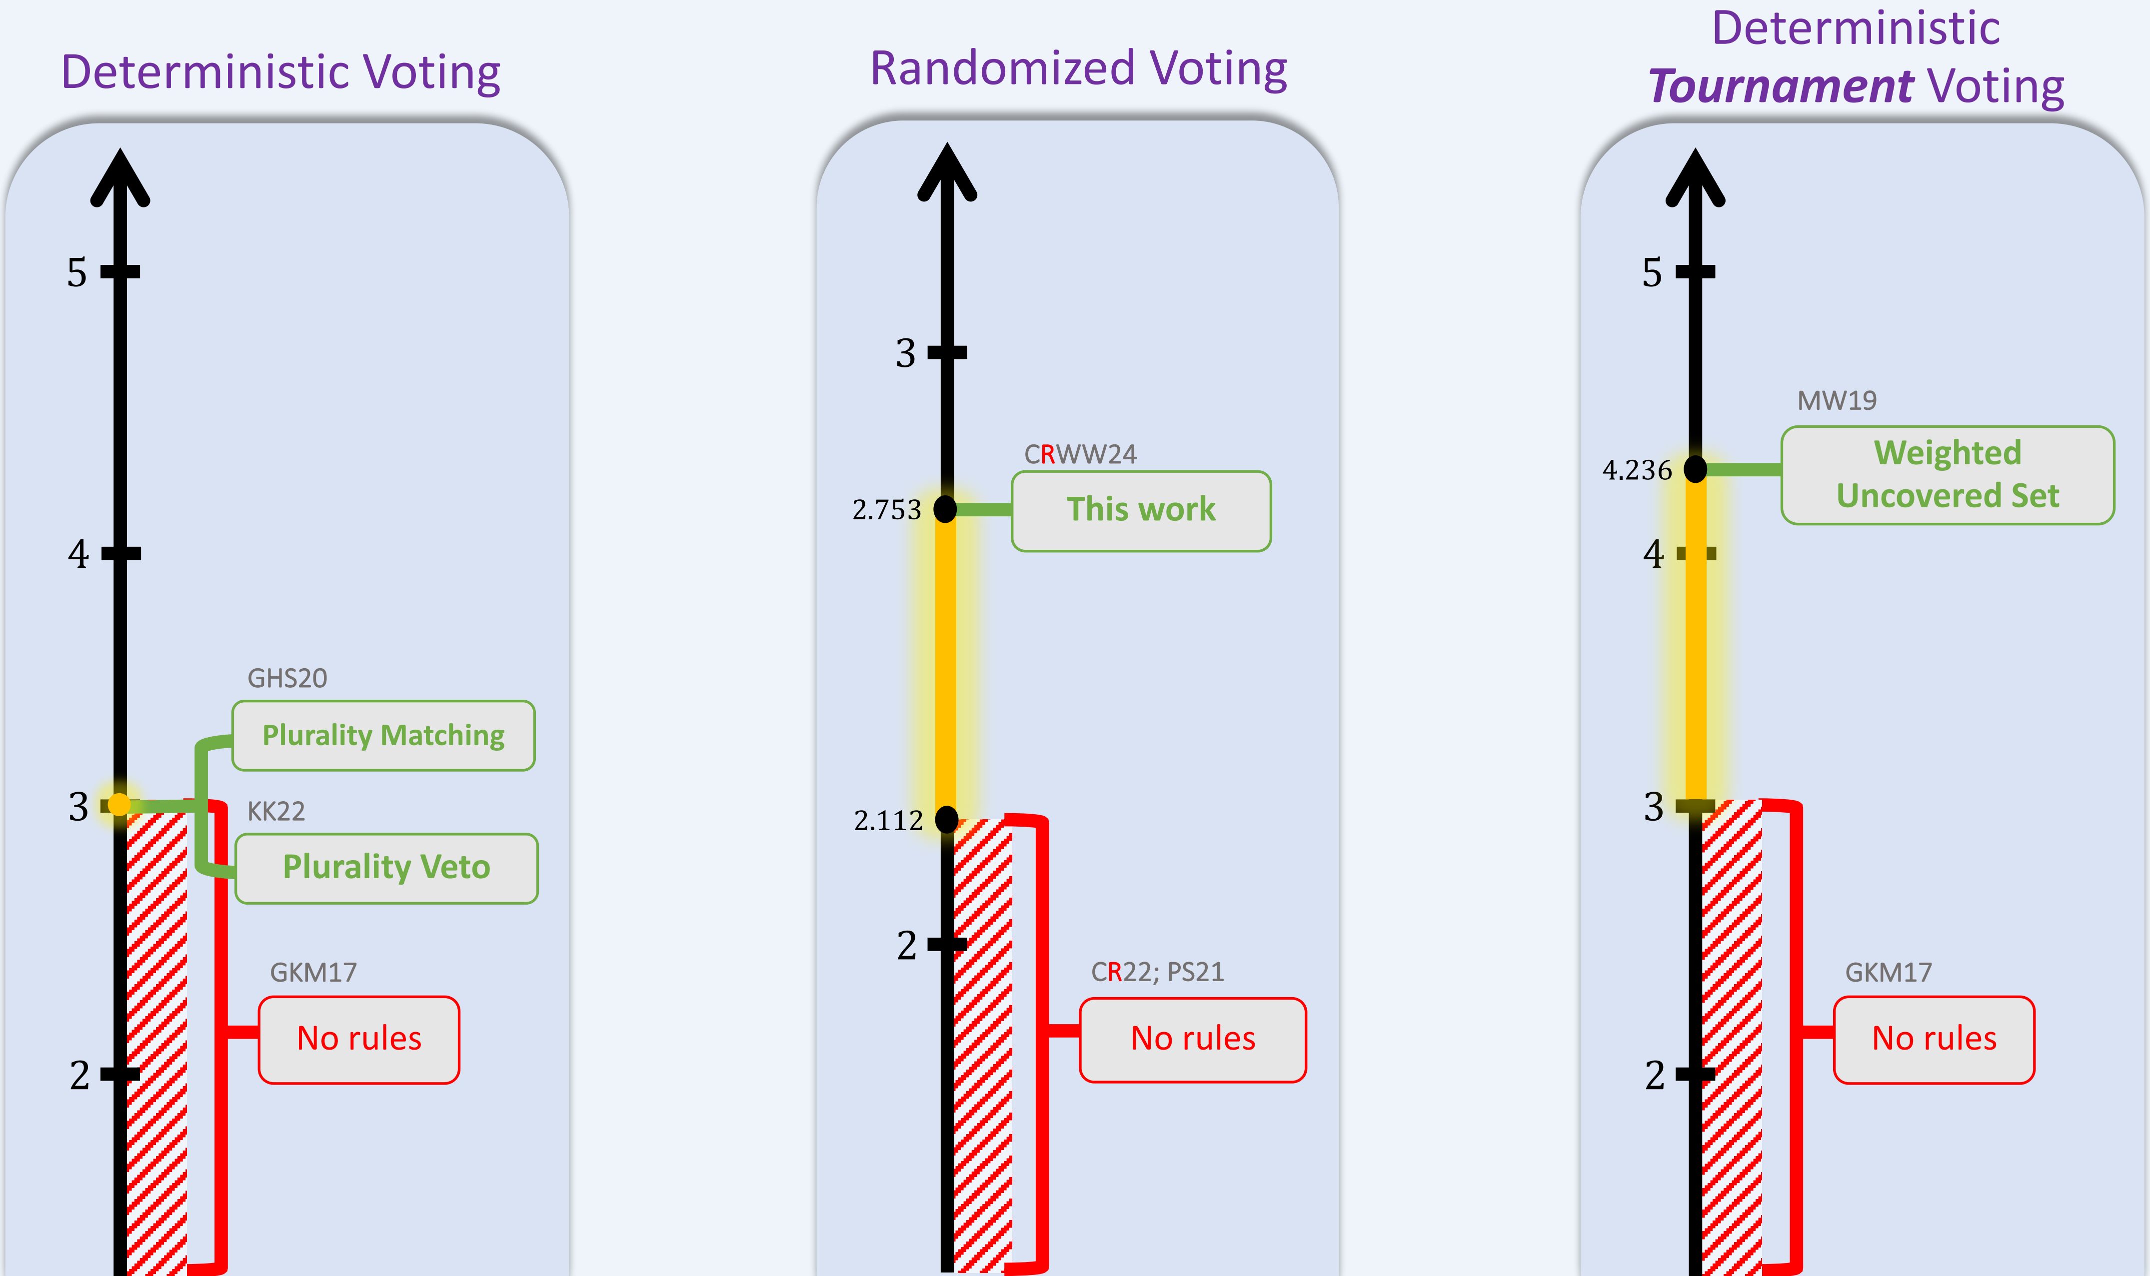

Takeaways

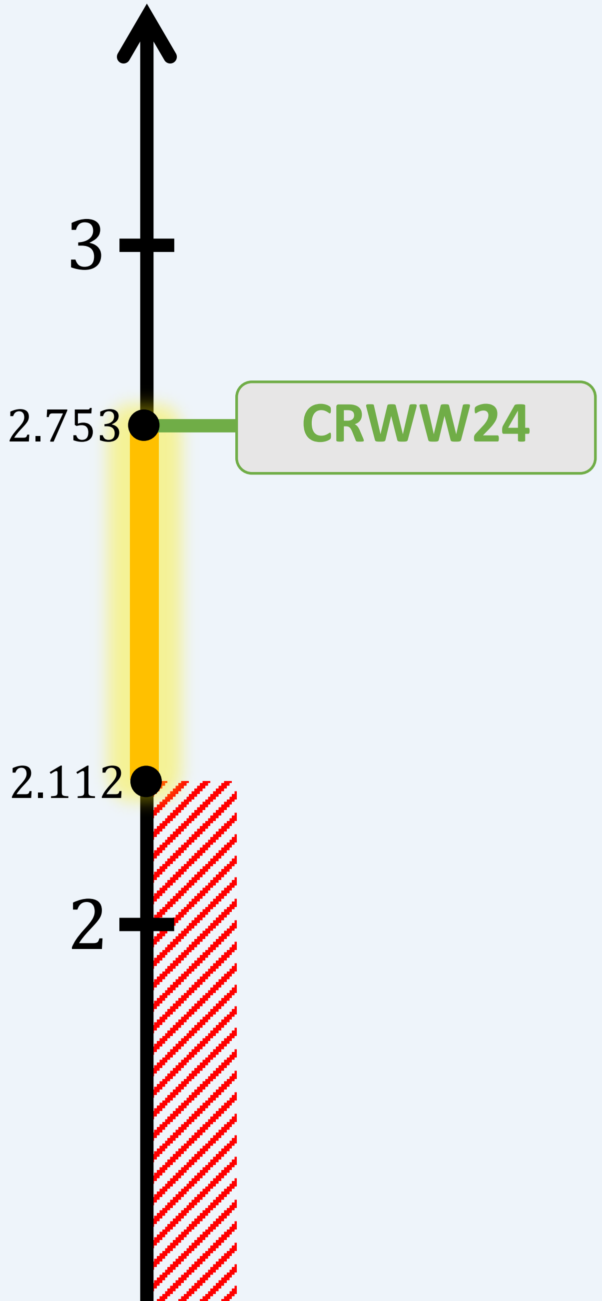

- Metric distortion lens is a new way to deepen our understanding of existing voting rules
- Clean mathematical framework motivates creation of interesting new rules
- Big problems still wide open: could there be a *simple* randomized rule with **optimal distortion**?
  - *Distortion in Social Choice Problems: The First 15 years and Beyond* (Anshelevich–Filos-Ratsikas–Shah–Voudouris 2021)
  - Our paper!

Thank You!

$$2.753 \approx 3 - \max_{B \in [\frac{1}{2}, 1]} \frac{4 \ln \frac{2}{3} + 4 \ln(1 + B)}{2 - \ln 3 + \ln \frac{1+B}{1-B}}$$

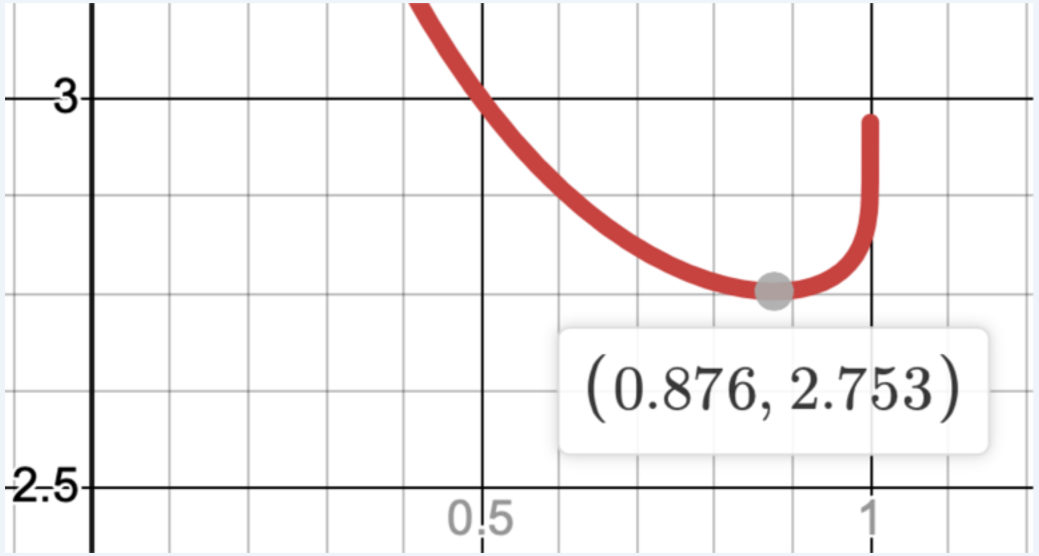

$$2.112 \approx$$

$$1 + \max_{\substack{a,b,c \in (0,\frac{1}{2}) \\ a+b+c \geq 1}} \frac{2(1-2c(1-c))}{\frac{-b(1-b)+(1-a)(1-2c(1-c))}{a} + \frac{(1-b)^2 - 2c(1-c)(1-b) + b(1-2c)(a+b)}{1-a}}$$