

Multiwinner Voting

Edith Elkind

University of Oxford
Alan Turing Institute



Voting: The Basics

- Input:
 - a set of alternatives (candidates) $C = \{c_1, \dots, c_m\}$
 - a set of voters $V = \{1, \dots, n\}$
 - for each voter, their preferences over C
 - ranked ballots: each ballot is a total order over C
 - approval ballots: each ballot is a subset of C
- Output:
 - a winner
 - a set of winners
 - a total ranking of the alternatives

What Does (Computational) Voting Theory Study?


- Catalogue and formalise existing voting rules
- Analyse their computational complexity
- Formulate desiderata (axioms) for voting rules
- Determine which rules satisfy which axioms
- Characterise rules by axioms
- Propose new voting rules that are
 - tractable
 - satisfy useful axioms

Voting: Examples

- What movie should the Simpson family watch?



: Frozen > Paddington > Minions



: Paddington > Minions > Frozen



: Minions > Paddington > Frozen



Voting: Examples

- Which DPhil applicant(s) should the algorithmic game theory group at Oxford accept?
 - Paul: $X > Y > Z > T$
 - Elias: $Y > X > Z > T$
 - Edith: $Z > Y > T > X$



Voting: Examples

- United Kingdom (specific precinct)
 - candidates: Conservatives (C), Labour (L), Liberal Democrats (LD)
 - 60 000 voters
 - 25 000 voters prefer C to LD to L: $C > LD > L$
 - 20 000 voters: $L > LD > C$
 - 11 000 voters: $LD > L > C$
 - 4 000 voters: $LD > C > L$

Example: Competition for a Fellowship

- Candidates: 50 students
- Voters: 15 panel members
 - each panel member reports a ranking of the candidates (or perhaps all deserving candidates)
- Goal: select 10 students who will get a fellowship

Example: Ranking of the Universities

- A **panel of experts** is supposed to rank UK universities
 - Expert 1: Cambridge > Oxford > UCL > Imperial
 - Expert 2: Oxford > Cambridge > Imperial > UCL
 - Expert 3: UCL > Cambridge > Oxford > Imperial
 - Expert 4: Oxford > Imperial > Cambridge > UCL
 - Expert 5: Imperial > Cambridge > UCL > Oxford
- Goal: produce a total ranking of 4 universities

Example: Ranking of the Universities

- A panel is supposed to rank UK universities
- Rankings are based on **5 different criteria**:
 - reputation ranking
 - grant income
 - student satisfaction
 - number of research papers published
 - average salary after graduation
- Rankings:
 - **criterion 1**: Cambridge > Oxford > Imperial > UCL
 - **criterion 2**: Oxford > Cambridge > UCL > Imperial
 - **criterion 3**: UCL > Cambridge > Oxford > Imperial
 - **criterion 4**: Oxford > Imperial > Cambridge > UCL
 - **criterion 5**: Imperial > Cambridge > UCL > Oxford
- Should all criteria have the **same weight**?

Warm-up:
single-winner voting rules
(ranked ballots)



Plurality

- Plurality:
 - each voter names his favorite candidate
 - candidates with the largest number of votes win
 - if two or more candidates get the highest score, the winner is chosen using some tie-breaking rule
 - but we can also think of Plurality as outputting the set of candidates tied for winning
- For 2 candidates,
Plurality selects the majority winner

Political Voting

- United Kingdom (specific precinct)
 - candidates: Conservatives (C), Labour (L), Liberal Democrats (LD)
 - 60 000 voters
 - 25 000 voters prefer C to LD to L: $C > LD > L$
 - 20 000 voters: $L > LD > C$
 - 11 000 voters: $LD > L > C$
 - 4 000 voters: $LD > C > L$
 - Plurality outcome: C wins with 25000 votes

Plurality: Good and Bad

- Plurality is obviously the best voting rule if there are only **2** candidates
- Plurality is **anonymous**: only the multiset of votes matters, not the voters' identities
- Plurality is **neutral**: candidates' identities do not matter (unless we use a non-neutral tie-breaker)
- However, for **3** candidates it may behave in an undesirable way
 - e.g., the **majority** of voters may **prefer some other alternative** to the current winner

Plurality: Example Revisited

- United Kingdom elections:
 - 25 000 voters: $C > LD > L$
 - 20 000 voters: $L > LD > C$
 - 11 000 voters: $LD > L > C$
 - 4 000 voters: $LD > C > L$
- Outcome under Plurality:
 - C wins with 25000 votes
- Undesirable property:
 - 31 000 voters prefer L to C ,
 - 35 000 voters prefer LD to C

History of Voting Rules, in One Slide

1. Someone comes up with a voting rule R
2. Someone else identifies a scenario S where this voting rule makes a bad choice
3. (optional) Someone formulates an axiom A regarding choices in S and similar scenarios
4. Someone comes up with a voting rule R' that makes a good choice in S (satisfies A)
(rinse and repeat)

Two-Round Elections

1. All voters vote for their favorite candidate
2. All but the two highest-scoring candidates are eliminated
3. The voters are asked to vote again over the remaining candidates

This rule is known as Plurality with Runoff;
used in France for presidential elections

Plurality With Runoff: Example

- United Kingdom elections:
 - 25 000 voters: C > LD > L
 - 20 000 voters: L > LD > C
 - 11 000 voters: LD > L > C
 - 4 000 voters: LD > C > L
- 1st round: C: 25 000, L: 20 000, LD: 15 000
 - noone has more than 30 000 votes, so LD is eliminated
- 2nd round: C: 25 000+4 000, L: 20 000 + 11 000
 - L gets the majority of votes, so it wins

Multi-Round Elections

1. All voters vote for their favorite candidate
2. If some candidate gets more than 50% of the votes, he is declared the winner
3. Otherwise, the candidate with the smallest number of votes is eliminated
4. The voters are asked to vote again over the remaining candidates
5. The process repeats until some candidate gets a majority of votes

Single Transferable Vote



- Multi-round elections often produce a more **appealing** outcome than Plurality
- However, they are **hard to implement**:
 - voters have to come to voting booths **many times**
- Single Transferable Vote: an **implementation** of multi-round elections in a **single round** of voting
 - each voter submits a total ranking of candidates
 - the election authority **simulates** multi-round elections based on the information in the ballots (assuming that all voters always vote for their most preferred **available** candidate)

How Good are Plurality With Runoff and STV?

- United Kingdom elections:
 - 25 000 voters: $C > LD > L$
 - 20 000 voters: $L > LD > C$
 - 11 000 voters: $LD > L > C$
 - 4 000 voters: $LD > C > L$
- Plurality chooses C , STV chooses L
- Yet, 40 000 voters prefer LD to L and 35 000 voters prefer LD to C
- Under both Plurality and STV, more than 50% of voters would have preferred a different candidate

Condorcet Winners


- Suppose that each of the n voters has a ranking of all m candidates
- Definition: a candidate c wins a **pairwise election** against a candidate d if more than half of the voters rank c above d
- A candidate is said to be a **Condorcet winner** if he wins in all pairwise elections he participates in



a is the Condorcet winner



Condorcet Consistency

- A voting rule is said to be **Condorcet-consistent** if it selects the Condorcet winner whenever it exists
- United Kingdom elections:
 - 25 000 voters: $C > LD > L$
 - 20 000 voters: $L > LD > C$
 - 11 000 voters: $LD > L > C$
 - 4 000 voters: $LD > C > L$
- Plurality chooses C , STV chooses L
- LD is the Condorcet winner
 - even though it has the smallest number of voters who rank it first
- Hence, neither Plurality nor STV are Condorcet-consistent



Do Elections Always Have Condorcet Winners?

- 2 voters rank **a** above **b**
- 2 voters rank **b** above **c**
- 2 voters rank **c** above **a**
- No Condorcet winner!



- Definition: **G** is a **pairwise majority graph** for an election **E** with a candidate set **C** if its vertex set is **C** and there is an edge from **a** to **b** iff majority of voters prefer **a** to **b**
 - **weighted** pairwise majority graph: the weight of **(a, b)** is the **number of voters** who prefer **a** to **b**
- Theorem: any directed graph with no **2**-cycles can arise as a pairwise majority graph
 - useful for hardness reductions from graph problems

Implementing Directed Graphs

- Theorem: any directed graph with no 2-cycles can arise as a pairwise majority graph
- Proof:
 - fix a directed graph $G = (V, E)$
 - for each edge (a, b) we create a pair of votes v_{ab}, v'_{ab} :
 - v_{ab} ranks the candidates as $a > b > \dots$
 - v'_{ab} is obtained by reversing v_{ab} and then swapping a and b
 - we claim that the majority graph of this election is G :
 - if (x, y) is not in E , one voter in each pair (v_{ab}, v'_{ab}) ranks x above y and the other ranks y above x
 - if (x, y) is in E , both v_{xy} and v'_{xy} rank x above y ; in every other pair one voter ranks x above y and one voter ranks y above x

Condorcet-Consistent Rules: Copeland

- A Condorcet-consistent rule must elect a Condorcet winner when one exists
 - how can we extend this principle if there is no Condorcet winner?
- Copeland rule: each candidate gets
 - 1 point for each pairwise election he wins
 - 0.5 points for each pairwise election he ties
 - the candidate with the largest number of points wins
- In an m -candidate election, if a Condorcet winner exists, he gets $m-1$ point, all other candidates get at most $m-2$ points

Condorcet-Consistent Rules: Maximin

- Maximin rule: the score of each candidate is the number of votes he gets in his worst pairwise election
 - the candidate with the highest score wins
- In an n -voter election,
 - if a Condorcet winner exists
 - his Maximin score is greater than $n/2$,
 - everyone else's Maximin score is less than $n/2$

Consistency

- Consider two elections $E = (C, V)$ and $E' = (C, V')$, where the voter sets V and V' are disjoint
- Suppose a voting rule R outputs a on E and on E'
- **Consistency** axiom: in this case, R should output a on $(C, V \cup V')$
- Claim: STV violates consistency
 - E : 4 x abc , 2 x bac , 4 x cba (a wins under STV)
 - E' : 4 x abc , 6 x bac , 3 x cab (a wins under STV)
 - in the merged election, b wins under STV
- Claim: Copeland and Maximin violate consistency

Scoring Rules

- A **scoring rule** for an election with **m** candidates is given by a vector $(s_1, \dots, s_m)$, $s_1 \geq \dots \geq s_m$
 - each candidate gets s_i points from each voter who ranks him i -th
 - candidate with the **maximum number of points** wins
- Plurality is a scoring rule with score vector $(1, 0, \dots, 0)$
- Borda: $(m-1, m-2, \dots, 2, 1, 0)$
- k-approval: $(\underbrace{1, \dots, 1}_k, 0, \dots, 0)$
 - equivalent to allowing voters to vote for **k** candidates

Scoring Rules: Example

- United Kingdom elections:
 - 25 000 voters: $C > LD > L$
 - 20 000 voters: $L > LD > C$
 - 11 000 voters: $LD > L > C$
 - 4 000 voters: $LD > C > L$
- Plurality: C wins with 25 000 points
- Borda:
 - C gets $2 \times 25\,000 + 1 \times 4\,000 = 54\,000$ points
 - L gets $2 \times 20\,000 + 1 \times 11\,000 = 51\,000$ points
 - LD gets $1 \times 45\,000 + 2 \times 15\,000 = 75\,000$ points
- 2-approval:
 - C : 29 000, L : 31 000, LD : 60 000

Scoring Rules: Pro and Contra

- Scoring rules are easy to understand and implement
- They take into account preferences other than just the voter's top choice
- Every scoring rule is consistent
 - suppose $s(a, E) > s(b, E)$ and $s(a, E') > s(b, E')$
 - then
$$s(a, E \cup E') = s(a, E) + s(a, E') > s(b, E) + s(b, E') = s(b, E \cup E')$$
- However, no scoring rule is Condorcet-consistent
 - Borda: abcde, abedc, bcdea
 - a is the Condorcet winner,
yet a gets 8 points, while b gets 10

More Rules: Bucklin's rule

- How do we choose k for k -approval?
- One possible answer: adaptively
- Let k^* be the smallest value of k such that there is a candidate ranked in top k positions by more than $n/2$ voters
- **Bucklin rule**: output all k^* -approval winners
- Alternative interpretation:
 - for $k = 1, \dots, m$ do
 - ask each voter to name their top k candidates
 - stop when some candidate is named by a **majority**
 - report all such candidates

More Rules: Schulze's Rule (1997)

- Consider the weighted majority graph
 - the weight of the edge ab is the number of voters who prefer a to b
 - only keep edges whose weight is $\geq n/2$
- Strength of a path from a to b :
min weight along that path
- $p[a, b]$: strength of the strongest path from a to b
- a is a winner if $p[a, b] \geq p[b, a]$ for all b
 - always exists (not obvious, but true)
 - polynomial-time computable

More Axioms: Monotonicity


- **Monotonicity**: if **c** wins, and some voter moves **c** higher in her ranking, without changing the order of other candidates, then **c** still wins

+: scoring rules, Copeland, Maximin, Bucklin, Schulze

-: Plurality with Runoff, STV

Example (STV):

a moves to the top
in the first 2 votes



Social Welfare Functions: Rules That Map Ranking to Rankings

Social Welfare Functions

- $L(C)$: the set of all rankings over a candidate set C
 - $|L(C)| = |C|!$
- Definition: a social welfare function maps a collection of rankings over candidates (a profile) to a single joint ranking:
 - $F: L(C)^n \rightarrow L(C)$
- Score-based voting rules can be used to construct social welfare functions: order candidates by score
 - not just scoring rules, but also Copeland, Maximin, etc.
- Kemeny rule:
 - for two votes u, v , define the swap distance between u and v as $d(u,v) = \# \{(a,b): a >_u b, b >_v a\}$
 - Kemeny score of a ranking: total swap distance to votes
 - output all rankings with minimum Kemeny score

Ranking of the Universities: Borda

- A **panel of experts** is supposed to rank UK universities
 - Expert 1: Cambridge > Oxford > UCL > Imperial
 - Expert 2: Oxford > Cambridge > Imperial > UCL
 - Expert 3: UCL > Cambridge > Oxford > Imperial
 - Expert 4: Oxford > Imperial > Cambridge > UCL
 - Expert 5: Imperial > Cambridge > UCL > Oxford
- Goal: produce a total ranking of 4 universities
- Borda rule:
 - each university gets **4 - i** points from each expert who ranks it in position **i**
 - Cambridge: **10**, Oxford: **9**, UCL: **5**, Imperial: **6**

Ranking of the Universities: Kemeny

- A **panel of experts** is supposed to rank UK universities
 - Expert 1: Cambridge > Oxford > UCL > Imperial
 - Expert 2: Oxford > Cambridge > Imperial > UCL
 - Expert 3: UCL > Cambridge > Oxford > Imperial
 - Expert 4: Oxford > Imperial > Cambridge > UCL
 - Expert 5: Imperial > Cambridge > UCL > Oxford
- Goal: produce a total ranking of 4 universities
- Kemeny rule:
 - need to score each of the 24 possible rankings
 - e.g., Oxford > Cambridge > UCL > Imperial
scores 1 + 1 + 3 + 2 + 5

Kemeny Rule: Complete Example

- 5 votes $a > b > c$
- 3 votes $b > a > c$
- 2 votes $c > a > b$
- $s_K(a > b > c) = 5 \times 0 + 3 \times 1 + 2 \times 2 = 7$
- $s_K(a > c > b) = 5 \times 1 + 3 \times 2 + 2 \times 1 = 13$
- $s_K(b > a > c) = 5 \times 1 + 3 \times 0 + 2 \times 3 = 11$
- $s_K(b > c > a) = 5 \times 2 + 3 \times 1 + 2 \times 2 = 17$
- $s_K(c > a > b) = 5 \times 2 + 3 \times 3 + 2 \times 0 = 19$
- $s_K(c > b > a) = 5 \times 3 + 3 \times 2 + 2 \times 1 = 23$
- Output: $a > b > c$

Transitive Majority Relation

- Definition: for odd n , the **majority relation** of a profile of rankings $P = (R_1, \dots, R_n)$ is the set R_{maj} of pairs of candidates (a, b) such that the **majority** of voters in P prefer a to b
 - i.e., the set of **edges** of the majority graph
- Definition: the majority relation R_{maj} of a profile P is **transitive** if $(a, b) \in R_{\text{maj}}$ and $(b, c) \in R_{\text{maj}}$ implies $(a, c) \in R_{\text{maj}}$
 - equivalently, **no directed cycles** in the majority graph
 - because if a graph has a directed cycle, it has one of length **3**
- Claim: if a profile has transitive majority, it has a **Condorcet winner**
 - start with a candidate a ; if a is a Condorcet winner, we are done
 - otherwise, $(b, a) \in R_{\text{maj}}$ for some b
 - by transitivity b has **higher outdegree** in the majority graph
 - if b is a Condorcet winner, we are done, else continue
 - we stop after at most $m - 1$ steps, as the **max** outdegree is $m - 1$

Transitive Majority: Ranking

- Claim (previous slide): if a profile has transitive majority, it has a Condorcet winner
- In the language of graphs: if an m -vertex graph has no directed cycles, it has a vertex of outdegree $m - 1$
- It follows by induction that its vertices have degrees $m - 1, m - 2, \dots, 0$
- Rank the candidates by their outdegree (Copeland score)
- Claim: a is ranked above b iff (a, b) is in R_{maj}
 - if a is ranked in position i , its outdegree is $m - i$
 - thus, it is majority-preferred to any candidate below her in the ranking

Kemeny for Transitive Majority

- By the previous slide, If a profile has a **transitive majority** relation, then R_{maj} can be viewed as a **ranking**
- Claim: in this case, the **Kemeny** rule outputs this ranking
- Proof:
 - let n_{ab} be the number of voters who prefer **a** to **b**
 - consider a ranking **R**
 - each pair **{a, b}** contributes at least $\min\{n_{ab}, n_{ba}\}$ to $s_K(R)$
 - if $R = R_{maj}$, each pair **{a, b}** contributes exactly $\min\{n_{ab}, n_{ba}\}$
- Intuition: when majority is not transitive, **Kemeny** outputs a ranking that is “closest” to the (weighted) majority relation
- Fact: computing Kemeny rankings is **NP-hard**, even for **n=4**

Axioms for Social Welfare Functions

- **Anonymity**: permuting the rankings in the input should not change the output ranking
- **Neutrality**: renaming all candidates produces the same ranking wrt the new names
 - e.g., suppose the rule outputs $a > b > c$ on $(a > b > c, b > a > c, c > a > b)$
 - rename the candidates as $a \rightarrow b, b \rightarrow c, c \rightarrow a$
 - then the input becomes $(b > c > a, c > b > a, a > b > c)$ and the rule should output $b > c > a$
- **Monotonicity**: if some voter moves c up in their ranking, in the overall ranking c does not go down

Axioms for Social Welfare Functions

- **Pareto efficiency**: if all voters rank **a** above **b**, in the final ranking **a** should appear above **b**
- **Unanimity**: if all voters submit the same ranking, the rule should output this ranking
 - follows from Pareto efficiency
- **Independence of irrelevant alternatives (IIA)**: if **a** is ranked above **b** in the current profile, and we permute the candidates in each vote without changing the relative order of **a** and **b**, then **a** should be ranked above **b** in the resulting profile

Dictatorship

- There is a very simple rule that produces a ranking of alternatives and satisfies **most** of our criteria: **dictatorship**
- This rule simply **copies** the ranking of some fixed voter
- Satisfies monotonicity, neutrality, Pareto efficiency, IIA (not anonymity though)
- No voter benefits from misreporting
- Is usually **not** an **acceptable** voting rule for obvious reasons

Arrow's Theorem [1951]

- Suppose there are at least 3 candidates
- Then any voting rule that produces a ranking of all candidates and is simultaneously:
 - Pareto efficient and
 - independent of irrelevant alternatives

is a dictatorship

“There is no perfect voting rule”

Multiwinner Rules With Ranked Ballots

k-Borda

$k = 2$


- Output k candidates with highest Borda scores
- May output **several committees** due to tie-breaking

– $k = 3$

– Borda scores are
 $s(a) = 15$, $s(b) = 14$,
 $s(c) = 11$, $s(d) = 11$,
 $s(e) = 9$

– output:

$\{a, b, c\}, \{a, b, d\}$




	4	3	2	1	0
V_1 :					
V_2 :					
V_3 :					
V_4 :					
V_5 :					
V_6 :					

Single Transferable Vote (STV)

- Quota $q = n/k$
- Candidates ranked **first**
 q times get elected
- For each elected candidate,
 q supporters are **removed**
— which ones?
- If there are no such candidates,
delete one with the **smallest**
number of 1st place votes
- **Transfer** the votes from elected
and removed candidates
- Repeat until k candidates
are **elected**

$$C = \{ \text{Candidate 1}, \text{Candidate 2}, \text{Candidate 3}, \text{Candidate 4}, \text{Candidate 5} \}$$
$$V = (v_1, \dots, v_6)$$



STV: Example

$$k = 2 \rightarrow \text{quota} = 6/2 = 3$$

- Quota $q = n/k$
- Candidates ranked **first** q times get elected
- For each elected candidate, q supporters are **removed**
 - which ones?
- If there are no such candidates, delete one with the **smallest number** of 1st place votes
- **Transfer** the votes from elected and removed candidates
- Repeat until k candidates are **elected**

$$C = \{ \text{Male Candidate}, \text{Female Candidate} \}$$
$$V = (v_1, \dots, v_6)$$

V_1 :



>



V_2 :



>

V_3 :



>

V_4 :



V_5 :



>



V_6 :




>



Chamberlin-Courant Rule

- Idea: each voter evaluates a committee according to its **best** member
- The **score** of voter i for candidate c :
 $sc(i, c) = |C| - t$ if i ranks c in position t (**Borda** scores)
- The **score** of voter i for committee S :
 $sc(i, S) = \max \{sc(i, c) : c \text{ in } S\}$
- Chamberlin-Courant rule:
given $v_1, \dots, v_n$, output a committee in
 $\operatorname{argmax}_{S \subseteq C, |S|=k} \sum_{i=1, \dots, n} sc(i, S)$ (**utilitarian**)
 $\operatorname{argmax}_{S \subseteq C, |S|=k} \min_{i=1, \dots, n} sc(i, S)$ (**egalitarian**)

Chamberlin-Courant Rule: Example



- Egalitarian Chamberlin-Courant score: $\min\{3, 4, 3, 2, 4\} = 2$
- Utilitarian Chamberlin-Courant score: $3 + 4 + 3 + 2 + 4 = 16$

Axioms for Multiwinner Voting: Narrow Top Consistency

- Narrow top consistency:
 - let T be the set of candidates ranked first by at least one voter
 - if $|T| = k$, then T should be a winning committee
- Satisfied by Chamberlin-Courant (both util and egal)
- Failed by k -Borda and STV
 - 5 votes: $a \succ b \succ x$
 - 1 vote: $x \succ b \succ a$
 - $k = 2$

Axioms for Multiwinner Voting: Committee Monotonicity

- Committee monotonicity:
 - suppose candidate **a** is on a winning committee at a profile **P** for committee size **k**
 - then **a** is on a winning committee at **P** for committee size **k+1**
- Satisfied by **k-Borda**
- Failed by **STV** and **Chamberlin-Courant**
- Example for **CC**:
 - **n/2** votes of the form **a > b > ... > c**
 - **n/2** votes of the form **c > b > ... > a**
 - **k=1: {b}; k=2: {a, c}**

Axioms for Multiwinner Voting: Solid Coalitions

- A group of voters N' forms a j -solid coalition in favour of a set of candidates S , $|S|=j$, if:
 - $|N'| = jn/k$
 - all voters in N' rank S above $C \setminus S$
- **Solid Coalitions**: if there is a j -solid coalition on favour of S , all members of S get elected
- Satisfied by **STV**
 - by pigeonhole principle, some candidate in S gets n/k first-place votes from voters in N'
 - no vote in N' gets transferred outside of S
 - induction
- Failed by **k -Borda** and **Chamberlin-Courant**

Chamberlin-Courant: NP-hardness

- Theorem: it is NP-complete to decide if a given profile P admits a size- k committee with egal-CC score $\geq s$
- Proof:
 - reduction from Vertex Cover
 - graph (V, E) and the target cover size $t \rightarrow (P, k, s)$
 - set of candidates $C = V$
 - one voter per edge: edge $\{u, v\} \rightarrow$ a voter $u \succ v \succ V \setminus \{u, v\}$
 - $k = t, s = |C| - 2$
 - cover $S, |S| = t \iff$ size- k committee with egal-CC score $\geq s$

Chamberlin-Courant: NP-hardness

- Theorem: it is NP-complete to decide if a given profile P admits a size- k committee with util-CC score $\geq s$
- Can we use the same construction?
 - cover $\rightarrow$ committee S with $sc(i, S) \geq |V| - 2$ for each voter i
 - thus, util-CC score of S is $\geq |E|(|V| - 2)$
 - if there is no size- k vertex cover, then for each size- k committee S we have $sc(i, S) \leq |V| - 3$ for some voter i
 - but what if $sc(j, S) = |V| - 1$ for each $j \neq i$?
 - in that case, the util-CC score of S may be as high as $(|E| - 1)(|V| - 1) + |V| - 3 = |E|(|V| - 1) - 2$
 - this is greater than $|E|(|V| - 2)$ for $|E| > 2$

Chamberlin-Courant: NP-hardness

- Can we tweak our reduction?
- Idea: use **dummy** candidates
 - set of candidates $C = V \cup D$, $D = \{d_1, \dots, d_q\}$
 - edge $\{u, v\} \rightarrow$ voter $u \succ v \succ d_1 \succ \dots \succ d_q \succ V \setminus \{u, v\}$
 - let S be a size- k subset of V
 - max possible util-CC score is $M = |E|(q + |V| - 1)$
 - if S is a vertex cover, its util score is at least $X = M - |E|$
 - if S is not a vertex cover, its util-CC score is at most $Y = M - (q + |V| - 1) + (|V| - 3) = M - q - 2$
 - $X > Y$ if $q > |E| - 2$
 - but what if S contains d_1 ?

Chamberlin-Courant: NP-hardness

- Can we tweak our reduction?
- Idea: use $k + 1$ sets of **dummy** candidates
 - $C = V \cup D^{(1)} \cup D^{(2)} \cup \dots \cup D^{(k+1)}$, $|D^{(j)}| = q$ for each j
 - edge $\{u, v\} \rightarrow k + 1$ voters:
 - voter $u \succ v \succ D^{(1)} \succ D^{(2)} \dots \succ D^{(k+1)} \succ V \setminus \{u, v\}$
 - voter $u \succ v \succ D^{(2)} \succ D^{(1)} \dots \succ D^{(k+1)} \succ V \setminus \{u, v\}$
 - ...
 - voter $u \succ v \succ D^{(k+1)} \succ D^{(k)} \dots \succ D^{(1)} \succ V \setminus \{u, v\}$
 - key observation: $|S| = k$, so $S \cap D^{(j)} = \emptyset$ for some $j = 1, \dots, k + 1$
 - max possible util-CC score is $M = |E|(k + 1)(q(k + 1) + |V| - 1)$
 - if S is a vertex cover, its util-CC score is at least $X = M - |E|(k + 1)$
 - if S is not a vertex cover, its util-CC score is at most $Y = M - (q(k + 1) + |V| - 1) + (qk + |V| - 3) = M - q - 2$
 - $X > Y$ for $q > |E|(k + 1) - 2$ ✓

Multiwinner Rules With Approval Ballots

Challenge: choosing invited speakers

- A workshop PC has 9 members
 - 6 algorithms researchers
 - 3 game theory researchers
- 3 slots for invited speakers, 5 candidates

PC members


speakers

●	●	●	●		●			
●		●	●	●	●			
●	●		●	●	●			
			●			●		●
	●					●	●	●

should there be a game theory speaker?

Voting with approval ballots


- A set of alternatives C
- n voters $\{1, \dots, n\}$
- Each voter **approves** a subset of candidates $A_i \subseteq C$
- Goal: select k winners



1: c_1, c_2
2: c_2
3: c_2
4: c_1
5: c_3

Approval voting (AV)

- Each candidate gets **one** point from each voter who approves her
- **k** candidates with the **highest score** are selected
 - ties broken **deterministically**



for $k=3$
AV outputs
 $\{C_1, C_2, C_3\}$

Speaker selection under AV

- A workshop PC has 9 members
 - 6 algorithms researchers
 - 3 game theory researchers
- 3 slots for invited speakers, 5 candidates

PC members


speakers

●	●	●	●		●			
●		●	●	●	●			
●	●		●	●	●			
			●			●		●
	●					●	●	●

no game theory speaker is selected

Representation

- 5 voters get 3 representatives,
4 voters get 0 representatives
- Intuition: each cohesive
group of voters
of size n/k “deserves”
at least one representative




for $k=3$
AV outputs
 $\{C_1, C_2, C_3\}$

Justified Representation

[ABCEFW, AAI'15/SCW'17]


- Definition: a committee W provides **justified representation (JR)** for a list of ballots $(A_1, \dots, A_n)$ and committee size k if there is no set of voters X with $|X| \geq n/k$ and $\bigcap_{i \in X} A_i \neq \emptyset$ such that W contains no candidates from $\bigcup_{i \in X} A_i$.
- Approval Voting fails JR



Applications

- Representative governing bodies

- Search/advertising




- Choice of training data for ML models

- Blockchain proof-of-stake (Polkadot)


Greedy Approval Voting (GrAV)

- Greedy algorithm (GrAV)
 - proceed in k steps
 - at each step, select a candidate that **covers** as many uncovered voters as possible
- Claim: GrAV **always** outputs a committee that provides JR.



Proportional Approval Voting (PAV)

- Thiele, early 20th century
- Voter i derives “utility” of 1 from having 1 approved candidate in W ,
 $1 + 1/2$ from having 2,
 $1 + 1/2 + 1/3$ from having 3, etc.
- Goal: select a size- k committee W that maximizes the sum of voters’ “utilities”



for $k=2$


AV outputs $\{c_1, c_2\}$,

PAV outputs

$\{c_1, c_3\}$ or $\{c_2, c_3\}$

PAV satisfies JR

- Theorem: PAV satisfies JR
- Proof idea:
 - if not, there is some $c \in C \setminus W$ that could **increase** the total utility by $\geq n/k$
 - we can show that for some candidate $a \in W$ his **marginal contribution** is $< n/k$
 - hence, swapping a with c would increase the **PAV score**



PAV satisfies JR: Details

- Proof:

- $MC(a, i) := u_i(W) - u_i(W \setminus a)$: marginal contribution of a for i

- $MC(a) := \sum_i MC(a, i)$: marginal contribution of a

- $\sum_a MC(a) =$


- $MC(a) < n/k$ for some a in W

$$MC(a, 1) = 1/4$$

$$MC(a, 2) = 1/3$$

$$MC(a, 3) = 1/5$$

$$MC(a) = 1/4 + 1/3 + 1/5$$




GrAV vs PAV?

- Both GrAV and PAV satisfy JR
- Have we solved the issue of **fair representation**?
- Axiomatically?
- Algorithmically?

Is JR enough?

- Is it OK to choose $\{c_1, c_4\}$???
- Perhaps a **very large** coherent group of voters “deserves” **several** representatives?
- Idea: if n/k voters who agree on a candidate “deserve” **one** representative, then maybe $\ell \cdot n/k$ voters who agree on ℓ candidates “deserve” ℓ representatives?



Extended Justified Representation

- Definition: a committee W provides **extended justified representation (EJR)** if for every $\ell > 0$ and every ℓ -large, ℓ -cohesive set of voters X it holds that **at least one** voter in X receives ℓ representatives
- $\ell = 1$: justified representation
- Note for game theorists: this is similar to, but weaker than **core stability**

GrAV vs PAV: representation

- Observation: GrAV fails EJR
 - once a voter is **covered**, GrAV ignores it
- Theorem: PAV satisfies EJR
 - **local swap** argument

GrAV vs PAV: complexity


- Observation: GrAV is **poly-time** computable
- Theorem: PAV is **NP-hard** to evaluate
[AGGMMW'14, FSL'14]

Challenge [ABCEFW'15]:

Is there an **easy to compute**
voting rule that satisfies **EJR**?

First attempt: Greedy

- Thiele, early 20th century
- **Sequential** version of PAV
- Initialize:
 $W = \emptyset$
- Repeat k times:
 - add to W a candidate c that increases the **PAV score** as much as possible
- Was briefly used in Sweden
- Polynomial-time, but fails **JR**



for $k=2$

PAV outputs $\{c_2, c_3\}$,
seqPAV outputs
 $\{c_1, c_2\}$ or $\{c_1, c_3\}$

Generalizing PAV and seqPAV: arbitrary weights

- PAV and seqPAV both use weight vector $(1, 1/2, 1/3, \dots)$
- We can use an **arbitrary** weight vector $(w_1, w_2, \dots)$ with $w_1 = 1, w_1 \geq w_2 \geq \dots$ instead:
 $(w_1, w_2, \dots)$ -AV and $(w_1, w_2, \dots)$ -seqAV
- $(1, 0, \dots)$ -AV:
 - choose k candidates to **cover** as many **voters** as possible
- $(1, 0, \dots)$ -seqAV:
 - choose k candidates to **cover** as many **voters** as possible **at each step** (GrAV)

Satisfying EJR?

- Theorem: $(1, 0, \dots)$ -seqAV satisfies JR
- Theorem: $(w_1, w_2, \dots)$ -seqAV fails JR if $w_2 > 0$
- Theorem: $(1, 0, \dots)$ -seqAV fails EJR
- Theorem: $(w_1, w_2, \dots)$ -AV fails EJR if $(w_1, w_2, \dots) \neq (1, 1/2, 1/3, \dots)$

Second attempt: Local search

- PAV satisfies EJR, but is NP-hard to compute
 - recall: if a committee fails EJR, we can find a committee with a better PAV score by a local swap
- LS-PAV [AEHLSS'18]
 - start with arbitrary W
 - if there is a local swap that improves the PAV score, perform it
 - LS-PAV satisfies EJR
 - each iteration is easy to compute
 - but the # of iterations may be super-poly [EK'24]

One more step... [AHLSS'18]

- PAV satisfies EJR, but is NP-hard to compute
 - recall: if a committee fails EJR, we can find a committee with a ~~better~~ PAV score by a local swap
 - much better: by at least n/k^2
- ϵ -LS-PAV
 - start with some W
 - if there is a local swap that improves the PAV score by at least ϵ , perform it

for $\epsilon = n/k^2$,
 ϵ -LS-PAV satisfies EJR
and is in P

Are we done?

- Is ϵ -LS-PAV a usable rule?
- Can you explain it to voters?
 - computing PAV scores?
 - using local search?
 - ignoring small improvements?
- Can it be extended beyond the multi-winner voting setting?
- In a way that preserves its properties?

Beyond multiwinner voting: Participatory budgeting

- A set of P projects
- A budget b
- Each project p in P has a cost c_p
- Voters report (approval) preferences over projects
- Goal: select a budget-feasible set of projects to fund based on voters' preferences
- Porto Alegre'89, Paris, New York, Madrid...
- Note: MW voting = PB with unit costs



EJR for participatory budgeting

- Definition: Consider a subset of projects T
A group of voters $S \subseteq N$ is T -cohesive if
 - all voters in S approve all projects in T and
 - $c(T) \leq |S| \times b / |N|$
- Definition: an outcome W provides extended justified representation (EJR) if for every $T \subseteq P$ and every T -cohesive group of voters S some voter in S approves $\geq |T|$ projects in W


Method of Equal Shares [PS'20,PPS'21]

- The budget b is **split** among the $n(=|N|)$ voters
 - each voter gets b/n units of money
- Projects are added one by one
- (Some of) voters who approve project p may **pool their money** to buy p (at cost c_p)
 - cost c_p needs to be **split equally** among those who pay for p
 - except if someone's about to **run out** of money
- At each step, we select a project that **minimises the max payment** per voter

MES: Example (general costs)

- 4 voters, budget per voter: \$10
- Projects: $c(A) = \$8$, $c(B) = \$21$, $c(C) = \$16$
- Approvals

- voter 1: A, B
- voter 2: A, B
- voter 3: B, C
- voter 4: C




- Step 1:
 - A can be shared as (V1: \$4, V2: \$4): $\pi(A) = \$4$
 - B can be shared as (V1: \$7, V2: \$7, v3: \$7): $\pi(B) = \$7$
 - C can be shared as (V3; \$8, V4: \$8): $\pi(C) = \$8$
 - A gets selected; budgets of V1 and V2 are reduced to \$6

MES: Example (general costs)

- 4 voters, budget per voter: \$10
- Projects: $c(A) = \$8$, $c(B) = \$21$, $c(C) = \$16$
- Approvals


- voter 1: A, B
- voter 2: A, B
- voter 3: B, C
- voter 4: C



- Step 1: A is selected, V1 and V2 pay \$4 each
- Step 2:
 - B can be shared as (V1: \$6, V2: \$6, V3: \$9): $\pi(B) = \$9$
 - C can be shared as (V3: \$8, V4: \$8): $\pi(C) = \$8$
 - C gets selected; budgets of V3 and V4 are reduced to \$2

Theory meets practice

- PAV fails EJR in the context of PB
- MES satisfies EJR in the context of PB
[Peters, Pierczynski, Skowron'20]
- Implemented in practice
 - Wieliczka, April 2023, Aarau, June 2023, Swiecie, Oct 2023
 - further locations under discussion




MES: Room for improvement

- **MES** may fail to **exhaust** the budget
 - i.e., it may select a set of projects W such that there is a project p in $P \setminus W$ with $c_p \leq b - c(W)$
- Current approach: run **MES** with budgets $b, b+n, b+2n, \dots$ to get $W_b, W_{b+n}, W_{b+2n}, \dots$
 - until cost at step t **exceeds** b
 - output solution at step $t-1$
 - restart **from scratch** every time
- Expensive
 - at least for experiments and fine-tuning)

Local search again [Kraicz, E., WINE'23]

- Observation: output of **MES** is actually a pair (W, X) , where $X = (x_{i,p})$ encodes how much voter i pays for project p
- Idea: define a **local search** procedure on the space of pairs (W, X)
 - output is **MES-like**
 - convergence in **poly-many** steps
 - when **ls-MES** converges, the outcome satisfies **EJR**
- Advantage: **iterating** through budgets quickly

Empirically?



Measure: number of changes in outcome
(project added/removed)

What's next?

- Temporal dimension [KTE'22, TOE'24a,b,c]
- Negative votes?
- Unequal entitlements?
- Richer ballot formats?
- Project dependencies?
- Beyond PB: fair participatory urban planning
 - public transport routes [BEL'24]
 - cycle paths
 - road repair