Computational Social Choice 2020

Ulle Endriss
Institute for Logic, Language and Computation
University of Amsterdam

[http://www.illc.uva.nl/~ulle/teaching/comsoc/2020/]

Plan for Today

So far we have only studied voting rules designed to elect *one winner* (ties were considered a nuisance, not a desideratum).

Today we are going to discuss two kinds of voting rules designed to elect at set of (exactly) k *winners* (so tie-breaking is still an issue):

- multiwinner voting rules with *approval ballots*
- multiwinner voting rules with *ranked ballots*

We shall restrict attention to rules under which each voter reports her preferences on *individual* alternatives (not k-sets of alternatives).

Much of this and more is reviewed by Faliszewski et al. (2017).

P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Multiwinner Voting: A New Challenge for Social Choice Theory. In *Trends in COMSOC*. AI Access, 2017.

Application Scenarios

All of these scenarios can be modelled as multiwinner elections:

- A hiring committee has to shortlist k out of m job candidates to invite to interviews (after which one of them will get an offer).
- An online retailer needs to pick k out of m products to display on the company's front page, given (likely) customer preferences.
- In a national election, k out of m candidates running need to be chosen to form the new parliament, based on voter preferences.

Exercise: *What are good rules? What properties should they satisfy?*

Exercise: *Explain the difference between a multiwinner voting rule and an irresolute (single-winner) voting rule.*

Multiwinner Voting with Approval Ballots

Fix a finite set $A = \{a, b, c, \ldots\}$ of *alternatives* with $|A| = m \geqslant 2$ and a positive integer $k \leqslant m$. Let $A[k] = \{S \subseteq A \mid \#S = k\}$.

Each member of a set $N = \{1, \ldots, n\}$ of *voters* supplies us with an *approval ballot* $A_i \subseteq A$, yielding a *profile* $\boldsymbol{A} = (A_1, \ldots, A_n)$.

A multiwinner *voting rule* for approval ballots for N, A, and k maps any given profile to one or more winning *committees* of size k each:

$$F : (2^A)^n \to 2^{A[k]} \setminus \{\emptyset\}$$

Such a rule is called *resolute* in case $|F(\boldsymbol{A})| = 1$ for all profiles $\boldsymbol{A}$.

Example: The basic rule of *approval voting* (AV) elects committees S with maximal approval score $\sum_{x \in S} \sum_{i \in N} \mathbb{1}_{x \in A_i} = \sum_{i \in N} |S \cap A_i|$.

Remark: For *participatory budgeting* with a budget of k and projects of cost 1, both *greedy approval* and *max-approval* reduce to AV.

Proportional Justified Representation

We would like to be able to guarantee some form of *proportionality:* sufficiently large and cohesive groups need sufficient representation.

One way of attempting to formalise this intuition:

A rule F satisfies *proportional justified representation* (PJR) if, for every profile $\boldsymbol{A}$, coalition $C \subseteq N$, and $\ell \in \mathbb{N}$ with $|\bigcap_{i \in C} A_i| \geqslant \ell$ and $\frac{|C|}{\ell} \geqslant \frac{n}{k}$, it is the case that $|S \cap \bigcup_{i \in C} A_i| \geqslant \ell$ for all $S \in F(\boldsymbol{A})$.

If this holds at least for $\ell = 1$, we speak of *justified representation*.
If $\max_{i \in C} |S \cap A_i| \geqslant \ell$, we speak of *extended justified representation*.

<u>Exercise:</u> *What do you think about these definitions? Reasonable?*

H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. Justified Representation in Approval-based Committee Voting. *Soc. Choice & Welf.*, 2017.

L. Sánchez-Fernández, E. Elkind, M. Lackner, N. Fernández, J.A. Fisteus, P. Basanta Val, and P. Skowron. Proportional Justified Representation. AAAI-2017.

Counterexamples

The rule of *basic AV* does not satisfy even the weakest JR axiom:

> Suppose $k = 3$. If 51% approve $\{a, b, c\}$ and 49% approve $\{d\}$, then AV elects $\{a, b, c\}$, even though the 49% "deserve" d.

You may feel that a more appropriate definition of JR would require $\bigcap_{i \in C} A_i \neq \emptyset$ and $|C| \geqslant \frac{n}{k}$ to imply $S \cap \bigcap_{i \in C} A_i \neq \emptyset$ for all $S \in F(\boldsymbol{A})$.

But this axiom of *strong justified representation* is violated by *all* rules:

> Suppose $k = 3$. Suppose 2 voters each approve $\{a\}$ and $\{d\}$, while 1 voter each approves $\{b\}$, $\{c\}$, $\{a, b\}$, $\{b, c\}$, $\{c, d\}$. Then each $x \in \{a, b, c, d\}$ is approved by a coalition of $3 \geqslant \frac{9}{3}$, but we cannot elect all four alternatives.

Proportional Approval Voting

The rule of *proportional approval voting* (PAV) returns committees S that maximse the score $\sum_{i \in N} 1 + \cdots + \frac{1}{|S \cap A_i|}$.

Idea: Diminishing marginal utility of getting an extra representative.

Proposed by Danish mathematician Thorvald N. Thiele in the 1890s.

Generalisation: The *Thiele rule* with weights $\boldsymbol{w} = (w_1, w_2, \ldots)$ returns committees S that maximise the score $\sum_{i \in N} w_1 + \cdots + w_{|S \cap A_i|}$.

So why is PAV a good name for this rule . . . ?

Proportionality of PAV

Theorem 1 (Sánchez-Fernández et al., 2017) *PAV satisfies PJR.*

Remark: Minor variant of a result originally due to Aziz et al. (2017).

Proof sketch: For the sake of contradiction, let us assume there is an $S \in F(\boldsymbol{A})$ with $|S \cap \bigcup_{i \in C} A_i| < \ell \leqslant |\bigcap_{i \in C} A_i|$ despite $\frac{|C|}{\ell} \geqslant \frac{n}{k}$.

So $x^\star \notin S$ for some $x^\star \in \bigcap_{i \in C} A_i$. As $|S \cap A_i| < \ell$ for each $i \in C$, *adding $x^\star$ to S* would yield a *marginal PAV score $\geqslant$* $|C| \cdot \frac{1}{\ell} \geqslant \frac{n}{k}$.

Calculate sum of marginal scores in S: Voter i contributes $\frac{1}{|S \cap A_i|}$ for each $x \in S \cap A_i$. So sum across all $i \in N$, $x \in S$ is $\leqslant \sum_i \frac{|S \cap A_i|}{|S \cap A_i|} = n$. Hence, the *marginal score of the worst $x \in S$* must be $\leqslant \frac{n}{k}$.

So *swapping that x for $x^\star$* would give at least as high a score. A small refinement of this idea yields a full proof (see original papers). ✓

Crucially, the same authors also prove that PAV is *the only* Thiele rule that satisfies PJR. So *PJR characterises PAV* (in this space of rules).

Load-Balancing Approval Voting

Another important family of rules are the *load-balancing rules* going back to the work of Lars Edvard Phragmén (1863–1937).

We saw the generalisation of one of them for participatory budgeting.

These rules also satisfy PJR.

M. Brill, R. Freeman, S. Janson, and M. Lackner. Phragmén's Voting Methods and Justified Representation. AAAI-2017.

Proportionality and Strategyproofness

Suppose voters care about the number of elected alternatives they like. So a reasonable definition of *strategyproofness* for resolute rules F would be to require $|F(\boldsymbol{A}) \cap A_i| \geqslant |F(\boldsymbol{A}_{-i}, A_i') \cap A_i|$ for all $\boldsymbol{A}$, i, A_i'.

Theorem 2 (Peters, 2018) *If $k \geqslant 3$, k divides n, and $m \geqslant k+1$, then no resolute and weakly efficient multiwinnner voting rule for approval ballots satisfies both JR and strategyproofness.*

Holds also for weaker forms of proportionality and strategyproofness. Weak efficiency just requires $F(\boldsymbol{A}) \subseteq \bigcup_{i \in N} A_i$ unless $|\bigcup_{i \in N} A_i| < k$.

Proof obtained with the help of automated reasoning techniques. For a generalisation to irresolute rules, refer to Kluiving et al. (2020).

D. Peters. Proportionality and Strategyproofness in Multiwinner Elections. AAMAS-2018.

B. Kluiving, A. de Vries, P. Vrijbergen, A. Boixel, and U. Endriss. Analysing Irresolute Multiwinner Voting Rules with Approval Ballots via SAT Solving. ECAI-2020.

Multiwinner Voting with Ranked Ballots

Fix a finite set $A = \{a, b, c, \ldots\}$ of *alternatives* with $|A| = m \geqslant 2$ and a positive integer $k \leqslant m$. Let $A[k] = \{S \subseteq A \mid \#S = k\}$.

Each member of a set $N = \{1, \ldots, n\}$ of *voters* supplies us with a *ranked ballot* $R_i \in \mathcal{L}(A)$, yielding a *profile* $\boldsymbol{R} = (R_1, \ldots, R_n)$.

A multiwinner *voting rule* for ranked ballots for N, A, and k maps any given profile to one or more winning *committees* of size k each:

$$F : \mathcal{L}(A)^n \to 2^{A[k]} \setminus \{\emptyset\}$$

Such a rule is called *resolute* in case $|F(\boldsymbol{R})| = 1$ for all profiles $\boldsymbol{R}$.

Approaches Based on Sequential Elimination

Need to elect k committee members from a pool A of m alternatives, based on the ranked preferences reported by n voters.

Earlier we discussed *STV* as a single-winner voting rule, but in fact it is mostly used for multiwinner elections:

- If some candidate $x^\star$ is ranked first at least $q = \lfloor \frac{n}{k+1} \rfloor + 1$ times, then elect $x^\star$ and eliminate both $x^\star$ and q voters ranking $x^\star$ first.

  Otherwise, eliminate a plurality loser from the profile.

  Repeat until all k seats are filled.

Involves tie-breaking at all levels ($\hookrightarrow$ *parallel-universe tie-breaking*).

Could also use rules other than plurality to eliminate weak candidates.

From Simple Voting to Multiwinner Voting

Need to elect k committee members from a pool A of m alternatives, using a standard voting rule F. Three approaches come to mind:

- *Rank-and-cut:* If F really is a social welfare function returning a ranking (e.g., Kemeny or Slater), elect its k top elements.
- *Score-and-cut:* If F comes with a notion of score for an alternative (e.g., Borda or Copeland), elect the k top-scoring alternatives.
- *Choose-and-repeat:* If F is resolute, elect the winner $x^\star$ under F, and repeat with $A := A \setminus \{x^\star\}$ until all seats are filled.

  Alternatively: If F is irresolute, in each round, choose *all* winners.

Of course, there are tie-breaking issues for all of these.

Exercise: *What do you think? Are these approaches any good?*

Example

Suppose we want to use the *plurality rule* to elect $k = 2$ winners:

3 voters:	$a \succ c \succ b$
2 voters:	$b \succ c \succ a$
1 voter:	$c \succ b \succ a$

We might proceed as follows:

- *Score-and-cut:* a gets 3, b gets 2, c gets 1. So $\{a, b\}$ wins.
- *Choose-and-repeat:* a wins first round, then c. So $\{a, c\}$ wins.

Thus: these really are very different rules!

Some Common Rules

Simple extensions to multiwinner voting rules are common in practice:

- Choose-and-repeat + plurality is known as *sequential plurality*. For $k = 3$ used to elect English bishops.
- Score-and-cut + plurality is known as *single nontransferable vote*. For $k = 3$ used to elect rectors of public universities in Brazil.
- Score-and-cut + k-approval PSR (where k is the committee size) is known as *bloc voting*. For $k = 3$ used to elect Irish bishops.

S. Barberà and D. Coelho. How to Choose a Non-controversial List with k Names. *Social Choice and Welfare*, 2008.

E. Elkind, P. Faliszewski, P. Skowron, and A. Slinko. Properties of Multiwinner Voting Rules. *Social Choice and Welfare*, 2017.

The Increasing-Committee-Size Paradox

Suppose we use *bloc voting* (score-and-cut + k-approval):

1 voter:	$a \succ b \succ c \succ d$
1 voter:	$b \succ a \succ d \succ c$
1 voter:	$a \succ c \succ d \succ b$
2 voters:	$b \succ d \succ c \succ a$

For committee size $k = 2$, candidates a and b are the winners.

But if we increase the committee size to $k = 3$, then a will lose!

The above is a simplified variant of a paradox due to Staring (1986), who presents a profile where the two committees even are disjoint.

Exercise: *Show that no choose-and-repeat rule has this problem.*

M. Staring. Two Paradoxes of Committee Elections. *Mathematics Mag.*, 1986.

Condorcet Committees

Not clear or uncontroversial how to extend the Condorcet Principle to multiwinner elections. One proposal is due to Gehrlein (1985):

- Committee $S \subseteq A$ is a weak *Condorcet committee* under profile $\boldsymbol{R}$ if $|N^{\boldsymbol{R}}_{x \succ y}| \geqslant |N^{\boldsymbol{R}}_{y \succ x}|$ for all $x \in S$ and $y \in A \setminus S$.

So a committee that is *not* a weak Condorcet committee is *unstable:* a majority of voters would want to replace one of its members.

Call a multiwinner voting rule *stable* if it elects a weak Condorcet committee whenever such a committee exists.

Remark: Of course, weak Condorcet committees need not exist.

W.V. Gehrlein. The Condorcet Criterion and Committee Selection. *Mathematical Social Sciences*, 1985.

Unstable and Stable Rules

Proposition 3 (Barberà and Coelho, 2008) *No PSR combined with either score-and-cut or choose-and-repeat is stable.*

Proof: Omitted. Similar to the proof (by inspection of a problematic profile) we have seen for all PSR's failing the Condorcet Principle.

Proposition 4 (Barberà and Coelho, 2008) *The Kemeny rule combined with rank-and-cut is stable.*

Proof: Omitted, but easy and unsurprising.

Proposition 5 (Barberà and Coelho, 2008) *Every stable rule is subject to the increasing-committee-size paradox.*

Proof: Omitted (uses example where CCs for $k = 2, 3$ are disjoint).

S. Barberà and D. Coelho. How to Choose a Non-controversial List with k Names. *Social Choice and Welfare*, 2008.

Complex Rules: Minimising Misrepresentation

Suppose that, given profile $\boldsymbol{R}$, we have elected committee S of size k. Use a function $r : N \to S$ to assign each voter to "her" *representative*. The *misrepresentation* of i is how many candidates she prefers to $r(i)$.

The *Chamberlin-Courant rule* chooses a committee S that minimises total misrepresentation if voters pick their favourite representatives:

$$\text{elect } S \subseteq A \text{ minimising } \min_{r:N \to S} \sum_{i \in N} \#\{x \in A \mid (x, r(i)) \in R_i\}$$

The *Monroe rule* does the same, but requires each committee member to represent the same number of voters (± 1): $\lfloor \frac{n}{k} \rfloor \leqslant |r^{-1}(x)| \leqslant \lceil \frac{n}{k} \rceil$

Procaccia et al. (2008) showed that, for both rules, deciding whether a given bound on misrepresentation can be respected is *NP-complete*.

A.D. Procaccia, J.S. Rosenschein, and A. Zohar. On the Complexity of Achieving Proportional Representation. *Social Choice and Welfare*, 2008.

Summary

This has been a brief introduction to the topic of designing voting rules to elect committees of a given fixed size k.

- multiwinner voting rules with *approval ballots*
- multiwinner voting rules with *ranked ballots*

The focus has been on systematic approaches to designing new rules and on properties of such rules (such as proportionality).

Multiwinner voting reduces to standard *single-winner voting* for $k = 1$ and is equivalent to unit-cost *participatory budgeting* with budget k.

What next? More on preferences over sets of alternatives.