Computational Social Choice 2024

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

[http://www.illc.uva.nl/~ulle/teaching/comsoc/2024/]

Plan for Today

So far we mostly worked with the "*standard model*" of voting theory, where preferences are rankings and we want to elect a single alternative.

Today we will briefly review some *alternative models* of voting:

- variable electorates
- weak orders
- incomplete preferences
- approval sets
- multiwinner voting
- apportionment
- participatory budgeting
- liquid democracy

The focus won't be on results, but on appreciating the rich *design space* available to us when setting up a system for taking collective decisions.

The Standard Model

Given a (fixed) set A of *alternatives* and a (fixed) set $N = \{1, \ldots, n\}$ of *voters*, we studied *voting rules* of this form:

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

In other words:

- input: profile of strict linear orders
- output: nonempty set (ideally: singleton)

Variable Electorates

In our formal model of voting, the number n of voters was always fixed.

But all real-world voting rules we discussed in fact work for electorates of all possible sizes. Could enrich the formal model to account for this:

$$F : \bigcup_{N \subseteq \mathbb{N}} \mathcal{L}(A)^N \to 2^A \setminus \{\emptyset\}$$

Here $\mathbb{N}$ is the "universe" of voters who might vote on a given day; N is any finite subset; and $\mathcal{L}(A)^N$ is the set of functions from N to $\mathcal{L}(A)$.

Exercise: *Explain this new definition of voting rule!*

In this model we can, for instance, define the *reinforcement axiom*, which can differentiate between PSRs and Condorcet extensions:

> *If two disjoint electorates elect overlapping sets of alternatives, then their union should elect the intersection of those sets.*

Remark: Could vary the set of alternatives as well ("*variable agenda*").

Preferences as Weak Orders

By modelling preferences as strict linear orders, we presuppose that a voter will never like two alternatives equally much. Unrealistic.

Could instead work with *weak orders:* rankings of clusters of equally preferred alternatives. Note that strict linear orders are a special case.

Exercise: *When we move from strict to weak orders, how does this affect the impossibilities we observed? Do things get better or worse?*

Incomplete Preferences

You might *prefer* a over b, you might *disprefer* a to b, you might be *indifferent* between them ... or you might be *unable to compare*.

Thus, sometimes preferences will be *incomplete*. Possible reasons:

- voter is unaware of all altenatives
- space of alternatives is huge
- comparing two alternatives is costly
- voter only cares about ranking most preferred alternatives

Exercise: *What would be a natural generalisation of the Borda rule when each voter only ranks her most preferred alternatives?*

Exercise: *What would be a natural generalisation of the Slater rule when each voter only ranks some pairs of alternatives?*

Z. Terzopoulou. *Collective Decisions with Incomplete Individual Opinions*. PhD thesis, ILLC, University of Amsterdam, 2021.

Preferences as Approval Sets

We already saw *approval voting:* Instead of asking voters to rank the alternatives, we ask them to indicate which alternatives they approve.

Now preferences/ballots are *approval sets* $A_i \subseteq A$.

Exercise: *Approval voting is different from k-approval. Explain!*

We can also mix ranked preferences and approval preferences:

- Preferences might be *rankings* with an *approval threshold*.
  Exercise: *How would you define monotonicity in this case?*
- Preferences might be *rankings* but ballots might be *approval sets*.
  Exercise: *How would you define strategyproofness in this case?*

Multiwinner Voting

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

But sometimes we in fact want to elect *multiple winners* . . .

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.

M. Lackner and P. Skowron. *Multi-Winner Voting with Approval Preferences*. Springer, 2023.

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.

In all cases, could use *ranked preferences* or *approval preferences*.

Exercise: *Difference between multiwinner and irresolute voting rule?*

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

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 $\sum_{x \in S} \sum_{i \in N} \mathbb{1}_{x \in A_i} = \sum_{i \in N} |S \cap A_i|$.

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 $\frac{|C|}{n} \geqslant \frac{\ell}{k}$, 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*.

Exercise: *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 $\frac{|C|}{n} \geqslant \frac{1}{k}$ 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 (and $\frac{3}{9} \geqslant \frac{1}{3}$), 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 $\sum_{i \in N} w_1 + \cdots + w_{|S \cap A_i|}$.

Fact: PAV satisfies PJR, but other Thiele rules do not. (proof omitted)

Apportionment

In general elections, we vote for parties, not people. Imagine there are 150 seats, and party p gets 21.87% of the vote. *How many seats for p?* This is known as the problem of *apportionment*.

Exercise: *Explain how this is a special case of multiwinner voting!*

In the Netherlands, we use the method of *D'Hondt* for apportionment:

- Find d such that $\lfloor \#party_1 \,/\, d \rfloor + \cdots + \lfloor \#party_m \,/\, d \rfloor \; = \; \#seats$
- Award $\lfloor \#party_i \,/\, d \rfloor$ seats to party i

Other methods exist. This one tends to favour larger parties.

M.L. Balinski and H.P. Young. *Fair Representation: Meeting the Ideal of One Man, One Vote*. 2nd edition, Bookings Institution Press, 2001.

M. Brill, J.-F. Laslier, and P. Skowron. Multiwinner Approval Rules as Apportionment Methods. *Journal of Theoretical Politics*, 2018.

D. Peters. Online calculator: `https://pref.tools/apportionment/`, 2023.

Participatory Budgeting

A generalisation of multiwinner voting is *participatory budgeting*:

> The city wants to consult residents on how to spend some of its *budget*. There are several *projects*, each with a *cost*. People *vote*. Need a *rule* to choose which projects to fund.

Exercise: *Explain how multiwinner voting is a special case of this!*

S. Rey. *Variations on Participatory Budgeting*. PhD thesis, ILLC, University of Amsterdam, 2023.

Liquid Democracy

The idea of *liquid democracy* has been proposed as a compromise between *direct democracy* and *representative democracy*.

J. Green-Armytage. Direct Voting and Proxy Voting. *Constitutional Political Economy*, 2015.

C. Blum and C.I. Zuber. Liquid Democracy: Potentials, Problems, and Perspectives. *Journal of Political Philosophy*, 2016.

J. Behrens. The Origins of Liquid Democracy. *Liquid Democracy Journal*, 2017.

The Basic Model of Liquid Democracy

The *voters* in $N = \{1, \ldots, n\}$ need to choose an *alternative* from A.

Each voter $i \in N$ either (i) *reports a preference* R_i (e.g., from $\mathcal{L}(A)$) or (ii) *delegates her right to vote* to another voter $j \in N \setminus \{i\}$.

$\Rightarrow$ *delegation graph* $\langle N, \to \rangle$, with the sinks being the *casting voters*

If $\langle N, \to \rangle$ is *acyclic*, we can construct a *profile* $\boldsymbol{R} = (R_1, \ldots, R_n)$, by setting $R_i := R_{i^\star}$ for voter i and the casting voter $i^\star$ with $i \to^* i^\star$. We can then apply our favourite *voting rule* to this preference profile.

When there are *cycles* (which is usually considered highly undesirable), the simplest solution is to assume that the voters involved *abstain*.

Research Questions in Liquid Democracy

What is a reasonable model for how delegation graphs are formed?

- How do voters choose? Link between preferences and delegation?
- Delegation on everything / specific issues / policy areas?

What can be said about the structure of delegation graphs?

- How should we deal with cycles? Interpret them as abstentions?
- Should we be concerned about extreme concentrations of power?
- Should we impose restrictions on delegation for better control?

What voting rule should we use to aggregate cast preferences?

- Should the structure of the graph matter (or just # leaves)?
- Normative characterisation? Epistemic characterisation?

How does liquid democracy perform relative to other approaches?

- Compared to direct / representative democracy?
- Compared to proxy voting without transitivity?

The Bigger Picture

Liquid democracy (as in: transitive proxy voting) is but a particularly salient example for a much broader research agenda currently forming:

- Behrens et al. (2014), creators of the *LiquidFeedback* platform, discuss challenges such as the fair *elicitation of proposals*.
- Brill (2018) outlines a broader research agenda for building new theoretical foundations for *participatory decision making*.
- Grandi (2017) reviews research at the interface of social choice theory with *social network analysis* more generally. This includes research on (discrete) *opinion diffusion* on social networks.

J. Behrens, A. Kistner, A. Nitsche, and B. Swierczek. *The Principles of LiquidFeedback*. Interaktive Demokratie e.V., 2014.

M. Brill. Interactive Democracy. AAMAS-2018 (Blue Sky Ideas Track).

U. Grandi. Social Choice and Social Networks. In U. Endriss (ed.), *Trends in Computational Social Choice*. AI Access, 2017.

Summary

This has been a discussion of various alternative models: we varied the input, the output, and the environment in which voting takes place.

What next? Explainability and logical modelling in social choice.