Fairness in Multiwinner Elections

Julian Chingoma, Jan Maly and Simon Rey

Lecture One

{j.z.chingoma, j.f.maly, s.rey}@uva.nl

June, 2022

Institute of Logic, Language and Computation (ILLC)

University of Amsterdam

Project Overview

Website: All organisational information is available on the website. Plus, the lecture slides: http://simonrey.fr/en/teaching/MWV2022

Project Plan:

- ⋆ Week 1: Four lectures on Tuesday to Friday at 13:00-15:00 (room F2.19).
- ⋆ Week 2: Student presentations on a chosen paper.
- ⋆ Weeks 3-4: Students work on their final paper (with supervision from instructors). Week 3 also includes a plenary session (date, time and venue: TBD) to discuss progress on the final papers.
- ⋆ For both the presentations and final papers, you are free to either pair up (can be different pairs for the two deliverables), or work on your own.
- ⋆ Presentations lengths are 20 minutes (incl. 5 minutes of questions) for individuals, and 40 minutes (incl. 10 minutes of questions) for the pairs.

Final grade: [Presentation + Final Paper] $\rightarrow$ Pass/Fail.

Today's Lecture Plan

Lecture Plan:

- ⋆ Multiwinner election model.
- ⋆ Party-list profiles for apportionment.
- ⋆ Apportionment methods.
- ⋆ Proportionality properties in apportionment setting.
- ⋆ The Proportional Approval Voting (PAV) rule.

Computational Social Choice

Collective decision-making tasks where a group of agents submit some opinion over a set of alternatives, and the goal is to aggregate these opinions into a single collective choice.

Real-world examples:

- ⋆ Electing a single winning candidate in a presidential election.
- ⋆ Matching of agents from different sets, such as matching schools to children, or matching kidney donors to patients.
- ⋆ Division of resources amongst agents (divisible or indivisible), such as sharing a cake, or a bunch of toys, amongst children.
- ⋆ Choosing the set of public projects (e.g., a swimming pool or a park bench), with each having some cost, to be implemented within a given budget (Participatory Budgeting).

Multiwinner Elections

Multiwinner Elections

We focus on scenarios where agents vote to elect multiple winning candidates (some natural number k to be precise). Why study multiwinner elections? Many practical scenarios which can be modelled using multiwinner elections:

- ⋆ An job panel is given a list of potential employees and must be produce a shortlist of k applicants to continue to the next interview stage.
- ⋆ An airline must choose k movies to make available for passengers to view during their flights.
- ⋆ A parliamentary election where k parliament seats are to distributed amongst multiple political parties.

Q: What properties should these rules aim for?

Multiwinner Elections

We focus on scenarios where agents vote to elect multiple winning candidates (some natural number *k* to be precise). Why study multiwinner elections? Many practical scenarios which can be modelled using multiwinner elections:

- ⋆ An job panel is given a list of potential employees and must be produce a shortlist of *k* applicants to continue to the next interview stage.
- ⋆ An airline must choose *k* movies to make available for passengers to view during their flights.
- ⋆ A parliamentary election where *k* parliament seats are to distributed amongst multiple political parties.

Q: What properties should these rules aim for?

For the parliamentary application, the focus is on ***fairness***, and ensuring the views of smaller groups are well-represented in the collective choice. Specifically, we deal with ***proportionality***.

Motivation

There are many design choices one must make when modelling multiwinner elections, and the following are examples to take into consideration:

- ⋆ Voters submit preferences on ***individual*** candidates, and not over sets of candidates. There is work done on *preference extensions*.
- ⋆ We focus on ***approval ballots***, but there are other types such as *rankings*.
- ⋆ The committees we search for are of a ***fixed size*** k. But there is work on *variable-sized multiwinner voting*.

Approval-based Multiwinner Election Model

We define the general approval-based multiwinner model.

- A set of candidates $C = \{a, b, c, \ldots\}$ of size $m \geqslant 2$.
- A set of agents $N = \{1, \ldots, n\}$ of size n.
- Each agent $i \in N$ submits an *approval ballot* $A_i \subseteq C$.
- A vector of these n approval ballots is called an *approval profile* $\boldsymbol{A} = (A_1, \ldots, A_n)$.
- Take $\mathcal{P}_k(C)$ as the set of all committees W of size k (all feasible outcomes).
- An election instance is denoted by the pair $(\boldsymbol{A}, k)$.
- An approval-based multiwinner voting rule $\mathcal{R}$ maps election instances to a non-empty set $\mathcal{R}(\boldsymbol{A}, k)$ of size k committees from $\mathcal{P}_k(C)$. The elements of $\mathcal{R}(\boldsymbol{A}, k)$ are called *winning committees* and denoted as W. Generally, we deal with *irresolute* rules so there may be multiple, tied, winning committees.

Apportionment

Parliamentary Elections

We aim to distribute seats to each party so as to represent society's views in a proportional manner. This task is referred to as ***apportionment***.

Example: let's say some seats must be distributed amongst 5 parties, and 100 agents submit a vote (each agent votes for one party). How should the seats be distributed?

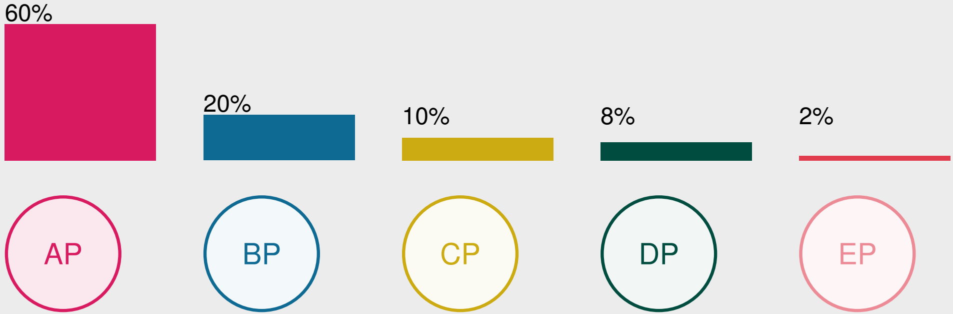

Parliamentary Elections

With 50 seats available, we aim for seat distributions that look, more or less, like this. Every party has a representative in the parliament.

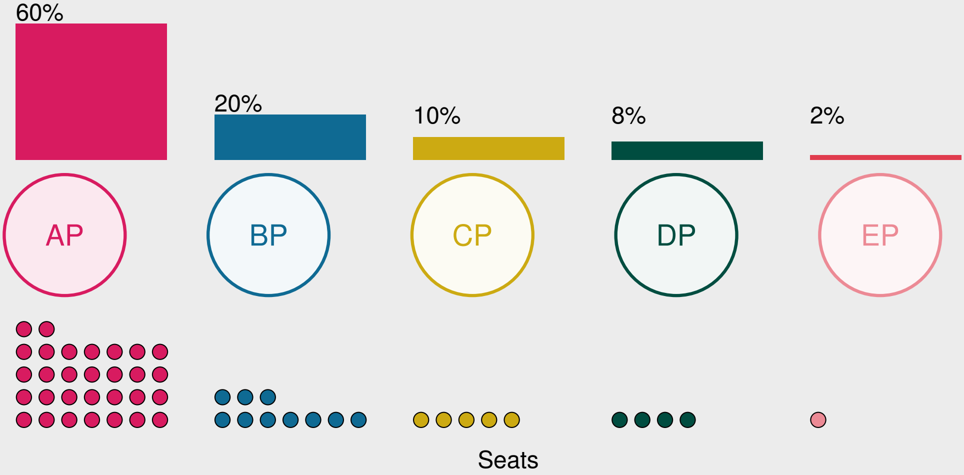

Parliamentary Elections

What if only 10 seats are available? Is it 'fair' if some parties get assigned no seats at all, like DP and EP below?

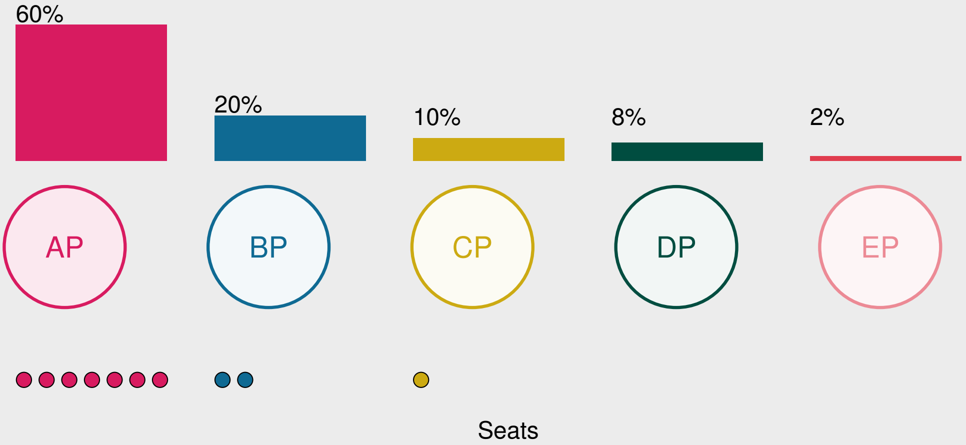

Party-list Multiwinner Elections

How do we model apportionment within the approval-based multiwinner model?

Party-list profile: An approval profile $\boldsymbol{A} = (A_1, \ldots, A_n)$ is a party-list profile if for every pair of agents $i, j \in N$, it holds that either $A_i = A_j$, or that $A_i \cap A_j = \emptyset$. An election instance $(\boldsymbol{A}, k)$ is a party-list election if $\boldsymbol{A}$ is a party-list profile, and for every voter $i \in N$, we have that $|A_i| \geqslant k$.

- ⋆ We restrict agents to vote for one party but this does not have to be the case.
- ⋆ Apportionment can also be extended to *bi-apportionment*, where two dimensions are taken into account (e.g., political and geographical).

Apportionment Methods

D'Hondt Method

First proposed by Thomas Jefferson (3rd United States president), and later developed, independently, by Victor D'Hondt (Belgian professor and lawyer).

We can divide agents and candidates into p disjoint groups, $N = N_1 \cup \ldots \cup N_p$ and $C = P_1 \cup \ldots \cup P_p$. So, for $i \in \{1, \ldots, p\}$, all of the agents in N_i support party P_i.

Definition

The D'Hondt method proceeds in k rounds, in each round, one of the k seats is allocated to some party. For some round r, let $s_i(r)$ be the number of seats that have been allocated to party P_i; then we have that $\sum_{i \in \{1,\ldots,p\}} s_i(r) = r - 1$. The r-th seat is allocated to the party with the highest ratio $|N_i| / s_i(r)+1$. If necessary, use some tiebreaking on parties.

This rule belongs to the class of divisor methods; other divisor methods use different ratios to determine the distribution of seats.

D'Hondt Method

How are the seats allocated by the D'Hondt method for the previous example with 5 parties?

	AP	BP	CP	DP	EP
$\lvert N_i \rvert$	**60**	**20**	**10**	8	2
$\lvert N_i \rvert/2$	**30**	**10**	5	4	1
$\lvert N_i \rvert/3$	**20**	6⅔	3⅓	2⅔	⅔
$\lvert N_i \rvert/4$	**15**	5	2⅔	2	½
$\lvert N_i \rvert/5$	**12**	4	2	1⅗	⅖
$\lvert N_i \rvert/6$	**10**	3⅓	1⅔	1⅓	⅓
$\lvert N_i \rvert/7$	**8 4/7**	2 6/7	1 3/7	1 1/7	2/7
$\lvert N_i \rvert/8$	7½	2½	1¼	1	¼
seats	7	2	1	0	0

Largest Remainder Method

We define the *Largest Remainder Method (LRM)*. It is also known as the Hamilton method as it was proposed by Alexander Hamilton (United States politician). It was used in practice until it was replaced by the D'Hondt/Jefferson method.

Definition

LRM first allocates $\lfloor |N_i| \cdot k / n \rfloor$ seats to every party P_i. Then, it allocates the remaining r seats, with $r < p$, to the r parties with the largest remainders $|N_i| \cdot k / n - \lfloor |N_i| \cdot k / n \rfloor$ (each party gets at most one seat in this second phase).

LRM

Now, let's check how LRM assigns seats with the same example as before.

The row with the standard quota calculation indicates the seats that are initially allocated to each party.

	AP	BP	CP	DP	EP
$\lvert N_i \rvert$	60	20	10	8	2
$\lfloor \lvert N_i \rvert \cdot k/n \rfloor$	6	2	1	0	0
remainder	0	0	0	0.8	0.2
seats	6	2	1	1	0

With a single seat remaining, the remainder values are calculated and as party DP has the highest such value, DP gets the last seat while EP ends up with none.

Properties For Apportionment Methods

Proportionality properties

Here are some examples of desirable properties. The first two are related to proportionality while the latter two conditions are important multiwinner properties.

1. *Lower quota*: Each party P_i with $|N_i|$ votes gets at least $\lfloor |N_i| \cdot k / n \rfloor$ seats.
2. *Upper quota*: Each party P_i with $|N_i|$ votes gets at most $\lceil |N_i| \cdot k / n \rceil$ seats.
3. *Population monotonicity*: An increase in a party's support does not harm it.
4. *Committee monotonicity*: Increase in available seats does not harm any party.

How do the rules measure up against the properties?

- ⋆ The D'Hondt method satisfies lower quota, but not upper quota.
- ⋆ LRM satisfies both lower and upper quota, but neither population nor committee monotonicity.
- ⋆ The D'Hondt method satisfies both population and committee monotonicity.

We will focus on the D'Hondt method, but first, let's detour to the general model.

Thiele Methods and PAV

w-Thiele Methods

Thorvald Thiele (Danish mathematician) proposed the class of Thiele methods where an agent's satisfaction with a committee is based on the number of committee members that they approve of.

Definition (w-Thiele Method)

Given a profile $\boldsymbol{A} = (A_1, \ldots, A_n)$, the following is the *$w$-Thiele method* that is induced by some weight function $w : \mathbb{N}_{>0} \to \mathbb{R}^+$:

$$\operatorname{argmax}_{W \in \mathcal{P}_k(C)} \sum_{i \in N} \sum_{j=1}^{|W \cap A_i|} w(j)$$

- ⋆ Standard *Multiwinner Approval (AV)* rule with $w_{AV}(j) = 1$ for all j.
- ⋆ Approval-based *Chamberlin Courant (CC)* rule with $w_{CC}(j) = 1$ if $j = 1$, and $w_{CC}(j) = 0$ otherwise.

PAV

We focus on the w-Thiele method that plays a prominent role in the proportionality literature. This is the *Proportional Approval Voting (PAV)* rule.

PAV is the w_{PAV}-Thiele method with $w_{PAV}(j) = 1/j$ for all j.

This rule captures the effect of *diminishing returns*, and balances the continued satisfaction of larger groups with trying to satisfy the smaller groups.

- PAV is defined for the general approval-based model. Could we still use PAV for the apportionment task?
- Yes, just use the party-list profiles and this yields some apportionment method.
- The number of seats allocated to a party P_i is then equal to the number of candidates from P_i, that are members of a winning committee W that PAV returns.

PAV

Take the following party-list profile with 5 parties. Here, $k = 10$ seats are to be distributed.

$$C = \{a_1, \ldots, a_k, b_1, \ldots, b_k, c_1, \ldots, c_k, d_1, \ldots, d_k, e_1, \ldots, e_k\}$$

100 agents have their following approvals:

- ⋆ $A_i = \{a_1, \ldots, a_k\}$ for agents $i \in \{1, \ldots, 60\}$.
- ⋆ $A_i = \{b_1, \ldots, b_k\}$ for agents $i \in \{61, \ldots, 80\}$.
- ⋆ $A_i = \{c_1, \ldots, c_k\}$ for agents $i \in \{81, \ldots, 90\}$.
- ⋆ $A_i = \{d_1, \ldots, d_k\}$ for agents $i \in \{91, \ldots, 98\}$.
- ⋆ $A_i = \{e_1, \ldots, e_k\}$ for agents $i \in \{99, 100\}$.

PAV returns the committees with 7 candidates from $\{a_1, \ldots, a_k\}$, 2 candidates from $\{b_1, \ldots, b_k\}$, and 1 candidate from $\{c_1, \ldots, c_k\}$.

PAV

Note that the committees returned by PAV induce a seat distribution identical to that of the D'Hondt method. The following result establishes the connection between the two methods that were defined for different settings.

Theorem (Brill, Laslier, and Skowron, 2018)

PAV extends the D'Hondt method of apportionment to the general approval-based multiwinner setting.

We shall briefly sketch a proof of this. To do so, we need a few more definitions.

Sequential Thiele Methods

First, we define the sequential variants of the w-Thiele methods.

Definition (*seq-w*-Thiele Method)

For every w-Thiele method induced by a weight function w, we can define a sequential variant *seq-w-Thiele method* to return a committee W.

- Start with an empty committee $W_0 = \emptyset$.
- In every round r, the *seq-w*-Thiele method set $W_r = W_{r-1} \cup \{c\}$ by adding a candidate $c \in C$ that maximises:

$$\sum_{i \in N} \sum_{j=1}^{|A_i \cap W_{r-1} \cup \{c\}|} w(j).$$

  Use some tiebreaking, if necessary.
- Return the committee W_k.

We refer to the *seq-w_{PAV}*-Thiele Method as *sequential PAV*.

PAV and D'Hondt

Now, with sequential PAV defined, we can use the following result.

Proposition (Brill, Laslier, and Skowron, 2018)

Let w be a weight function. Then the apportionment method induced by the w-Thiele method coincides with the apportionment method induced by the seq-w-Thiele method.

Now, all we need to show is that the D'Hondt method coincides with the apportionment method defined by the sequential PAV rule.

PAV and D'Hondt

For that, let's precisely formulate divisor methods using *divisor sequences*.

Take $d = (d(0), d(1), d(2), \ldots)$ to be a sequence such that $0 < d(i) \leqslant d(i+1)$ for all $i \in \mathbb{N} \cup \{0\}$. A divisor method $M(d)$ based on d works as follows: Start with no seats allocated, and in iterations, it allocates a seat to party P_i that maximises $|N_i|/d(s_i)$ where s_i is the number of seats assigned to P_i in prior iterations. The D'Hondt method $M(d_{dh})$ is the divisor method based on $d_{dh} = (1, 2, 3, \ldots)$.

Now, we will use the previous proposition to show that every w-Thiele method with a non-increasing weight function induces some divisor method. For a weight function w (recall we deal with positive weight functions), take the sequence $d_w = (d_w(0), d_w(1), \ldots)$ such that $d_w(s) = 1/w(s+1)$ for all $s \in \mathbb{N} \cup \{0\}$. So PAV's weight function $w_{PAV} = 1/j$ for all $j \in \mathbb{N}_{>0}$, gives a sequence $d_{w_{PAV}} = (1, 2, 3, \ldots)$.

PAV and D'Hondt

Theorem (Brill, Laslier, and Skowron, 2018)

The w-Thiele method which is induced by some non-increasing weight function w coincides with the divisor method based on $d_w = (1/w(1), 1/w(2), 1/w(3), \ldots)$.

Proof sketch.

- ⋆ In iterations, both rules assign seats. The $M(d_w)$ divisor method gives a seat to a party that maximises $|N_i|/d_w(s_i)$ where s_i is the number of seats previously allocated to party P_i.
- ⋆ At the same iteration, the *seq*-w-Thiele method on the party-list election instance, maximises *marginal increase in total agent satisfaction*, and for a party P_i, the equals $|N_i| \cdot w(s_i + 1)$. So, take $d_w(s) = 1/w(s+1)$.
- ⋆ Observe that the sequences $(|N_i|/d_w(s_i+1))_{i \in [p]}$ and $(|N_i| \cdot w(s_i + 1))_{i \in [p]}$ coincide.

□

PAV and D’Hondt

Proportional apportionment rules seem like good candidates to generalise to ensure proportionality in the more general setting.

- The D’Hondt method performs well so this is good news for both PAV and sequential PAV. But, these two are not the only D’Hondt extensions.
- There exists other rules, such as *Phragmën’s Sequential Rule* (see Section 2.5 in the book), that extend the D’Hondt method.
- However, it turns out that PAV is the only extension of the D’Hondt method that also provides other properties considered desirable for more general multiwinner elections.

PAV and D'Hondt

Here are the other multiwinner election properties:

- ⋆ Neutrality: treat all candidates the same.
- ⋆ Anonymity: treat all agents the same.
- ⋆ Consistency: if a committee W is preferred to W', in both elections $(\boldsymbol{A}, k)$ and $(\boldsymbol{A'}, k)$, then W should be preferred to W' in the election $(\boldsymbol{A} + \boldsymbol{A'}, k)$.
- ⋆ Continuity: if a committee W is preferred to W' for some election $(\boldsymbol{A'}, k)$, then there exists a positive integer t such that W is preferred to W' in election instance $(\boldsymbol{A} + t\boldsymbol{A'}, k)$.

Theorem (Lackner and Skowron, 2021)

PAV is the unique extension of the D'Hondt method that satisfies neutrality, anonymity, consistency and continuity.

References

1. Markus Brill, Paul Gölz, et al. (2020). “Approval-based apportionment”. In: *Proceedings of the AAAI Conference on Artificial Intelligence*
2. Friedrich Pukelsheim (2017). *Proportional Representation: Apportionment Methods and Their Applications*. 2nd ed. Springer International Publishing
3. Michel L Balinski and H Peyton Young (2010). *Fair representation: meeting the ideal of one man, one vote*. Brookings Institution Press
4. Thorvald N. Thiele (1895). “Om flerfoldsvalg”. In: *Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandlinger*
5. Markus Brill, Jean-François Laslier, and Piotr Skowron (2018). “Multiwinner approval rules as apportionment methods”. In: *Journal of Theoretical Politics*
6. Martin Lackner and Piotr Skowron (2021). “Consistent approval-based multi-winner rules”. In: *Journal of Economic Theory*

Wrapping Up

Lecture Recap

Today's summary:

- ⋆ The multiwinner election model and the apportionment setting.
- ⋆ Some apportionment methods: the D'Hondt Method and the Largest Remainder Method.
- ⋆ Lower and upper quota (plus some monotonicity properties).
- ⋆ The PAV rule in relation to the apportionment setting.

Next time

Next time:

- ⋆ More on the general model and why PAV is a good rule.
- ⋆ Some important proportionality axioms for the general model