Computational Social Choice 2022

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

[http://www.illc.uva.nl/~ulle/teaching/comsoc/2022/]

Plan for Today

Last time we introduced the *axiomatic method* for JA and discussed the *List-Pettit Theorem*, establishing an *impossibility* for three specific axioms (*anonymity*, *neutrality*, *independence*) together with *collective rationality*.

Today we further explore the axiomatic method:

- Approaches to *circumventing the impossibility*, by relaxing some of our requirements and/or considering special cases
- Systematic understanding of the power of the axioms, particularly the independence axiom, by introducing the concept of a *winning coalition*
- Axiomatic *characterisations* of aggregation rules, specifically *quota rules*

Much of this material is covered in the general references listed below.

C. List. The Theory of Judgment Aggregation: An Introductory Review. *Synthese*, 2012.

U. Endriss. Judgment Aggregation. In F. Brandt *et al.* (eds.), *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

Reminder

Last time we proved:

Theorem 1 (List and Pettit, 2002) *No judgment aggregation rule for an agenda Φ with $\{p, q, p \wedge q\} \subseteq \Phi$ that is anonymous, neutral, and independent can guarantee outcomes that are complete and consistent.*

Exercise: *What do you do in such a situation?*

C. List and P. Pettit. Aggregating Sets of Judgments: An Impossibility Result. *Economics and Philosophy*, 2002.

Circumventing the Impossibility

If we are prepared to relax some of our requirements, we may be able to circumvent the impossibility and successfully aggregate judgments.

Next, we will explore some such possibilities:

- Relax the *universal domain assumption:* maybe not *every* logically possible profile will materialise in practice?
- Relax *collective rationality:* we won't compromise on collective consistency, but we might want to relax collective *completeness*.
- Relax the *axioms* (we won't treat this systematically):
  - *Anonymity:* maybe some agents are smarter than others?
  - *Neutrality:* maybe it actually is ok to treat, say, atomic propositions differently from conjunctions?
  - *Independence:* there *are* logical dependencies between propositions; so why not allow them to affect aggregation?

Domain Restriction 1: Unidimensional Alignment

Call a profile *unidimensionally aligned* if we can order the agents such that, for each (positive) proposition $\varphi \in \Phi$, the agents *accepting* φ are either all to the *left* or all to the *right* of those *rejecting* φ. Example:

	1	2	3	4	5	(Majority)
p	Yes	Yes	No	No	No	(No)
q	No	No	No	No	Yes	(No)
$p \to q$	No	No	Yes	Yes	Yes	(Yes)

List (2003) showed that under this *domain restriction* we can satisfy all our axioms and be consistent (and complete if n is odd):

Proposition 2 (List, 2003) *For any unidimensionally aligned profile, the majority rule will return a consistent outcome.*

C. List. A Possibility Theorem on Aggregation over Multiple Interconnected Propositions. *Mathematical Social Sciences*, 2003.

Proof

For ease of exposition, suppose the number n of individuals is odd.

Here is again our example, for illustration:

	1	2	3	4	5	(Majority)
p	Yes	Yes	No	No	No	(No)
q	No	No	No	No	Yes	(No)
$p \to q$	No	No	Yes	Yes	Yes	(Yes)

Call the $\lceil \frac{n}{2} \rceil$th agent in this order the *median agent*.

(1) By definition, for each φ in the agenda, at least $\lceil \frac{n}{2} \rceil$ agents (a majority) accept φ *iff* the median agent does.

(2) As the judgment set of the median agent is consistent, so is the collective judgment set under the majority rule. ✓

Domain Restriction 2: Value Restriction

A set $X \subseteq \Phi$ is called *minimally inconsistent* if it is inconsistent and every proper subset $Y \subset X$ is consistent.

Call a profile $\boldsymbol{J}$ *value-restricted* if for every *mi-set* $X \subseteq \Phi$ there exist distinct $\varphi_X, \psi_X \in X$ such that $\{\varphi_X, \psi_X\} \subseteq J_i$ for *no* agent $i \in N$.

Proposition 3 (Dietrich and List, 2010) *For any value-restricted profile, the majority rule will return a consistent outcome.*

F. Dietrich and C. List. Majority Voting on Restricted Domains. *Journal of Economic Theory*, 2010.

Proof

Assume profile $\boldsymbol{J} = (J_1, \ldots, J_n)$ is *value-restricted*.

For the sake of contradiction, suppose $J := F_{\mathrm{maj}}(\boldsymbol{J})$ is *inconsistent*.
Then there exists a set $X \subseteq J$ that is *minimally inconsistent*.

By value restriction of $\boldsymbol{J}$, there exist two formulas $\varphi_X, \psi_X \in X$ such that *no agent accepts both* of them in $\boldsymbol{J}$.

Note that from $\varphi_X, \psi_X \in X$ and $X \subseteq J$, we get $\varphi_X, \psi_X \in J$.

Hence, there must have been a *strict majority* for both φ_X and ψ_X, meaning that at least *one agent* must have *accepted both*.

Contradiction! ✓

Relaxing Axioms: Premise-Based Aggregation

Most pragmatic aggregation rules circumvent the impossibility theorem by sacrificing *independence*. Premise-based rules also relax *neutrality*.

Exercise: *Why does it say "[p]remise-based rules" (plural) above?*

Recall the *premise-based rule* F_{pre} for premises Φ_p and conclusions Φ_c:

$$F_{\mathrm{pre}}(\boldsymbol{J}) \quad = \quad \Delta \,\cup\, \{\varphi \in \Phi_c \mid \Delta \models \varphi\},$$
$$\text{where } \Delta = \{\varphi \in \Phi_p \mid \#N_\varphi^{\boldsymbol{J}} > \frac{n}{2}\}$$

Recall: If we assume that

- the set of *premises* is the set of *literals* in the agenda,
- the agenda Φ is is *closed under propositional letters*, and
- the number n of individuals is *odd*,

then $F_{\mathrm{pre}}(\boldsymbol{J})$ will always be *consistent* and *complete*.

Relaxing Neutrality: Unanimous Rule with Defaults

Consider this relaxed form of neutrality of an aggregation rule F:

> For any φ, ψ in the agenda Φ with $\varphi \models \psi$ and any profile $\boldsymbol{J}$ with $N_\varphi^{\boldsymbol{J}} = N_\psi^{\boldsymbol{J}}$ we should have $\varphi \in F(\boldsymbol{J}) \Rightarrow \psi \in F(\boldsymbol{J})$.

Let $F^\star$ be the rule that accepts, for every pair $(\varphi, \neg\varphi)$ in the agenda, φ if it has unanimous support and (the "default") $\neg\varphi$ otherwise.

Proposition 4 (Terzopoulou and Endriss, 2021) *For the agenda* $\Phi = \{p, \neg p, q, \neg q, p \wedge q, \neg(p \wedge q)\}$, *the rule* $F^\star$ *satisfies anonymity, relaxed neutrality, independence, completeness, and consistency.*

So by relaxing neutrality, we circumvent the List-Pettit impossibility.

Proof of this specific result is immediate. But the cited paper explores the idea of relaxing neutrality much more systematically.

Z. Terzopoulou and U. Endriss. Neutrality and Relative Acceptability in Judgment Aggregation. *Social Choice and Welfare*, 2021.

Relaxing Completeness and Anonymity: Oligarchies

We could relax completeness to *deductive closure*, requiring merely that $\varphi \in \Phi$ and $J \models \varphi$ should imply $\varphi \in J$ for the outcome J.

The *oligarchic rule* for coalition $C \subseteq N$ is the rule that accepts φ *iff* everyone in C does. Special cases:

- dictatorial rule: $|C| = 1$
- unanimous rule: $C = N$

It is easy to check that any oligarchic rule satisfies:

- collective *consistency* and *deductive closure*
- *neutrality* and *independence*
- but *not anonymity* (except for $C = N$)
- and also *not completeness* (except for $|C| = 1$)

Gärdenfors (2006) gives a more precise axiomatic characterisation.

P. Gärdenfors. A Representation Theorem for Voting with Logical Consequences. *Economics and Philosophy*, 2006.

Relaxing Completeness: Supermajority Rules

Recall *uniform quota rules* F_λ with $\lambda \in \{0, 1, \ldots, n{+}1\}$:

$$F_\lambda(\boldsymbol{J}) \quad = \quad \{\varphi \in \Phi \mid \#N_\varphi^{\boldsymbol{J}} \geqslant \lambda\}$$

Uniform quota rules with $\lambda > \frac{n}{2}$ are known as *supermajority rules*.

Recall: High quotas can restore consistency (sacrificing completeness).

Restricted Agendas

Suppose the agenda Φ consists of *literals only*. Then the majority rule always returns an outcome that is consistent (and complete for odd n).

This is a trivial example for an *agenda property* ensuring consistency.

More interesting examples to be discussed later on in the course.

Domain Restrictions vs. Agenda Properties

Note the difference:

- *Domain restrictions* apply to *profiles* (for arbitrary agendas).

  Examples: unidimensionally aligned / value-restricted profiles

- *Agenda properties* apply to *agendas* (restricting their structure).

  Example: literals-only agendas

Axiomatic Characterisation of Rules

The *axiomatic method* is not only good for impossibility results.

We can also get *characterisation results*, i.e., unique characterisations of (families of) aggregation rules in terms of axioms:

- *useful* to argue for a rule in terms of fundamental principles
- literature *still sparse:* for many rules we don't have characterisations
- but for *quota rules*, clear picture with nice and easy results ($\hookrightarrow$)

Axioms

Last time's axioms, and a new one:

- *Anonymity:* Treat all agents symmetrically!

  For any profile $\boldsymbol{J}$ and any permutation $\pi : N \to N$, we should have $F(J_1, \ldots, J_n) = F(J_{\pi(1)}, \ldots, J_{\pi(n)})$.

- *Neutrality:* Treat all propositions symmetrically!

  For any φ, ψ in the agenda Φ and any profile $\boldsymbol{J}$ with $N_\varphi^{\boldsymbol{J}} = N_\psi^{\boldsymbol{J}}$ we should have $\varphi \in F(\boldsymbol{J}) \Leftrightarrow \psi \in F(\boldsymbol{J})$.

- *Independence:* Should be able to decide on one issue at a time!

  For any φ in the agenda Φ and any profiles $\boldsymbol{J}$ and $\boldsymbol{J'}$ with $N_\varphi^{\boldsymbol{J}} = N_\varphi^{\boldsymbol{J}'}$ we should have $\varphi \in F(\boldsymbol{J}) \Leftrightarrow \varphi \in F(\boldsymbol{J'})$.

- *Monotonicity:* Additional support should not harm a formula!

  For any profile $\boldsymbol{J}$, agent i, judgment set $J_i' \in \mathcal{J}(\Phi)$, and formula $\varphi \in J_i' \setminus J_i$, we should have $\varphi \in F(\boldsymbol{J}) \Rightarrow \varphi \in F(\boldsymbol{J}_{-i}, J_i')$.

Winning Coalitions

Alternative Definition: Rule F is *independent* if there exists a family of sets (*winning coalitions*) of agents $\mathcal{W}_\varphi \subseteq 2^N$, one for each $\varphi \in \Phi$, such that for all profiles $\boldsymbol{J} \in \mathcal{J}(\Phi)^n$ we have $\varphi \in F(\boldsymbol{J})$ *iff* $N_\varphi^{\boldsymbol{J}} \in \mathcal{W}_\varphi$.

Recall: $N_\varphi^{\boldsymbol{J}} = \{i \in N \mid \varphi \in J_i\}$ is the set of supporters of φ in $\boldsymbol{J}$.

Now suppose F is independent and defined by $\{\mathcal{W}_\varphi\}_{\varphi \in \Phi}$. Then:

- F is *anonymous* *iff* $\mathcal{W}_\varphi$ is closed under equinumerosity: $C \in \mathcal{W}_\varphi$ and $|C| = |C'|$ entail $C' \in \mathcal{W}_\varphi$ for all $C, C' \subseteq N$ and all $\varphi \in \Phi$.
- F is *monotonic* *iff* $\mathcal{W}_\varphi$ is upward closed: $C \in \mathcal{W}_\varphi$ and $C \subseteq C'$ entail $C' \in \mathcal{W}_\varphi$ for all $C, C' \subseteq N$ and all $\varphi \in \Phi$.
- F is *complete* *iff* $\mathcal{W}_\varphi$ is maximal: $C \in \mathcal{W}_\varphi$ or $\overline{C} \in \mathcal{W}_{\sim\varphi}$ for all C, φ.
- F is *complement-free* *iff* $C \notin \mathcal{W}_\varphi$ or $\overline{C} \notin \mathcal{W}_{\sim\varphi}$ for all C, φ.

What about neutrality?

A Subtlety about Neutrality

Recall the formal definition of neutrality:

- For any φ, ψ in the agenda Φ and any profile $\boldsymbol{J}$ with $N_\varphi^{\boldsymbol{J}} = N_\psi^{\boldsymbol{J}}$ we should have $\varphi \in F(\boldsymbol{J}) \Leftrightarrow \psi \in F(\boldsymbol{J})$.

Intuitively, this says that all formulas should be treated symmetrically. Thus, we (almost) get:

- For any independent rule F, it is the case that F is *neutral* *iff* $\mathcal{W}_\varphi = \mathcal{W}_\psi$ for all formulas $\varphi, \psi \in \Phi$.

But note that neutrality does not "bite" for trivial agendas such as $\Phi = \{p, \neg p\}$: it holds vacuously, as there exists no admissible profile in which the same agents accept p and $\neg p$. But for nontrivial agendas, above characterisation holds (see reference below for details).

U. Endriss. Judgment Aggregation. In F. Brandt *et al.* (eds.), *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

Reminder

A *quota rule* F_q is defined by a function $q : \Phi \to \{0, 1, \ldots, n{+}1\}$:

$$F_q(\boldsymbol{J}) \quad = \quad \{\varphi \in \Phi \mid \#N_\varphi^{\boldsymbol{J}} \geqslant q(\varphi)\}$$

F_q is called *uniform* if $q \equiv \lambda$ for a fixed number $\lambda \in \{0, 1, \ldots, n{+}1\}$.

Axiomatic Characterisation of Quota Rules

We are now ready to prove this simple characterisation result:

Proposition 5 (Dietrich and List, 2007) *An aggregation rule is anonymous, independent, and monotonic iff it is a quota rule.*

Proof: Right-to-left direction immediate. ✓ Left-to-right direction easy when thinking about axioms in terms of winning coalitions. ✓

Thus, for nontrivial agendas (avoiding the subtlety with neutrality):

Corollary 6 *An aggregation rule is anonymous, neutral, independent, and monotonic (= ANIM) iff it is a uniform quota rule.*

F. Dietrich and C. List. Judgment Aggregation by Quota Rules: Majority Voting Generalized. *Journal of Theoretical Politics*, 2007.

Axiomatic Characterisation of the Majority Rule

Increasing the quota preserves complement-freeness, while lowering it preserves completeness. For odd n, we can get both for $\lambda = \lceil \frac{n}{2} \rceil = \lfloor \frac{n}{2} \rfloor$:

Proposition 7 *For odd n, an aggregation rule is ANIM, complete, and complement-free iff it is the (strict) majority rule.*

Remark: Note the close connection to May's Theorem!

For even n, neither the strict nor the weak majority rule are both complete and complement-free. Thus:

Proposition 8 *For even n, there can be no aggregation rule that is ANIM, complete, and complement-free.*

Summary

We saw some more uses of the *axiomatic method*, namely:

- ways of *circumventing the basic impossibility theorem* by relaxing domain assumptions, collective rationality, and axioms
- reformulation of basic axioms in terms of restrictions to families of *winning coalitions* (this will get used heavily in future lectures)
- axiomatic *characterisation* of *quota rules* (including *majority rule*)

What next? JA as a modelling language (specifically: preferences).