Computational Social Choice: Spring 2015

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

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*, 187(1):179–207, 2012.

U. Endriss. Judgment Aggregation. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A.D. Procaccia (eds.), *Handbook of Computational Social Choice*. CUP, 2015.

Reminder

Last time we proved:

Theorem 1 (List and Pettit, 2002) *No judgment aggregation rule for an agenda Φ with $\{p, q, p \wedge q\} \subseteq \Phi$ satisfies all of the axioms of anonymity, neutrality, independence, completeness, and consistency.*

- What to do?

C. List and P. Pettit. Aggregating Sets of Judgments: An Impossibility Result. *Economics and Philosophy*, 18(1):89–110, 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 (other) *axioms* (we won't treat this systematically):
  - *Anonymity:* maybe some agents are smarter than others?
  - *Neutrality:* maybe it is actually 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*, 45(1):1–13, 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 individual according to our left-to-right ordering (establishing unidimensional alignment) the *median individual*.

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

(2) As the judgment set of the median individual 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 $i \in \mathcal{N}$.

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*, 145(2):512–543, 2010.

Proof

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

For the sake of contradiction, suppose $J := F_{\text{maj}}(\boldsymbol{J})$ 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 $\boldsymbol{J}$.

Now 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*.

The *premise-based procedure* F_{pre} for premises Φ_p and conclusions Φ_c:

$$F_{\text{pre}}(\boldsymbol{J}) \quad = \quad \Delta \cup \{\varphi \in \Phi_c \mid \Delta \models \varphi\},$$

where $\Delta = \{\varphi \in \Phi_p \mid |N_\varphi^{\boldsymbol{J}}| > \frac{n}{2}\}$

and $\models$ denotes logical consequence

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 $F_{\text{pre}}(\boldsymbol{J})$ will always be *consistent* and *complete*. (You see why?)

Relaxing Completeness and Anonymity: Oligarchies

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

The *oligarchic rule* for the set of individuals $N \subseteq \mathcal{N}$ is the rule that accepts φ *iff* everyone in N does. Special cases:

- dictatorial rule: $|N| = 1$
- unanimous rule: $N = \mathcal{N}$

It is easy to check that any oligarchic rule satisfies:

- collective *consistency* and *deductive closure*,
- *neutrality* and *independence*,
- but *not anonymity* (except for $N = \mathcal{N}$),
- and also *not completeness* (except for $|N| = 1$).

Gärdenfors (2006) gives a more precise axiomatic characterisation.

P. Gärdenfors. A Representation Theorem for Voting with Logical Consequences. *Economics and Philosophy*, 22(2):181–190, 2006.

Relaxing Completeness: Supermajority Rules

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

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

Call a uniform quota rule with quota $\lambda > \frac{n}{2}$ a *supermajoritry rule*.

It is intuitively clear that a uniform quota rule with a sufficiently high quota will be consistent (to see this, just consider $\lambda = n$).

▶ But how high is high enough?

Safety of the Agenda

Call agenda Φ *safe* for rule F if $F(\boldsymbol{J})$ is consistent for all $\boldsymbol{J} \in \mathcal{J}(\Phi)^n$.

Proposition 4 (Dietrich and List, 2007) *Let Φ be an agenda and let k be the size of the largest minimally inconsistent subset of Φ. Then every uniform quota rule F_λ with quota $\lambda > \frac{k-1}{k} \cdot n$ is safe for Φ.*

Recall: X is a *mi-set* *iff* X is inconsistent but all $Y \subset X$ are consistent.

Proof: Suppose F_λ is *not safe* for Φ: $F_\lambda(\boldsymbol{J})$ is inconsistent for some $\boldsymbol{J}$.

Consider any mi-set $X \subseteq F_\lambda(\boldsymbol{J})$. We know $|X| \leqslant k$.

Each $\varphi \in X$ is accepted by $\geqslant \lambda$ agents. So #acceptances $\geqslant |X| \cdot \lambda$.

Thus, some *agent accepts* $\geqslant \frac{|X| \cdot \lambda}{n} > \frac{k-1}{k} \cdot |X|$, i.e., *all formulas* in X.

Contradiction! ✓

F. Dietrich and C. List. Judgment Aggregation by Quota Rules: Majority Voting Generalized. *Journal of Theoretical Politics*, 19(4):391–424, 2007.

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: agendas with "small" minimally inconsistent subsets

Axiomatic Characterisation of Rules

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

We can also obtain *characterisation results:* 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 (some invented very recently) we do not yet have axiomatic characterisations
- but for *quota rules*, clear picture with nice and easy results

Axioms

Last time's axioms, and a new one:

- *Anonymity:* Treat all individuals symmetrically!

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

- *Neutrality:* Treat all propositions symmetrically!

  Formally: for any φ, ψ in the agenda Φ and any profile $\boldsymbol{J}$, if for all $i \in \mathcal{N}$ we have $\varphi \in J_i \Leftrightarrow \psi \in J_i$, then $\varphi \in F(\boldsymbol{J}) \Leftrightarrow \psi \in F(\boldsymbol{J})$.

- *Independence:* Only the "pattern of acceptance" should matter!

  Formally: for any φ in the agenda Φ and any profiles $\boldsymbol{J}$ and $\boldsymbol{J'}$, if $\varphi \in J_i \Leftrightarrow \varphi \in J'_i$ for all $i \in \mathcal{N}$, then $\varphi \in F(\boldsymbol{J}) \Leftrightarrow \varphi \in F(\boldsymbol{J'})$.

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

  Formally: for any $\varphi \in \Phi$ and profiles $\boldsymbol{J}, \boldsymbol{J'}$, if $\varphi \in J'_{i^\star} \setminus J_{i^\star}$ for some $i^\star$ and $J_i = J'_i$ for all $i \neq i^\star$, then $\varphi \in F(\boldsymbol{J}) \Rightarrow \varphi \in F(\boldsymbol{J'})$.

Winning Coalitions

Alternative Definition: Rule F is *independent* if there exists a family of sets (*winning coalitions*) of agents $\mathcal{W}_\varphi \subseteq 2^{\mathcal{N}}$, 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: $N_\varphi^{\boldsymbol{J}} = \{i \in \mathcal{N} \mid \varphi \in J_i\}$ 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 $C, C' \subseteq \mathcal{N}$ 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 $C, C' \subseteq \mathcal{N}$ and all $\varphi \in \Phi$.
- F is *complete* *iff* $\mathcal{W}_\varphi$ is maximal: $C \in \mathcal{W}_\varphi$ or $\overline{C} \in \mathcal{W}_\varphi$ for all C, φ.
- F is *complement-free* *iff* $C \notin \mathcal{W}_\varphi$ or $\overline{C} \notin \mathcal{W}_\varphi$ 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 $\boldsymbol{J}$, if for all $i \in \mathcal{N}$ we have $\varphi \in J_i \Leftrightarrow \psi \in J_i$, then $\varphi \in F(\boldsymbol{J}) \Leftrightarrow \psi \in F(\boldsymbol{J})$.

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 (details are worked out the chapter below).

U. Endriss. Judgment Aggregation. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A.D. Procaccia (eds.), *Handbook of Computational Social Choice*. CUP, 2015.

Aside: Systematicity

Older papers in JA often work with the axiom of *systematicity*:

- For any two formulas $\varphi, \psi \in \Phi$ and any two profiles $\boldsymbol{J}$ and $\boldsymbol{J'}$, if $\varphi \in J_i \Leftrightarrow \psi \in J'_i$ for all $i \in \mathcal{N}$, then $\varphi \in F(\boldsymbol{J}) \Leftrightarrow \psi \in F(\boldsymbol{J'})$.

Intuitively, this is *neutrality* + *independence* (but for trivial agendas this, again, is not exactly correct).

Downsides of using systematicity:

- does not differentiate between two conceptually different aspects
- does not correspond to classical axioms in preference aggregation

Upside of using systematicity:

- avoids subtlety: easier to work with winning coalitions

Open Research Questions: Neutrality deserves to be studied more. What is the cleanest definition? What are natural relaxations of the axiom (reflecting objective differences between different formulas)?

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

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

Proof: Immediate from characterisation using 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.*

Observe: High/low quotas help compl.-freeness/completness. Thus:

Proposition 7 *For n even, there can be no aggregation rule that is ANIM, complete, and complement-free.*

Proposition 8 *For n odd, an aggregation rule is ANIM, complete, and complement-free iff it is the (strict) majority rule.*

F. Dietrich and C. List. Judgment Aggregation by Quota Rules: Majority Voting Generalized. *Journal of Theoretical Politics*, 19(4):391–424, 2007.

Summary

The main theme has been the *axiomatic method*. Specifically:

- 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* (will get used heavily in future lectures)
- Axiomatic *characterisation* of *quota rules*

What next?

We will see more *specific aggregation rules*, including rules that are *hard to compute* (so far, everything was computationally easy!).