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

This will be a first introduction to *judgment aggregation* (JA), starting from an example originally discussed in legal theory:

- motivating example: *doctrinal paradox*
- *formal model* for judgment aggregation
- some *specific aggregation rules* to use in practice
- introduction to the *axiomatic method*
- the *impossibility theorem* of List and Pettit

Most of this material is covered in the expository papers cited below.

We'll also talk about *homework* requirements and the *mini-projects*.

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.

Example: The Doctrinal Paradox

A court with three judges is considering a case in contract law.

Legal doctrine stipulates that the defendant is *liable* (r) iff the contract was *valid* (p) and has been *breached* (q): $r \leftrightarrow p \wedge q$.

	p	q	r
Judge 1	Yes	Yes	Yes
Judge 2	No	Yes	No
Judge 3	Yes	No	No

Exercise: *Should the court pronounce the defendant liable?*

L.A. Kornhauser and L.G. Sager. The One and the Many: Adjudication in Collegial Courts. *California Law Review*, 1993.

Why Paradox?

So why is this example usually referred to as a “paradox”?

	p	q	$p \wedge q$
Agent 1	Yes	Yes	Yes
Agent 2	No	Yes	No
Agent 3	Yes	No	No
Majority	Yes	Yes	No

Explanation 1: Two natural aggregation rules, the *premise-based rule* and the *conclusion-based rule*, produce *different* outcomes.

Explanation 2: Each individual judgment is *logically consistent*, but the collective judgment returned by the (natural) *majority rule* is *not*.

In philosophy, this is also known as the *discursive dilemma* of choosing between *responsiveness* to the views of decision makers (by respecting majority decisions) and the *consistency* of collective decisions.

The Model

Notation: Let $\sim\varphi := \varphi'$ if $\varphi = \neg\varphi'$ and let $\sim\varphi := \neg\varphi$ otherwise.

An *agenda* Φ is a finite nonempty set of propositional formulas (w/o double negation) closed under complementation: $\varphi \in \Phi \;\Rightarrow\; \sim\varphi \in \Phi$.

A *judgment set* J on an agenda Φ is a subset of Φ. We call J:

- *complete* if $\varphi \in J$ or $\sim\varphi \in J$ for all $\varphi \in \Phi$
- *complement-free* if $\varphi \notin J$ or $\sim\varphi \notin J$ for all $\varphi \in \Phi$
- *consistent* if there exists an assignment satisfying all $\varphi \in J$

Let $\mathcal{J}(\Phi)$ be the set of all complete and consistent subsets of Φ.

Now a finite set of *agents* $N = \{1, \ldots, n\}$, with $n \geqslant 2$, express judgments on the formulas in Φ, producing a *profile* $\boldsymbol{J} = (J_1, \ldots, J_n)$.

A (resolute) *aggregation rule* for an agenda Φ and a set of n agents is a function mapping any profile of complete and consistent individual judgment sets to a single collective judgment set: $F : \mathcal{J}(\Phi)^n \to 2^\Phi$.

Variants of the Model

Several variants of this *formula-based model* of JA have been proposed.

In the model of *binary aggregation with integrity constraints* we have:

- issues (w/o internal structure) instead of agenda formulas
- one integrity constraint describing dependencies between issues

Example: Rather than working with agenda items p, q, and $p \wedge q$, we could use p, q, and r and impose the integrity constraint $r \leftrightarrow p \wedge q$.

It is also possible to combine the two ideas and allow for an external integrity constraint on top of complex agenda formulas.

These different models are essentially equivalent, although there are subtle differences. We will switch models when convenient.

See the paper cited below for a systematic overview of model variants.

U. Endriss, R. de Haan, J. Lang, and M. Slavkovik. The Complexity Landscape of Outcome Determination in Judgment Aggregation. *JAIR*, 2020.

Useful Notation

Let $N_\varphi^{\boldsymbol{J}}$ denote the *coalition* of *supporters* of φ in $\boldsymbol{J}$, i.e., the set of all those agents who accept formula φ in profile $\boldsymbol{J} = (J_1, \ldots, J_n)$:

$$N_\varphi^{\boldsymbol{J}} \quad := \quad \{i \in N \mid \varphi \in J_i\}$$

The Majority Rule

The (strict) *majority rule* F_{maj} takes a (complete and consistent) profile of judgment sets as input and returns the set of those propositions that are accepted by more than half of the agents:

$$\begin{aligned} F_{\mathrm{maj}} &: \mathcal{J}(\Phi)^n \rightarrow 2^{\Phi} \\ F_{\mathrm{maj}} &: \boldsymbol{J} \mapsto \{\varphi \in \Phi \mid \# N_{\varphi}^{\boldsymbol{J}} > \tfrac{n}{2}\} \end{aligned}$$

Example: Majority Rule

Suppose three agents ($N = \{1, 2, 3\}$) express judgments on the propositions in the agenda $\Phi = \{p, \neg p, q, \neg q, p \vee q, \neg(p \vee q)\}$.

For simplicity, we only show the positive formulas in our tables:

	p	q	$p \vee q$	formal notation
Agent 1	Yes	No	Yes	$J_1 = \{p, \neg q, p \vee q\}$
Agent 2	Yes	Yes	Yes	$J_2 = \{p, q, p \vee q\}$
Agent 3	No	No	No	$J_3 = \{\neg p, \neg q, \neg(p \vee q)\}$

For this example: $F_{\text{maj}}(\boldsymbol{J}) = \{p, \neg q, p \vee q\}$ [complete and consistent!]

Recall: F_{maj} does *not* guarantee *consistent* outcomes for all agendas.

Exercise: *Show that F_{maj} guarantees complement-free outcomes.*

Exercise: *Show that F_{maj} guarantees complete outcomes iff n is odd.*

Quota Rules

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

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

A quota rule F_q is called *uniform* if the function q maps any given formula to the same number λ. Examples:

- The *unanimous rule* $F_n : \boldsymbol{J} \mapsto J_1 \cap \cdots \cap J_n$ accepts φ *iff* all do.
- The *constant rule* F_0 (F_{n+1}) accepts all (no) formulas.
- The *(strict) majority rule* F_{maj} is the quota rule with $\lambda = \lceil \frac{n+1}{2} \rceil$.
- The *weak majority rule* is the quota rule with $\lambda = \lceil \frac{n}{2} \rceil$.

Observe that for *odd* n the majority rule and the weak majority rule coincide. For *even* n they differ (and only the weak one is complete).

Quota Setting and Consistency

Intuition: high quotas good for consistency (but bad for completeness)

Recall that the *unanimous rule* is the uniform quota rule with $\lambda = n$.

Exercise: *Show that the unanimous rule guarantees consistency!*

Recall the *doctrinal paradox agenda* of $\{p, \neg p, q, \neg q, p \wedge q, \neg(p \wedge q)\}$.

Exercise: *For the doctrinal paradox agenda and n agents, what is the lowest uniform quota λ that still guarantees consistency?*

Premise-Based Aggregation

Suppose we can divide the agenda into *premises* and *conclusions:*

$$\Phi \quad = \quad \Phi_p \uplus \Phi_c \quad \text{(each closed under complementation)}$$

Then the *premise-based rule* F_{pre} for Φ_p and Φ_c is this function:

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

A common assumption is that *premises* = *literals*.

Exercise: *Show that this assumption guarantees consistent outcomes.*

Exercise: *Does it also guarantee completeness? What details matter?*

Remark: The *conclusion-based rule* is less attractive from a theoretical standpoint (as it is incomplete by design), but often used in practice.

Example: Premise-Based Aggregation

Suppose *premises* = *literals*. Consider this example:

	p	q	r	$p \vee q \vee r$
Agent 1	Yes	No	No	Yes
Agent 2	No	Yes	No	Yes
Agent 3	No	No	Yes	Yes
F_{pre}	No	No	No	No

So the *unanimously accepted* conclusion is *collectively rejected!*

Discussion: *Is this ok?*

Axiomatic Method

So how do you choose the right aggregation rule?

One way is to use the *axiomatic method:*

- identify normatively appealing properties of rules
- cast those properties into mathematically rigorous definitions
- explore the consequences: *characterisations* and *impossibilities*

Any such intuitively appealing and mathematically defined property is called an *axiom*. Note the difference to how the same term is used in mathematical logic: "*obviously desirable*" *vs.* "*obviously true*".

Basic Axioms

What makes for a "good" aggregation rule F? The following *axioms* all express intuitively appealing (but always debatable!) properties:

- *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 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 $\varphi \in F(\boldsymbol{J}) \Leftrightarrow \varphi \in F(\boldsymbol{J'})$.

Observe that the *majority rule* satisfies all of these axioms.

Exercise: *But so do some other rules! Can you think of examples?*

A Subtlety of Terminology

We had defined completeness, complement-freeness, and consistency as properties of *judgment sets*, but then we are also talked about *aggregation rules* having these properties. Formally:

- F is *complete* if $F(\boldsymbol{J})$ is complete for all $\boldsymbol{J} \in \mathcal{J}(\Phi)^n$
- F is *complement-free* if $F(\boldsymbol{J})$ is compl.-free for all $\boldsymbol{J} \in \mathcal{J}(\Phi)^n$
- F is *consistent* if $F(\boldsymbol{J})$ is consistent for all $\boldsymbol{J} \in \mathcal{J}(\Phi)^n$

Remark: Sometimes these properties of aggregation rules are also referred to as "axioms". In some technical sense they indeed are axioms, but I prefer to call them *collective rationality* requirements.

A Basic Impossibility Theorem

We saw that the majority rule is not consistent. Is there maybe some other "reasonable" aggregation rule that does not have this problem? *Surprisingly, no!* (at least not for certain agendas)

This is the main result in the original paper introducing the formal model of JA and proposing to apply the axiomatic method:

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

Remark: Similar impossibilities arise for other agendas with some minimal structural richness. (To be discussed later on in the course.)

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

Proof: Part 1

Recall: $N_\varphi^{\boldsymbol{J}}$ is the set of agents who accept formula φ in profile $\boldsymbol{J}$.

Let F be any aggregator that is independent, anonymous, and neutral.

We observe:

- Due to *independence*, whether $\varphi \in F(\boldsymbol{J})$ only depends on $N_\varphi^{\boldsymbol{J}}$.
- Then, due to *anonymity*, whether $\varphi \in F(\boldsymbol{J})$ only depends on $|N_\varphi^{\boldsymbol{J}}|$.
- Finally, due to *neutrality*, the manner in which the status of $\varphi \in F(\boldsymbol{J})$ depends on $|N_\varphi^{\boldsymbol{J}}|$ must itself *not* depend on φ.

Thus: If φ and ψ are accepted by the same number of agents, then we must either accept both of them or reject both of them.

Proof: Part 2

Recall: For all $\varphi, \psi \in \Phi$, if $|N^{\boldsymbol{J}}_\varphi| = |N^{\boldsymbol{J}}_\psi|$, then $\varphi \in F(\boldsymbol{J}) \Leftrightarrow \psi \in F(\boldsymbol{J})$.

First, suppose the number n of agents is *odd* (and $n > 1$):

Consider a profile $\boldsymbol{J}$ where $\frac{n-1}{2}$ agents accept p and q; one accepts p but not q; one accepts q but not p; and $\frac{n-3}{2}$ accept neither p nor q. That is: $|N^{\boldsymbol{J}}_p| = |N^{\boldsymbol{J}}_q| = |N^{\boldsymbol{J}}_{\neg(p \wedge q)}|$. Then:

- Accepting all three formulas contradicts consistency. ✓
- But if we accept none, completeness forces us to accept their complements, which also contradicts consistency. ✓

If n is *even*, we can get our impossibility even without having to make (almost) any assumptions regarding the structure of the agenda:

Consider a profile $\boldsymbol{J}$ with $|N^{\boldsymbol{J}}_p| = |N^{\boldsymbol{J}}_{\neg p}|$. Then:

- Accepting both contradicts consistency. ✓
- Accepting neither contradicts completeness. ✓

Note: Neutrality only has "bite" here because we also have $q \in \Phi$.

Summary

This has been an introduction to the field of *judgment aggregation:*

- discussion of how to map application scenarios into a *formal model*
- examples for *rules:* quota rules, premise-based rule (more to come)
- example for *axioms:* anonymity, neutrality, independence (*ditto*)
- examples for a result: basic *impossibility* theorem

What next? We'll delve deeper into the axiomatic method.

Homework

Most exercises will be of the problem-solving sort, requiring:

- a good *understanding* of the topic to see what the question is
- some *creativity* to find the solution
- *mathematical maturity*, to write it up correctly, often as a proof
- good *taste*, to write it up in a reader-friendly manner

Solutions must be *typed* up professionally (LaTeX strongly preferred).

Of course, solutions should be *correct*. But just as importantly, they should be *short* and *easy to understand*. (This is the advanced level: it's not anymore just about you getting it, it's now about your reader!) Good solutions will typically have *around 1 page* of text per exercise.

The usual rules on permissible *collaboration* apply: discussing with others to improve your *understanding* is fine (indeed, it is encouraged), but producing your *solutions* is something you do by yourself.

Mini-Projects

During the second part of the course you'll work on your mini-project in a team of *three students*. Possible types of projects include:

- identify an interesting relevant paper not covered in class and fill in some gaps, or come up with an extension or a generalisation
- apply an algorithmic technique to a problem that to date has only been considered by economists/political scientists/philosophers
- explore an application domain: could be a literature review, an idea for a new application, or an experimental study

All projects must be related to *judgment aggregation* and make use of some of the *COMSOC techniques* introduced during the course.

Deliverables: *talk* (exam week) + *paper* (due end of block)

Activities: sessions on *how to write a paper* and *how to give a talk*, various *progress meetings* (plenary and individual), paper *reviewing*