Computational Social Choice: Spring 2015

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

Opening Example

Suppose three robots are in charge of climate control for this building. They need to make judgments on p (the temperature is below 17°C), q (we should switch on the heating), and $p \rightarrow q$.

	p	$p \rightarrow q$	q
Robot 1:	Yes	Yes	Yes
Robot 2:	No	Yes	No
Robot 3:	Yes	No	No

- What should be the collective decision?

Plan for Today

This will be an introduction to *basic judgment aggregation* (JA), starting from an example first discussed in legal theory and culminating in the first (very simple) impossibility theorem for JA:

- more on this kind of example: the *doctrinal paradox*
- general *formal framework* for judgment aggregation
- a couple of *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.

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.

The Doctrinal Paradox

Suppose 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 it 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
Majority:	Yes	Yes	No

Paradox: Taking majority decisions on the *premises* (p and q) and then inferring the conclusion (r) yields a different result from taking a majority decision on the *conclusion* (r) directly.

L.A. Kornhauser and L.G. Sager. The One and the Many: Adjudication in Collegial Courts. *California Law Review*, 81(1):1–59, 1993.

Variants

Our judges were expressing judgments on *atoms* (p, q, r) and consistency of a judgment set was evaluated w.r.t. an *integrity constraint* ($r \leftrightarrow p \wedge q$).

Alternatively, we could allow judgments on *compound formulas*, like so:

	p	q	$p \wedge q$
Judge 1:	Yes	Yes	Yes
Judge 2:	No	Yes	No
Judge 3:	Yes	No	No
Majority:	Yes	Yes	No

	p	q	$r \leftrightarrow p \wedge q$	r
Judge 1:	Yes	Yes	Yes	Yes
Judge 2:	No	Yes	Yes	No
Judge 3:	Yes	No	Yes	No
Majority:	Yes	Yes	Yes	No

Thus, we can also work within a framework without integrity constraints ("legal doctrines"), where all inter-relations between propositions stem from the logical structure of those propositions themselves.

And we do not need to distinguish premises from conclusions either.

Why Paradox?

Again, what's paradoxical about our example?

	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 procedure* and the *conclusion-based procedure*, produce *different* outcomes.

Explanation 2: Each individual judgment set is logically consistent, but applying the natural *majority rule* to all propositions produces a collective judgment set that is *inconsistent* (majority rule does not "lift" consistency).

In the philosophical literature, the term *discursive dilemma* is used for the dilemma of choosing between *responsiveness* to the views of decision makers (by respecting majority decisions) and the *consistency* of collective decisions.

Formal Framework

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

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

Remark

Later on in the course we will also see a second formal framework (called *binary aggregation with integrity constraints*) that formalises the scenario of the original doctrinal paradox more directly:

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

The two frameworks are (in some sense) equivalent. Some phenomena are better studied in one, some better in the other framework.
In principle, everything could be done in either one of the frameworks (but so far not all details have been worked out in the literature).

Useful Notation

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

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

The Majority Rule

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

$$F_{\mathsf{maj}} \;:\; \mathcal{J}(\Phi)^n \;\to\; 2^{\Phi}$$

$$F_{\mathsf{maj}} \;:\; \boldsymbol{J} \;\mapsto\; \{\varphi \in \Phi \mid |N_{\varphi}^{\boldsymbol{J}}| > \tfrac{n}{2}\}$$

Example: Majority Rule

Suppose three agents ($\mathcal{N} = \{1, 2, 3\}$) 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	$J_3 = \{\neg p, \neg q, \neg(p \vee q)\}$

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

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

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

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

Premise-Based Procedures

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 procedure F_{pre}* for Φ_p and Φ_c is this function:

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

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

Discussion: Distinction between premises and conclusions meaningful for many concrete application scenarios, but less attractive in theory.

Remark: The *conclusion-based procedure* is even less attractive from a theoretical standpoint (as it is incomplete by design).

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

Thus: the *unanimously accepted* conclusion is *collectively rejected*.

Discussion: Is this ok?

Quota Rules

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

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

A quota rule F_q is called *uniform* if 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 $q = \lceil \frac{n+1}{2} \rceil$.
- The *weak majority rule* is the quota rule with $q = \lceil \frac{n}{2} \rceil$.

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).

Axiomatic Method

So how do you choose the right aggregation rule?

One way is to use the *axiomatic method*, as in economic theory:

- identify normatively appealing properties of aggregators
- 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: here, axioms need not always be satisfied.

Basic Axioms

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

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

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

(But so do some other rules! Can you think of some examples?)

A Subtlety of Terminology

Recall the definitions of completeness, complement-freeness and consistency, which are properties *of judgment sets*.

They give rise to three *axioms*, i.e., properties *of aggregation rules* F:

- 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: Whether to call these three requirements "axioms" actually is a matter of taste. Later on in the course we will prefer to call them *collective rationality* requirements.

Impossibility Theorem

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

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.*

Remark 1: Note that the theorem requires $n \geqslant 2$.

Remark 2: Similar impossibilities arise for other agendas with some minimal structural richness. (To be discussed later on in the course.)

Remark 3: This is the main result in the original paper introducing the formal framework of JA and proposing to apply the axiomatic method.

C. List and P. Pettit. Aggregating Sets of Judgments: An Impossibility Result. *Economics and Philosophy*, 18(1):89–110, 2002.

Proof: Part 1

Recall: $N_\varphi^{\boldsymbol{J}}$ is the set of individuals 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 individuals, then we must either accept both of them or reject both of them.

Proof: Part 2

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

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

Consider a profile $\boldsymbol{J}$ where $\frac{n-1}{2}$ individuals 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: $|N_p^{\boldsymbol{J}}| = |N_q^{\boldsymbol{J}}| = |N_{\neg(p \wedge q)}^{\boldsymbol{J}}|$. 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 $|N_p^{\boldsymbol{J}}| = |N_{\neg p}^{\boldsymbol{J}}|$. Then:

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

Remark: To be exact, you also need, say, $q \in \Phi$ for neutrality to bite.

Homework

The first homework set will come out this afternoon.

The usual *rules* apply:

- Cooperation to gain understanding is encouraged.
- The solutions you submit must be your own.

Remember what the *objectives* are:

- Partly this is about digesting the material taught and solving sometimes challenging problems.
- Partly this is about learning how to write good science: succinct, precise, readable, enlightening, elegant, ...

Spend roughly equal amounts of time on each of these two objectives.

Good solutions for Homework #1 should fit on one page.

What next?

Plan for (roughly) the next two lectures:

- Is the impossibility theorem the end of it?

  No: discussion of various ways of *circumventing the impossibility*

- More on the axiomatic method: *characterisation results*

- More *specific aggregation rules* to use in practice

And later on in the course:

- Refinement of the impossibility theorem: for *which agendas*, beyond those including $\{p, q, p \wedge q\}$, do we get this problem?