Computational Social Choice: Autumn 2013

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

Plan for Today

Preferences are not the only thing we may wish to aggregate.

Today's lecture will be an introduction to *judgment aggregation*, a framework where the views to be aggregated concern the truth or falsehood of formulas expressed in propositional logic. We will cover:

- Motivation: *doctrinal paradox* and *discursive dilemma*
- Definition of the *formal framework* and of basic *axioms*
- Embedding of *preference aggregation* into JA
- Basic impossibility result: *List-Pettit Theorem*
- Discussion of a few specific *aggregation procedures*

C. List. The Theory of Judgment Aggregation: An Introductory Review. *Synthese*, 187(1):179–207, 2012.

U. Endriss. Logic and Social Choice Theory. In A. Gupta and J. van Benthem (eds.), *Logic and Philosophy Today*, College Publications, 2011.

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.

The Discursive Dilemma

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

Alternatively, we could allow judgements on *compound formulas*. Examples:

	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

From now on we will work with a framework without integrity constraints ("legal doctrines"), where all inter-relations between propositions stem from the logical structure of those propositions themselves.

In the philosophical literature, the term *doctrinal paradox* is reserved for the first version of our paradox, and the more general term *discursive dilemma* is used when there is no external "doctrine" that is responsible for the problem.

Why Paradox?

Again, what's paradoxical about our example?

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

Explanation 1: Two seemingly reasonable methods of aggregation, the *premise-based procedure* and the *conclusion-based procedure*, produce different outcomes.

Explanation 2: Each individual judgment set is logically consistent, but applying the seemingly reasonable *majority rule* to all propositions yields a collective judgment set that is inconsistent. The majority rule *cannot lift consistency* from the individual to the collective level.

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 procedure* 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: $F : \mathcal{J}(\Phi)^n \to 2^\Phi$.

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	$p \vee q$	
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	$J_3 = \{\neg p, \neg q, \neg(p \vee q)\}$

The (strict) *majority rule* F_{maj} takes a (complete and consistent) profile and returns the set of propositions accepted by $> \frac{n}{2}$ agents.

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

In general, F_{maj} only ensures completeness and complement-freeness [and completeness only in case n is odd].

Example: Embedding Preference Aggregation

In *preference aggregation*, individuals express preferences (linear orders) over a set of alternatives $\mathcal{X}$ and we need to find a collective preference.

We can embed this into JA (suppose $\mathcal{X} = \{A, B, C\}$):

- Take atomic propositions to be $[A \succ A]$, $[A \succ B]$, ...
- Suppose all individuals accept these propositions:
  - Irreflexivity: $\neg[A \succ A]$, $\neg[B \succ B]$, $\neg[C \succ C]$
  - Completeness: $[A \succ B] \vee [B \succ A]$ etc.
  - Transitivity: $[A \succ B] \wedge [B \succ C] \rightarrow [A \succ C]$, etc.

Then the *Condorcet paradox* corresponds to this example in JA:

	$[A \succ B]$	$[A \succ C]$	$[B \succ C]$	*corresponding order*
Agent 1:	Yes	Yes	Yes	$A \succ B \succ C$
Agent 2:	No	No	Yes	$B \succ C \succ A$
Agent 3:	Yes	No	No	$C \succ A \succ B$
Majority:	Yes	No	Yes	*not a linear order*

Axioms

What makes for a "good" aggregation procedure F? The following *axioms* all express intuitively appealing (yet, 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 procedures! Can you think of some examples?)

Impossibility Theorem

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

Theorem 1 (List and Pettit, 2002) *No judgment aggregation procedure for an agenda Φ with $\{p, q, p \wedge q\} \subseteq \Phi$ that satisfies the axioms of anonymity, neutrality, and independence will always return a collective judgment set that is complete and consistent.*

Remark 1: Note that the theorem requires $|\mathcal{N}| > 1$.

Remark 2: Similar impossibilities arise for other agendas with some minimal structural richness. To be discussed tomorrow.

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

Proof: Part 1

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

<u>Thus:</u> 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 each accept exactly one of p and q; and $\frac{n-3}{2}$ accept neither p nor q.
That is: $|N_p^{\boldsymbol{J}}| = |N_q^{\boldsymbol{J}}| = |N_{\neg(p \wedge q)}^{\boldsymbol{J}}| = \frac{n+1}{2}$. 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 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. ✓

Change of Perspective

Does the impossibility theorem mean that all hope is lost? *No*.

- We could analyse in more depth for *what agendas* this problem can actually occur. And if it can, we could analyse *how to detect* such a situation. We will follow this route tomorrow.
- We could argue that it is ok to *weaken those axioms:*
  - *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?

Next we look at some practical aggregators that circumvent the noted impossibility (i.e., they all must violate at least one of the axioms).

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 k. Examples:

- The *unanimous rule* F_n accepts φ *iff* everyone does.
- 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).

Quota Rules with a High Quota

Clearly, a (uniform) quota rule with a sufficiently high quota will be consistent. Dietrich and List (2007) give lower bounds for the quota to ensure consistency as a function of n and the size of the largest *minimally inconsistent subset* of the agenda Φ. Example:

> Let $\Phi = \{p, \neg p, q, \neg q, p \wedge q, \neg(p \wedge q)\}$. The largest mi-subset is $\{p, q, \neg(p \wedge q)\}$. Any quota $> \frac{2}{3}$ will ensure consistency.

But: We (may) lose completeness of the collective judgment set.

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

Characterisation of Quota Rules

Quota rules are nice to demonstrate the axiomatic method . . .

One more axiom:

- *Monotonicity:* If an accepted proposition gets additional support, then we should continue to accept it!

  Formally: for any $\varphi \in \Phi$, profile $\boldsymbol{J}$, agent i, and judgment set J_i', $\varphi \in J_i' \setminus J_i$ entails $\varphi \in F(\boldsymbol{J}) \Rightarrow \varphi \in F(\boldsymbol{J}_{-i}, J_i')$.

We can now *characterise* the class of quota rules:

Proposition 1 (Dietrich and List, 2007) *An aggregation procedure is anonymous, independent and monotonic iff it is a quota rule.*

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

Proof

Claim: *anonymous* + *independent* + *monotonic* $\Leftrightarrow$ *quota rule*

Clearly, any quota rule has these properties (right-to-left). ✓

For the other direction (using the same technique as before):

- Independence means that acceptance of φ only depends on $N_\varphi^{\boldsymbol{J}}$.
- Anonymity means that, in fact, it only depends on $|N_\varphi^{\boldsymbol{J}}|$.
- Monotonicity means that acceptance of φ cannot turn to rejection as additional individuals accept φ.

Hence, it must be a quota rule. ✓

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

More Characterisations

Clearly, a quota rule F_q is uniform *iff* it is neutral. Thus:

Corollary 1 *An aggregation procedure is anonymous, neutral, independent and monotonic (= ANIM) iff it is a uniform quota rule.*

Now consider a uniform quota rule F_q with quota q. Two observations:

- For F_q to be *complete*, we need $q \leqslant \max\limits_{0 \leqslant x \leqslant n}(x, n-x) \Rightarrow q \leqslant \lceil \frac{n}{2} \rceil$.
- For F_q to be *compl.-free*, we need $q > \min\limits_{0 \leqslant x \leqslant n}(x, n-x) \Rightarrow q > \lfloor \frac{n}{2} \rfloor$.

For n *even*, no such q exists. Thus:

Proposition 2 *For n even, no aggregation procedure is ANIM, complete and complement-free.*

For n *odd*, such a q does exist, namely $q = \lceil \frac{n}{2} \rceil = \frac{n+1}{2}$. Thus:

Proposition 3 *For n odd, an aggregation procedure is ANIM, complete and complement-free iff it is the (strict) majority rule.*

The Premise-Based Procedure

Suppose we *can* divide the agenda into *premises* and *conclusions* (i.e., we are willing to give up *neutrality*):

$$\Phi \quad = \quad \Phi_p \uplus \Phi_c$$

The *premise-based procedure* PBP for Φ_p and Φ_c is this function:

$$\begin{aligned}\mathrm{PBP}(\boldsymbol{J}) \quad &= \quad \Delta \cup \{\varphi \in \Phi_c \mid \Delta \models \varphi\}, \\ & \quad\quad \text{where } \Delta = \{\varphi \in \Phi_p \mid |\{i \mid \varphi \in J_i\}| > \frac{n}{2}\}\end{aligned}$$

If we assume that

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

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

Example: Violation of Propositionwise Unanimity

Consider the following basic axiom:

- *Propositionwise Unanimity*: $\varphi \in J_i$ for all $i \in \mathcal{N} \Rightarrow \varphi \in F(\boldsymbol{J})$.

  Unanimous acceptance implies collective acceptance!

Curiously, the premise-based procedure violates this form of unanimity:

	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
PBP:	No	No	No	No

Discussion: Maybe this is ok?

Distance-based Aggregation

The standard *distance-based procedure* is defined as follows:

$$\mathrm{DBP}(\boldsymbol{J}) \quad = \quad \underset{J \in \mathcal{J}(\Phi)}{\operatorname{argmin}} \sum_{i=1}^{n} |(J \setminus J_i) \cup (J_i \setminus J)|$$

<u>That is:</u> find a *complete and consistent* judgment set that minimses the sum of the *Hamming distances* to the individual judgment sets.

- *irresolute* aggregation procedure
- generalises the idea underlying the *Kemeny* rule in voting
- conincides wih the *majority outcome* whenever that is consistent
- other options: *Slater*, *Tideman*

M.K. Miller and D. Osherson. Methods for Distance-based Judgment Aggregation. *Social Choice and Welfare*, 32(4):575–601, 2009.

Complexity of Winner Determination

How hard is it to compute the collective judgment set for the aggregators we have seen? (This is the *winner determination problem*.)

Fact 1 *Winner determination for any quota rule F_q is polynomial.*

Proposition 4 *Winner determination for the PBP is polynomial.*

Proof: counting (for premises) $+$ model checking (for conclusions) ✓

Theorem 2 *Winner determination for the Kemeny-DBP is Θ_2^p-compl.*

Proof: Omitted. [Θ_2^p is also known as “parallel access to NP”]

U. Endriss, U. Grandi, and D. Porello. Complexity of Judgment Aggregation. *Journal of Artificial Intelligence Research*, 45:481–514, 2012.

Representative-Voter Rules

The complexity of the DBP stems from the fact that we have to search through *all consistent judgment sets* to find the one that's closest to the profile. If we restrict this set, we can do better.

Idea: Only search through the *support*, i.e., judgment sets proposed by individuals. That is, identify "*the most representative voter*".

One possible implementation of this idea is the *average-voter rule:*

$$\mathrm{AVR}(\boldsymbol{J}) \quad = \quad \underset{J \in \textsc{Supp}(\boldsymbol{J})}{\operatorname{argmin}} \sum_{i=1}^{n} |(J \setminus J_i) \cup (J_i \setminus J)|$$

$$\text{where } \textsc{Supp}(\boldsymbol{J}) = \{J_1, J_2, \ldots, J_n\}$$

Fact 2 *Winner determination for the* AVR *is polynomial.*

U. Endriss and U. Grandi. Binary Aggregation by Selection of the Most Representative Voter. Proc. MPREF-2013.

Summary

This has been an introduction to judgment aggregation. Main topics:

- *axioms:* independence, neutrality, monotonicity, ...
- *List-Pettit Theorem:* no consistent aggregator is independent, neutral, and anonymous for the 'conjunctive agenda'
- *specific rules:* quota-based, premise-based, distance-based

What next?

In the next lecture, we will cover more advanced topics in judgment aggregation, particularly *agenda characterisation* results.