

Computational Social Choice: Autumn 2012

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

The Doctrinal Paradox

Consider a court with three judges. Suppose legal doctrine stipulates that the defendant is *liable* (r) iff there has been a valid *contract* (p) and that contract 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 issue-by-issue, here p and q , (and deciding on the case r accordingly) gives a different result from taking majority decisions case-by-case (that is, on 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.

Plan for Today

Preferences are not the only type of information we may wish to aggregate. *Judgment aggregation* deals with the aggregation of truth assignments to logically interrelated propositions.

Topics to be covered today:

- The Doctrinal Paradox (paradox of judgment aggregation)
- Formal framework, possible aggregation procedures, axioms
- An impossibility theorem
- Ways around the impossibility
- Links between preference aggregation and judgment aggregation

List (2011) covers these topics in detail. For a short introduction, refer to Section 5 of *Logic and Social Choice Theory*.

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.

Variants of the Paradox

In the example, individuals were expressing judgements on atomic propositions (p, q, r) and consistency of a judgement set was evaluated wrt. a background theory ($r \leftrightarrow p \wedge q$).

Alternatively, we could allow judgements directly on compound formulas. Or we could make the legal doctrine itself a proposition on which individuals can express a judgement.

	p	q	$p \wedge q$	r
Judge 1:	yes	yes	yes	yes
Judge 2:	no	yes	no	no
Judge 3:	yes	no	no	no
Majority:	yes	yes	no	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

Conclusion: We do not require the notion of a background theory (doctrine) to model the full extent of the problem.

Remark

Observe that the doctrinal paradox (all of its variants) satisfies the general definition of “paradox” given in the lecture of voting in combinatorial domains:

- In the original version, each individual judgment set “satisfies” the integrity constraint given by the legal doctrine, while the collective judgment set returned by the majority rule does not.
- In both of the variants, each individual judgment set satisfies the integrity constraint of logical consistency, while the collective judgment set returned by the majority rule does not.

U. Grandi and U. Endriss. Binary Aggregation with Integrity Constraints. Proc. IJCAI-2011.

Aggregation Procedures

Some of the aggregation procedures that have been considered in the literature (we shall review some of them in more depth later on):

- **Majority rule:** accept φ if a strict majority does
- **Quota rules:** accept φ if at least $q\%$ do
- **Premise-based procedure:** use majority rule on “premises” and logically infer status of conclusions
- **Conclusion-based procedure:** use majority rule on “conclusions” (and ignore premises)
- **Distance-based procedures:** choose a judgment set that minimises a suitable notion of distance from the profile

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

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

Axioms

What makes for a “good” aggregation procedure F ? The following axioms all express intuitively appealing properties:

- **Unanimity:** if $\varphi \in J_i$ for all i , then $\varphi \in F(\mathbf{J})$.
- **Anonymity:** for any profile \mathbf{J} and any permutation $\pi : \mathcal{N} \rightarrow \mathcal{N}$ we have $F(J_1, \dots, J_n) = F(J_{\pi(1)}, \dots, J_{\pi(n)})$.
- **Neutrality:** for any φ, ψ in the agenda Φ and profile $\mathbf{J} \in \mathcal{J}(\Phi)^{\mathcal{N}}$, if for all i we have $\varphi \in J_i \Leftrightarrow \psi \in J_i$, then $\varphi \in F(\mathbf{J}) \Leftrightarrow \psi \in F(\mathbf{J})$.
- **Independence:** for any $\varphi \in \Phi$ and profiles \mathbf{J} and \mathbf{J}' in $\mathcal{J}(\Phi)^{\mathcal{N}}$, if $\varphi \in J_i \Leftrightarrow \varphi \in J'_i$ for all i , then $\varphi \in F(\mathbf{J}) \Leftrightarrow \varphi \in F(\mathbf{J}')$.
- **Systematicity** = neutrality + independence

(Note that the majority rule satisfies all of these axioms.)

Impossibility Theorem

We have seen that the majority rule can produce inconsistent outcomes. Is there another aggregation procedure that satisfies our axioms but that does not have this problem? No! (at least not if our agenda satisfies some minimal structural richness condition)

Theorem 1 (List and Pettit, 2002) *No judgment aggregation procedure for an agenda Φ with $\{p, q, p \wedge q\} \subseteq \Phi$ that satisfies 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 in more later on.

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 the axioms, we may be able to circumvent the impossibility and successfully aggregate judgements. Next, we will explore some such possibilities:

- Relaxing the *input* conditions: drop the (implicit) universal domain assumption and design rules for restricted domains
- Relaxing the *output* conditions: drop the completeness requirement (dropping consistency works but is unattractive)
- Giving up *anonymity*: dictatorships will surely work, but maybe we can do a little better than that
- Weakening *systematicity*: maybe neutrality is after all rather inappropriate for logically interconnected propositions (?), and we already know that independence is a very demanding axiom

Proof

From anonymity, neutrality and independence: collective acceptance of φ can only depend on the *number* $\#[\varphi]$ of individuals accepting φ .

- Case where the number n of individuals is *even*:
Consider a scenario where $\#[p] = \#[\neg p]$.
As argued above, we need to accept either both or neither:
 - Accepting both contradicts consistency. ✓
 - Accepting neither contradicts completeness. ✓
- Case where the number n of individuals is *odd* (and $n > 1$):
Consider a scenario where $\frac{n-1}{2}$ accept p and q ; 1 each accept exactly one of p and q ; and $\frac{n-3}{2}$ accept neither p nor q .
That is: $\#[p] = \#[q] = \#[\neg(p \wedge q)]$. But:
 - Accepting all three formulas contradicts consistency. ✓
 - But if we accept none, completeness forces us to accept their complements, which also contradicts consistency. ✓

Unidimensional Alignment

Call a profile of individual judgement sets *unidimensionally aligned* if we can order the individuals such that for each proposition φ in the agenda the individuals 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)
$p \rightarrow q$	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 1 (List, 2003) *For any unidimensionally aligned profile, the majority rule will return a collective judgment set that is consistent.*

C. List. A Possibility Theorem on Aggregation over Multiple Interconnected Propositions. *Mathematical Social Sciences*, 45(1):1–13, 2003.

Proof

For simplicity, 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)
$p \rightarrow q$	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 judgement set of the median individual is consistent, so is the collective judgement set under the majority rule. ✓

Oligarchic Rules

As we have seen, supermajority rules (with suitable quota) can circumvent the impossibility, if we are prepared to give up completeness.

Instead, we may try replacing completeness by *deductive closure*:
 $\varphi \in \Phi$ and $J \models \varphi$ imply $\varphi \in \Phi$ for the (collective) judgement set J

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

- dictatorial rule: $|X| = 1$
- unanimous rule: $|X| = n$

It is easy to check that any oligarchic rule satisfies:

- *consistency* and *deductive closure* (if individuals do);
- the *universal domain* assumption, *neutrality*, and *independence*;
- but *not anonymity* (except if $|X| = n$).

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.

Supermajority Rules

Another option is to *drop completeness* from our list of requirements. If the collective judgement set need not be complete, then we can get aggregation procedures satisfying the remaining axioms:

- *Unanimous rule*: include φ in the collective judgement set iff φ is in every individual judgement set. Always works.
- Consider this variant of the doctrinal paradox:

	p	q	$r \leftrightarrow p \wedge q$	r
Judges 1–10:	yes	yes	yes	yes
Judges 11–20:	no	yes	yes	no
Judges 21–30:	yes	no	yes	no

Here the *4/5-supermajority rule*, accepting φ iff > 24 judges do, produces a consistent (but not necessarily complete) outcome. For general results of this sort, see Dietrich and List (2007).

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

The Premise-Based Procedure

Another option is to sacrifice *neutrality* ...

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

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

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

$$\text{PBP}(\mathbf{J}) = \Delta \cup \{\varphi \in \Phi_c \mid \Delta \models \varphi\},$$

where $\Delta = \{\varphi \in \Phi_p \mid \#\{i \mid \varphi \in J_i\} > \frac{n}{2}\}$

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 $\text{PBP}(\mathbf{J})$ will always be *consistent* and *complete*.

Discussion: Premise-Based Procedure

- The PBP satisfies all of our axioms *with respect to the set of premises*, but not with respect to the full agenda.
- Most notably, the PBP violates *neutrality* (premises and conclusions are treated differently). Maybe appropriate?
- Equating premises with literals seems overly simplistic. In general, we would need:
 - The premises should be (almost) logically independent from each other to avoid generating inconsistencies (we will see a precise characterisation in the next lecture).
 - For any collective judgment on the premises, the judgments on the conclusions must be logically determined.

These are *computationally demanding* properties to check.

Preference vs. Judgement Aggregation

Naturally, there are close links between PA and JA.

One can (and people do) argue over which is more general ...

For example, we can model the *Condorcet Paradox* in JA:

	$A \succ B$	$A \succ C$	$B \succ C$	
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]

And all agents agree on these propositions:

- $\neg[A \succ A], \neg[B \succ B], \neg[C \succ C]$
- $[A \succ B] \vee [B \succ A], [A \succ C] \vee [C \succ A], [B \succ C] \vee [C \succ B]$
- $[A \succ B] \wedge [B \succ C] \rightarrow [A \succ C]$, etc.

Distance-Based Procedures

Finally, we might be willing to sacrifice *independence* ...

A procedure that is more widely applicable than the premise-based procedure and that is intuitively appealing is *distance-based merging*:

$$\text{DBP}(\mathcal{J}) = \arg \min_{\mathcal{J} \in \mathcal{J}(\Phi)} \sum_{i=1}^n H(\mathcal{J}, \mathcal{J}_i)$$

Here the *Hamming distance* $H(\mathcal{J}, \mathcal{J}')$ between judgment sets \mathcal{J} and \mathcal{J}' is the number of positive agenda formulas on which they differ.

Remark: The DBP may return a set of tied winners.

The DBP behaves like the majority rule in case that is consistent, and makes a “reasonable” (consistent) choice otherwise. Variants are possible.

G. Pigozzi. Belief Merging and the Discursive Dilemma: An Argument-based Account of Paradoxes of Judgment. *Synthese*, 152(2):285–298, 2006.

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

Summary

This has been a brief introduction to judgment aggregation.

Each individual selects a (consistent) set of propositional formulas \Rightarrow how do we aggregate these choices into a consistent collective choice?

- Impossibility: anonymity, neutrality, independence \Rightarrow inconsistency
- Possibilities: domain restrictions or sacrificing one of completeness, anonymity, neutrality, and independence
- Relation to preference aggregation

What next?

Next we will discuss a number of advanced topics in judgment aggregation, particularly:

- deeper analysis of *impossibilities* and *characterisation* of agendas that allow for impossibilities
- a few *complexity* results