Computational Social Choice: Spring 2009

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

Plan for Today

Much of social choice theory is about the problems associated with *aggregating preferences* (linear orders, utility functions, ...).

Today we will look into the problem of *aggregating judgements:* truth assignments to logically interconnected propositions.

- *Doctrinal Paradox:* a first example demonstrating that JA is difficult (inspired by work in Law and Economics)
- *Impossibility Result:* a set of reasonable axioms and a theorem showing that there can be no JA procedure satisfying them all
- Conditions under which we can *circumvent impossibilities*
- *Procedures* for JA, each satisfying a subset of the axioms

Much of this lecture is based on the tutorial paper by List (2008).

Ch. List. *Judgment Aggregation: A Short Introduction.* Manuscript, London School of Economics, 2008.

The Doctrinal Paradox

Take a court with three judges. Suppose legal doctrine stipulates that the defendant is *liable* (C) iff there has been a valid *contract* (A) and that contract has been *breached* (B): $C \leftrightarrow A \wedge B$.

	A	B	C
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 A and B, (and deciding on the case C accordingly) gives a different result from taking majority decisions case-by-case (that is, on C 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 of the Paradox

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

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

	A	B	$A \wedge B$
Judge 1:	yes	yes	yes
Judge 2:	no	yes	no
Judge 3:	yes	no	no
Majority:	yes	yes	no

	A	B	$C \leftrightarrow A \wedge B$	C
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.

Judgement Aggregation: The Model

- Finite set of variables PS, propositional language $\mathcal{L}_{PS}$ over PS
- An *agenda* is a (finite) set of formulas $\Phi \subseteq \mathcal{L}_{PS}$ that is closed under complementation.
- *Judgement set:* subset $J \subseteq \Phi$ of formulas in the agenda
  - *consistent:* if $J \not\models \bot$
  - *complete:* if for each proposition $\varphi \in \Phi$, $\varphi \in J$ or $\overline{\varphi} \in J$
- Finite set of (at least two) *individuals* $I = \{1, \ldots, n\}$, each with a (usually consistent and complete) judgement set
- A *judgement aggregation rule* is a function mapping each profile of individual judgement sets to a collective judgement set.

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.

Axioms

Possible choices of axioms for judgement aggregation include:

- *Universal Domain* (UD): the rule should be defined for any profile of consistent and complete judgement sets
- *Anonymity* (AN): symmetry wrt. individuals
- *Neutrality* (NE): symmetry wrt. elements of the agenda
- *Independence* (IN): inclusion of a proposition φ (of the agenda) into the collective judgement set should depend solely on (non-)inclusion of φ in the individual judgement sets

Independence + neutrality is also known as *systematicity*.

Impossibility Theorem

The original impossibility theorem for judgement aggregation:

Theorem 1 (List and Pettit, 2002) *If the agenda contains at least P, Q, and $P \wedge Q$, then no rule producing consistent and complete judgement sets satisfies* (UD), (AN), (NE), *and* (IN).

Remark: The theorem also holds for other sufficiently complex agendas, e.g., any agenda containing at least P, Q, and $P \rightarrow Q$.

Now for the proof . . .

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

Proof

By anonymity and neutrality, collective acceptance of φ must be a function of the *number of individuals* accepting φ alone.

Write $\#[\varphi]$ for the number of individuals accepting φ.

- Suppose the number n of individuals is *even:*

  Due to the universal domain axiom, we must cater for the case where $\#[P \wedge Q] = \#[\neg(P \wedge Q)]$. As argued above, we need to accept either both or neither. Accepting both contradicts consistency. Accepting neither contradicts completeness. ✓

- Suppose the number n of individuals is *odd* (and $n > 1$):

  Suppose $\frac{n-1}{2}$ accept P and Q; 1 each accept exactly one of P and Q; and $\frac{n-3}{2}$ accept neither $\Rightarrow \#[P] = \#[Q] = \#[\neg(P \wedge Q)]$
  Accepting all three formulas contradicts consistency.
  But if we accept none, completeness forces us to accept their complements, which also contradicts consistency. ✓

Agenda Characterisation

Several variants of above impossibility theorem, for different sets of axioms, are discussed in the literature.

There are also so-called *agenda characterisation theorems*, which give sufficient and necessary conditions for an impossibility to arise.

Some suggestions for further reading are listed below.

K. Nehring and C. Puppe. Consistent Judgement Aggregation: The Truth-functional Case. *Social Choice and Welfare*, 31(1):41–57, 2008.

E. Dokow and R. Holzman. Aggregation of Binary Evaluations. *Journal of Economic Theory*. In press (2008).

E. Dokow and R. Holzman. Aggregation of Binary Evaluations for Truth-functional Agendas. *Social Choice and Welfare*, 32(2):221–241, 2009.

Ch. List and C. Puppe. *Judgement Aggregation: A Survey*. In P. Anand *et al.* (eds.), *Handbook of Rational and Social Choice*. OUP, 2009.

Circumventing the Impossibility Theorem

If we are prepared to relax some of the axioms, we may be able to circumvent the impossibility theorems and successfully aggregate judgements. Next, we will explore some such possibilities:

- Relaxing the *input* conditions: drop the universal domain axiom 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

Unidimensional Alignment

Call a profile of individual judgement sets *unidimensionally aligned* iff 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)
A	yes	yes	no	no	no	(no)
B	no	no	no	no	yes	(no)
$A \rightarrow B$	no	no	yes	yes	yes	(yes)

Theorem 2 (List, 2003) *If profiles are unidimensionally aligned, then the majority rule will produce a consistent outcome.*

Note that the other axioms are all satisfied by the majority rule also in the general case (completeness only if n is odd).

Ch. 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)
A	yes	yes	no	no	no	(no)
B	no	no	no	no	yes	(no)
$A \to B$	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. ✓

Interlude: Single-Peaked Preferences

Unidimensional alignment roughly corresponds to the case of single-peaked preferences in preference aggregation.

A profile of individual preferences over a set of alternatives A is called *single-peaked* iff there exists a "left-to-right" ordering $<$ on A such that for each individual's most preferred candidate x we have that y is preferred over z whenever $x < y < z$ or $z < y < x$.

On single-peaked domains, social choice works very well: the *Condorcet Paradox*, *Arrow's Theorem*, and the *Gibbard-Satterthwaite Theorem* all go away.

D. Black. *The Theory of Committees and Elections.* Cambridge, 1958.

Value Restriction

For simplicity, assume the agenda Φ doesn't contain contradictions.

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 of individual judgement sets *value-restricted* iff every minimally inconsistent $X \subseteq \Phi$ has a two-element subset $Y \subseteq X$ that is not a subset of any of the judgement sets.

Theorem 3 (Dietrich and List, 2007) *If profiles are value-restricted, then the majority rule will produce a consistent outcome.*

Remark: Unidimensional alignment entails value-restriction, so the former is more powerful a criterion (Dietrich and List, 2007).

F. Dietrich and Ch. List. *Majority Voting on Restricted Domains.* Working Paper, London School of Economics, 2007.

Proof

Assume the profile $\langle J_1, \ldots, J_n \rangle$ is value-restricted.

Now (for the sake of contradiction) suppose J is inconsistent. Then there exists a set $X \subseteq J$ that is minimally inconsistent.

By value restriction, there exists a set $Y = \{p, q\} \subseteq X$ such that $Y \not\subseteq J_i$ for all $i \in \{1, \ldots, n\}$.

On the other hand, due to $Y \subseteq J$, there must have been a (strict) majority for both p and q. Hence, there must exist at least one $i \in \{1, \ldots, n\}$ such that $Y \subseteq J_i \Rightarrow$ contradiction. ✓

Supermajority Rules

Or we could *drop completeness* from our list of requirements. If the collective judgement set need not be complete, we can get judgement aggregation rules 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 original doctrinal paradox:

	A	B	$C \leftrightarrow A \wedge B$	C
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 ≥ 25 judges do, produces a consistent (but not necessarily complete) outcome.
- For general results of this sort, see Dietrich and List (2007).

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

Oligarchic Rules

As we have seen, supermajority rules (with suitable quota) can circumvent 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 $X \subseteq I$ 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);
- *universal domain*, *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.

Premise-Based Procedure

For the original doctrinal paradox, the *premise-based procedure* consists in using the majority rule for A and B ("premises"), and then inferring the collective judgement on $A \wedge B$ ("conclusion").

	A	B	$A \wedge B$
Judge 1:	yes	yes	yes
Judge 2:	no	yes	no
Judge 3:	yes	no	no
Collective:	yes	yes	___

The premise-base procedure (for this agenda) satisfies consistency and completeness, but violates neutrality and independence.

General Premise-Based Procedures

How can we distinguish "premises" from "conclusions" in the general case? $\Rightarrow$ We can't. But we can do this:

(1) Label any logically independent subset Δ of the propositions in the agenda as "premises"

A set of formulas Δ is *logically independent* iff, for any $\Gamma \subseteq \Delta$, the set $\Gamma \cup \{\neg\varphi \mid \varphi \in \Delta \setminus \Gamma\}$ is consistent.

(2) Make collective judgements on each of these premises using the majority rule.

(3) Add any further propositions from the agenda that are logical consequences of these decisions to the collective judgement set.

This procedure satisfies *consistency* and *deductive closure*. If Δ is maximally logically independent, then it also satisfies *completeness*.

Logically Independent Agendas

A (very) special case is when some $\Delta \subseteq \Phi$ with $|\Delta| = \frac{1}{2} \cdot |\Phi|$ is *logically independent* (i.e., pick one from each pair of complements).

Then the *majority rule* will always produce a *consistent* outcome.

This roughly corresponds to the case of separable preferences discussed during the lecture on voting in combinatorial domains.

Distance-Based Procedures

Idea: *enforce consistency* by choosing collective judgement set *"closest"* to some ideal (possibly inconsistent) aggregated set

Assumption: For simplicity, assume the agenda Φ is such that any consistent and complete judgement set forces a *unique model* (e.g., assume Φ includes all atomic propositions).

Define a *distance-based procedure* in two steps:

- Fix a distance metric between models (and judgement sets), e.g., the *Hamming distance*
- Fix an objective function to optimise, e.g., (minimise) the *sum* of the individual distances to the collective choice

This procedure (Hamming/Σ) behaves like the majority rule in case that is consistent, and makes a "reasonable" choice otherwise.

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

Summary

This has been an introduction to judgement aggregation:

- Basic problem: each individual selects a (consistent) set of propositional formulas $\Rightarrow$ how do we aggregate these choices so as to obtain a consistent collective choice?
- *Doctrinal paradox* and an *impossibility theorem* (several further such results in the literature, some with necessary conditions)
- Aggregate anyway: *restricted domains*, *drop completeness*, *premise-based procedures*, *distance-based procedures*
- Related work: *belief merging* (see e.g. Konieczny and Pino Pérez, 2002) and *voting in combinatorial domains*

S. Konieczny and R. Pino Pérez. Merging Information under Constraints: A Logical Framework. *Journal of Logic and Computation*, 12(5):773–808, 2002.

References

A good starting point for learning about judgement aggregation is List's introductory paper (also the main reference for this lecture):

- Ch. List. *Judgment Aggregation: A Short Introduction*. Manuscript, London School of Economics, 2008.

Some additional material is covered in this survey paper:

- Ch. List and C. Puppe. *Judgement Aggregation: A Survey*. In P. Anand *et al.* (eds.), *Handbook of Rational and Social Choice*. Oxford University Press, 2009.

What next?

Next we will move on to problems related to *distributive justice*, *fair division*, and *multiagent resource allocation*.

- Rather than choosing one alternative for all individuals, now we need to *divide a common resource* and individuals have preferences over their lot (still a social choice problem!).
- Preferences will typically be modelled as *utility functions*, rather than as linear orders.

By restricting attention to more specific problems and allowing for richer preference structures, we will encounter fewer impossibilities.

Plan for the next few weeks:

- *axiomatic treatment* of different criteria for judging solutions
- procedures for different types of domains: *cake-cutting*, *distributed resource allocation*, *combinatorial auctions*