SYSU Lectures on the Theory of Aggregation
Lecture 1: Preference and Judgment Aggregation

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

[http://www.illc.uva.nl/~ulle/sysu-2014/]

Social Choice Theory

SCT studies collective decision making: how should we aggregate the preferences of the members of a group to obtain a "social preference"?

Expert 1: $\triangle \succ \bigcirc \succ \square$

Expert 2: $\bigcirc \succ \square \succ \triangle$

Expert 3: $\square \succ \triangle \succ \bigcirc$

Expert 4: $\square \succ \triangle \succ \bigcirc$

Expert 5: $\bigcirc \succ \square \succ \triangle$

?

Plan for this Lecture Series

This will be an introduction to the theory of aggregation, ranging from classical topics in social choice theory to some very recent results.

- Lecture 1: Preference and Judgment Aggregation
- Lecture 2: Binary Aggregation with Integrity Constraints
- Lecture 3: Graph Aggregation
- Lecture 4: Collective Annotation

Slides and additional reading materials are available online:

http://www.illc.uva.nl/~ulle/sysu-2014/

Plan for Today

Today's lecture will be an introduction to two standard domains: preference aggregation and judgment aggregation.

- Voting Theory: examples, rules, axioms, May's Theorem
- Preference Aggregation: Arrow's Theorem
- Remarks on the use of logic in social choice theory
- Judgment Aggregation: List-Pettit Theorem

We will revisit several of the topics touched upon today in more detail later on in the series.

Example: Electing a President

Remember Florida 2000 (simplified):

49%:	Bush $\succ$ Gore $\succ$ Nader
20%:	Gore $\succ$ Nader $\succ$ Bush
20%:	Gore $\succ$ Bush $\succ$ Nader
11%:	Nader $\succ$ Gore $\succ$ Bush

Questions:

- Who wins?
- Is that a fair outcome?
- What would your advice to the Nader-supporters have been?

Three Voting Rules

How should n *voters* choose from a set of m *alternatives*?

Here are three *voting rules* (there are many more):

- *Plurality:* elect the alternative ranked first most often (i.e., each voter assigns 1 point to an alternative of her choice, and the alternative receiving the most points wins)
- *Plurality with runoff:* run a plurality election and retain the two front-runners; then run a majority contest between them
- *Borda:* each voter gives $m-1$ points to the alternative she ranks first, $m-2$ to the alternative she ranks second, etc.; and the alternative with the most points wins

Example

Consider this election with nine *voters* having to choose from three *alternatives* (namely what drink to order for a common lunch):

4 Dutchmen:	Milk	$\succ$	Beer	$\succ$	Wine
3 Frenchmen:	Wine	$\succ$	Beer	$\succ$	Milk
2 Germans:	Beer	$\succ$	Wine	$\succ$	Milk

Which beverage *wins* the election for

- the plurality rule?
- plurality with runoff?
- the Borda rule?

Positional Scoring Rules

We can generalise the idea underlying the Borda rule as follows:

A *positional scoring rule* is given by a *scoring vector* $s = \langle s_1, \ldots, s_m \rangle$ with $s_1 \geqslant s_2 \geqslant \cdots \geqslant s_m$ and $s_1 > s_m$.

Each voter submits a ranking of the m alternatives. Each alternative receives s_i points for every voter putting it at the ith position.

The alternative(s) with the highest score (sum of points) win(s).

Examples:

- *Borda rule* = PSR with scoring vector $\langle m{-}1, m{-}2, \ldots, 0 \rangle$
- *Plurality rule* = PSR with scoring vector $\langle 1, 0, \ldots, 0 \rangle$
- *Antiplurality rule* = PSR with scoring vector $\langle 1, \ldots, 1, 0 \rangle$
- For any $k \leqslant m$, *k-approval* = PSR with $\langle \underbrace{1, \ldots, 1}_{k}, 0, \ldots, 0 \rangle$

The Condorcet Principle

The Marquis de Condorcet was a public intellectual living in France during the second half of the 18th century.

An alternative that beats every other alternative in pairwise majority contests is called a *Condorcet winner*.

There may be no Condorcet winner; witness the *Condorcet paradox:*

Ann:	$A \succ B \succ C$
Bob:	$B \succ C \succ A$
Cindy:	$C \succ A \succ B$

Whenever a Condorcet winner exists, then it must be *unique*.

A voting rule satisfies the *Condorcet principle* if it elects (only) the Condorcet winner whenever one exists.

M. le Marquis de Condorcet. *Essai sur l'application de l'analyse à la probabilté des décisions rendues a la pluralité des voix*. Paris, 1785.

All PSR's Violate the Condorcet Principle (!)

Consider the following example:

$$\begin{array}{ll} \text{3 voters:} & A \succ B \succ C \\ \text{2 voters:} & B \succ C \succ A \\ \text{1 voter:} & B \succ A \succ C \\ \text{1 voter:} & C \succ A \succ B \end{array}$$

A is the *Condorcet winner*; she beats both B and C $4:3$. But any *positional scoring rule* makes B win (because $s_1 \geqslant s_2 \geqslant s_3$):

$$\begin{array}{ll} A: & 3 \cdot s_1 + 2 \cdot s_2 + 2 \cdot s_3 \\ B: & 3 \cdot s_1 + 3 \cdot s_2 + 1 \cdot s_3 \\ C: & 1 \cdot s_1 + 2 \cdot s_2 + 4 \cdot s_3 \end{array}$$

Thus, *no positional scoring rule* for three (or more) alternatives will satisfy the *Condorcet principle*.

Another Example: Sequential Majority Voting

Yet another rule: *sequential majority voting* means running a series of pairwise majority contests, with the winner always getting promoted to the next stage. This is guaranteed to meet the Condorcet principle.

But there is another problem. Consider this example:

```
        o
       / \
      o   D
     / \
    o   A
   / \
  B   C
```

Take this profile with three agents:

Ann: $\mathtt{A} \succ \mathtt{B} \succ \mathtt{C} \succ \mathtt{D}$

Bob: $\mathtt{B} \succ \mathtt{C} \succ \mathtt{D} \succ \mathtt{A}$

Cindy: $\mathtt{C} \succ \mathtt{D} \succ \mathtt{A} \succ \mathtt{B}$

D *wins!* (despite being dominated by C)

This is a violation of the (weak) *Pareto principle:* if you can make a change that improves everyone's welfare, then do make that change.

Vilfredo Pareto was an Italian economist active around 1900.

Insights so far

Our examples have demonstrated:

- There are different *methods* of aggregation (voting rules).
  We need clear *criteria* for choosing one.

- There are all sorts of *paradoxes* (counterintuitive outcomes).
  We need to clearly specify *desiderata* for methods of aggregation to have a chance of understanding these problems.

Formal Framework: Voting Theory

Basic terminology and notation:

- finite set of *individuals* $\mathcal{N} = \{1, \ldots, n\}$, with $n \geqslant 2$
- (usually finite) set of *alternatives* $\mathcal{X} = \{x_1, x_2, x_3, \ldots\}$
- Denote the set of *linear orders* on $\mathcal{X}$ by $\mathcal{L}(\mathcal{X})$.
  Preferences (or *ballots*) are taken to be elements of $\mathcal{L}(\mathcal{X})$.
- A *profile* $\boldsymbol{R} = (R_1, \ldots, R_n) \in \mathcal{L}(\mathcal{X})^n$ is a vector of preferences.
- We shall write $N^{\boldsymbol{R}}_{x \succ y}$ for the set of individuals that rank alternative x above alternative y under profile $\boldsymbol{R}$.

For now we focus on the case of (irresolute) *voting rules*, mapping any given profile to a nonempty set of winning alternatives:

$$F : \mathcal{L}(\mathcal{X})^n \to 2^{\mathcal{X}} \setminus \{\emptyset\}$$

The Axiomatic Method

Many important classical results in social choice theory are *axiomatic*. They formalise desirable properties as "*axioms*" and then establish:

- *Characterisation Theorems*, showing that a particular (class of) mechanism(s) is the only one satisfying a given set of axioms
- *Impossibility Theorems*, showing that there exists *no* aggregation mechanism satisfying a given set of axioms

Anonymity and Neutrality

Two very basic axioms:

- F is *anonymous* if $F(R_1, \ldots, R_n) = F(R_{\pi(1)}, \ldots, R_{\pi(n)})$ for any profile $(R_1, \ldots, R_n)$ and any permutation $\pi : \mathcal{N} \to \mathcal{N}$.
- F is *neutral* if $F(\pi(\boldsymbol{R})) = \pi(F(\boldsymbol{R}))$ for any profile $\boldsymbol{R}$ and any permutation $\pi : \mathcal{X} \to \mathcal{X}$ (with π extended to profiles and sets of alternatives in the natural manner).

Positive Responsiveness

A (not necessarily resolute) voting rule satisfies *positive responsiveness* if, whenever some voter raises a (possibly tied) winner $x^\star$ in her ballot, then $x^\star$ will become the *unique* winner. Formally:

> F satisfies positive responsiveness if $x^\star \in F(\boldsymbol{R})$ implies $\{x^\star\} = F(\boldsymbol{R'})$ for any alternative $x^\star$ and any two *distinct* profiles $\boldsymbol{R}$ and $\boldsymbol{R'}$ with $N^{\boldsymbol{R}}_{x^\star \succ y} \subseteq N^{\boldsymbol{R}'}_{x^\star \succ y}$ and $N^{\boldsymbol{R}}_{y \succ z} = N^{\boldsymbol{R'}}_{y \succ z}$ for all $y, z \in \mathcal{X} \setminus \{x^\star\}$.

Recall: $N^{\boldsymbol{R}}_{x \succ y}$ is the set of voters ranking x above y in profile $\boldsymbol{R}$.

May's Theorem

Now we can fully characterise the plurality rule (which is often called the *simple majority rule* when there are only two alternatives):

Theorem 1 (May, 1952) *A voting rule for two alternatives satisfies anonymity, neutrality, and positive responsiveness if and only if it is the simple majority rule.*

Next: proof

K.O. May. A Set of Independent Necessary and Sufficient Conditions for Simple Majority Decisions. *Econometrica*, 20(4):680–684, 1952.

Proof Sketch

Clearly, simple majority does satisfy all three properties. ✓

Now for the other direction:

Assume the number of voters is *odd* (other case: similar) $\rightsquigarrow$ no ties.

There are two possible ballots: $a \succ b$ and $b \succ a$.
Anonymity $\rightsquigarrow$ only number of ballots of each type matters.

Denote as A the set of voters voting $a \succ b$ and as B those voting $b \succ a$. Distinguish two cases:

- Whenever $|A| = |B| + 1$ then only a wins. Then, by PR, a wins whenever $|A| > |B|$ (which is exactly the simple majority rule). ✓
- There exist A, B with $|A| = |B| + 1$ but b wins. Now suppose one a-voter switches to b. By PR, now only b wins. But now $|B'| = |A'| + 1$, which is symmetric to the earlier situation, so by neutrality a should win $\rightsquigarrow$ contradiction. ✓

Formal Framework: Preference Aggregation

Basic terminology and notation:

- finite set of *individuals* $\mathcal{N} = \{1, \ldots, n\}$, with $n \geqslant 2$
- (usually finite) set of *alternatives* $\mathcal{X} = \{x_1, x_2, x_3, \ldots\}$
- Denote the set of *linear orders* on $\mathcal{X}$ by $\mathcal{L}(\mathcal{X})$.
  Preferences (or *ballots*) are taken to be elements of $\mathcal{L}(\mathcal{X})$.
- A *profile* $\boldsymbol{R} = (R_1, \ldots, R_n) \in \mathcal{L}(\mathcal{X})^n$ is a vector of preferences.
- We shall write $N^{\boldsymbol{R}}_{x \succ y}$ for the set of individuals that rank alternative x above alternative y under profile $\boldsymbol{R}$.

Next we are interested in preference aggregation mechanisms that map any profile of preferences to a single collective preference.

The proper technical term is *social welfare function* (SWF):

$$F : \mathcal{L}(\mathcal{X})^n \to \mathcal{L}(\mathcal{X})$$

The Pareto Condition

A SWF F satisfies the (weak) *Pareto condition* if, whenever all individuals rank x above y, then so does society:

$$N^{\boldsymbol{R}}_{x \succ y} = \mathcal{N} \text{ implies } (x, y) \in F(\boldsymbol{R})$$

Independence of Irrelevant Alternatives (IIA)

A SWF F satisfies *IIA* if the relative social ranking of two alternatives only depends on their relative individual rankings:

$$N^{\boldsymbol{R}}_{x \succ y} = N^{\boldsymbol{R}'}_{x \succ y} \text{ implies } (x, y) \in F(\boldsymbol{R}) \Leftrightarrow (x, y) \in F(\boldsymbol{R}')$$

In other words: if x is socially preferred to y, then this should not change when an individual changes her ranking of z.

Arrow's Theorem

A SWF F is a *dictatorship* if there exists a "dictator" $i \in \mathcal{N}$ such that $F(\boldsymbol{R}) = R_i$ for any profile $\boldsymbol{R}$, i.e., if the outcome is always identical to the preference supplied by the dictator.

Theorem 2 (Arrow, 1951) *Any SWF for $\geqslant 3$ alternatives that satisfies the Pareto condition and IIA must be a dictatorship.*

We will see a proof on Thursday (lecture on graph aggregation).

K.J. Arrow. *Social Choice and Individual Values.* John Wiley and Sons, 2nd edition, 1963. First edition published in 1951.

Remarks

- Note that this is a *surprising* result!
- Note that the theorem does *not* hold for *two* alternatives.
- Note that the *opposite direction* clearly holds: any dictatorship satisfies both the Pareto condition and IIA.
- Common misunderstanding: the SWF being *dictatorial* does not just mean that the outcome coincides with the preferences of some individual (rather: it's *the same* dictator for any profile).
- Arrow's Theorem is often read as an *impossibility theorem:*

  There exists no SWF for $\geqslant 3$ *alternatives that is Paretian, independent, and nondictatorial.*

- Significance of the result: (a) the result itself; (b) *general* theorem rather than just another observation about a flaw of a specific procedure; (c) *methodology* (precise statement of "axioms").

Logic and Social Choice Theory

Examples for applications of logic in social choice theory:

- Formal minimalism (Pauly, Synthese 2008)
- Verification of proofs (e.g., Nipkow, JAR 2009)
- Automation of tasks (Tang & Lin, AIJ 2009; Geist & E., JAIR 2011)

Much of classical social choice theory has been modelled in logic:

- Classical first-order logic (Grandi & E., JPL 2013)
- Tailor-made modal logics (e.g., Ågotnes et al., JAAMAS 2010)

But all of these approaches have some shortcomings:

- modelling of *universal domain* assumption not elegant
- set of *individuals* fixed to *specific size* (or at least not to any *finite* set)
- gap between logical modelling and suitability for *automated reasoning*

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

Example: Judgment Aggregation

	p	$p \rightarrow q$	q
Judge 1:	True	True	True
Judge 2:	True	False	False
Judge 3:	False	True	False
		?	

Formal Framework: Judgment Aggregation

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:	True	False	True	$J_1 = \{p, \neg q, p \vee q\}$
Agent 2:	True	True	True	$J_2 = \{p, q, p \vee q\}$
Agent 3:	False	False	False	$J_3 = \{\neg p, \neg q, \neg(p \vee q)\}$

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].

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 3 (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.

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.

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 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. ✓

Summary

We have seen examples from different areas of aggregation theory:

- voting and preference aggregation
- judgment aggregation

We have seen several typical features of work in this field:

- paradoxical examples
- desirable properties and their formalisation as axioms
- characterisations (May) and impossibilities (Arrow, List-Pettit)

Full technical details on everything discussed today are available in the expository article cited below. Programme for the rest of the week:

- Lecture 2: Binary Aggregation with Integrity Constraints
- Lecture 3: Graph Aggregation
- Lecture 4: Collective Annotation

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