

Tutorial on Computational Social Choice

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

[<http://www.illc.uva.nl/~ulle/teaching/kutaisi-2011/>]


Table of Contents

Introduction 3
 The Axiomatic Method in Social Choice Theory 6
 Social Choice in Combinatorial Domains 23
 Judgment Aggregation 40
 Conclusion 53

Introduction

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”?



SCT is traditionally studied in Economics and Political Science, but now also by “us”: *Computational Social Choice*.

Tutorial Overview

This tutorial will provide an introduction to both classical social choice theory and computational social choice. We will focus on three topics (one per lecture), which together highlight the diverse ways in which logic has been applied in this field:

- The Axiomatic Method in Social Choice Theory
- Social Choice in Combinatorial Domains
- Judgment Aggregation

The tutorial is based on the review paper cited below.

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

The Axiomatic Method in Social Choice Theory

Outline

This will be an introduction to the “axiomatic method” in social choice theory, in which we formalise normative intuitions about the proper way of aggregating preferences by stating so-called “axioms” and then investigate the consequences of those axioms.

Material to be covered in this part:

- Types of aggregation rules
- Examples for axioms (desirable properties of aggregators)
- Arrow’s Impossibility Theorem (with proof)
- Gibbard-Satterthwaite Theorem (very briefly)

Three Voting Rules

Voting is the prototypical form of collective decision making.

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

- *Plurality*: elect the candidate ranked first most often (i.e., each voter assigns one point to a candidate of her choice, and the candidate receiving the most votes wins)
- *Borda*: each voter gives $m-1$ points to the candidate she ranks first, $m-2$ to the candidate she ranks second, etc., and the candidate with the most points wins
- *Approval*: voters can approve of as many candidates as they wish, and the candidate with the most approvals wins

Example

Suppose there are three *candidates* (A, B, C) and 11 *voters* with the following *preferences* (where boldface indicates *acceptability*, for AV):

5 voters think: **A** \succ B \succ C

4 voters think: **C** \succ B \succ A

2 voters think: **B** \succ **C** \succ A

Assuming the voters vote *sincerely*, who *wins* the election for

- the plurality rule?
- the Borda rule?
- approval voting?

Conclusion: We need to be very clear about what properties we are looking for. So let's formalise this ...

Formal Framework

Basic terminology and notation:

- finite set of *individuals* $\mathcal{N} = \{1, \dots, n\}$, with $n \geq 2$
- (usually finite) set of *alternatives* $\mathcal{X} = \{x_1, x_2, x_3, \dots\}$
- 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* $\mathbf{R} = (R_1, \dots, R_n) \in \mathcal{L}(\mathcal{X})^{\mathcal{N}}$ is a vector of preferences.

Social choice theory studies various forms of aggregation, e.g.:

- A *social choice function* (SCF) or *voting rule* is a function $F : \mathcal{L}(\mathcal{X})^{\mathcal{N}} \rightarrow 2^{\mathcal{X}} \setminus \{\emptyset\}$ mapping any given profile to a nonempty set of winners (F is called *resolute* if $|F(\mathbf{R})| = 1$ for any \mathbf{R}).
- A *social welfare function* (SWF) is a function $F : \mathcal{L}(\mathcal{X})^{\mathcal{N}} \rightarrow \mathcal{L}(\mathcal{X})$ mapping any given profile to a (single) collective preference order.

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

We will first see a few of these axioms ...

Remark: On the following slides we work with SWFs, but very similar definitions and results exist for SCFs.

Anonymity and Neutrality

Two very basic axioms:

- A SWF F is *anonymous* if *individuals* are treated symmetrically:

$$F(R_1, \dots, R_n) = F(R_{\pi(1)}, \dots, R_{\pi(n)})$$

for any profile \mathbf{R} and any permutation $\pi : \mathcal{N} \rightarrow \mathcal{N}$

- A SWF F is *neutral* if *alternatives* are treated symmetrically:

$$F(\pi(\mathbf{R})) = \pi(F(\mathbf{R}))$$

for any profile \mathbf{R} and any permutation $\pi : \mathcal{X} \rightarrow \mathcal{X}$
(with π extended to preferences and profiles in the natural manner)

Keep in mind:

- not every SWF will satisfy every axiom we state here
- axioms are meant to be *desirable* properties

The Pareto Condition

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

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

This is a standard condition going back to the work of the Italian economist Vilfredo Pareto (1848–1923).

Notation: Here and in the sequel, we write $N_{x \succ y}^{\mathbf{R}}$ for the set of individuals that rank alternative x above alternative y in profile \mathbf{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:

$$N_{x \succ y}^{\mathbf{R}} = N_{x \succ y}^{\mathbf{R}'} \text{ implies } (x, y) \in F(\mathbf{R}) \Leftrightarrow (x, y) \in F(\mathbf{R}')$$

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

Pareto and IIA look like basic desirable properties. Yet, surprisingly, satisfying both properties is *impossible* in the following sense:

Theorem 1 (Arrow, 1951) Any SWF for ≥ 3 alternatives that satisfies the *Pareto condition* and *IIA* must be a *dictatorship*.

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

Note that:

- The theorem does *not* hold for *two* alternatives.
- The *opposite direction* also holds: dictatorship \Rightarrow Pareto + IIA.

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

Proof

We'll sketch a proof adapted from Sen (1986), using the “decisive coalition” technique. Full details are in my review paper.

Claim: Any SWF for ≥ 3 alternatives that satisfies the *Pareto condition* and *IIA* must be a *dictatorship*.

So let F be a SWF for ≥ 3 alternatives that satisfies Pareto and IIA.

Call a coalition $G \subseteq \mathcal{N}$ *decisive* on (x, y) iff $G \subseteq N_{x \succ y}^{\mathbf{R}} \Rightarrow (x, y) \in F(\mathbf{R})$.

Proof Plan:

- Pareto condition = \mathcal{N} is decisive for all pairs of alternatives
- Lemma: G with $|G| \geq 2$ *decisive* for all pairs \Rightarrow some $G' \subset G$ as well
- Thus (by induction), there's a decisive coalition of size 1 (a *dictator*).

A.K. Sen. *Social Choice Theory*. In K.J. Arrow and M.D. Intriligator (eds.), *Handbook of Mathematical Economics*, Volume 3, North-Holland, 1986.

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

About Decisiveness

Recall: $G \subseteq \mathcal{N}$ *decisive* on (x, y) iff $G \subseteq N_{x \succ y}^{\mathbf{R}} \Rightarrow (x, y) \in F(\mathbf{R})$

Call $G \subseteq \mathcal{N}$ *weakly decisive* on (x, y) iff $G = N_{x \succ y}^{\mathbf{R}} \Rightarrow (x, y) \in F(\mathbf{R})$.

Claim: G weakly decisive on $(x, y) \Rightarrow G$ decisive on *any* pair (x', y')

Proof: Suppose x, y, x', y' are all distinct (other cases: similar).

Consider a profile where individuals express these preferences:

- Members of G : $x' \succ x \succ y \succ y'$
- Others: $x' \succ x$ and $y \succ y'$ and $y \succ x$ (note: x' vs. y' not fixed)

From G being weakly decisive for (x, y) : society ranks $x \succ y$

From Pareto: society ranks $x' \succ x$ and $y \succ y'$

Thus, from transitivity: society ranks $x' \succ y'$

Note that this works for any ranking of x' vs. y' by non- G individuals.

By IIA, it still works if individuals change their non- x' -vs.- y' rankings.

Thus, for *any* profile \mathbf{R} with $G \subseteq N_{x' \succ y'}^{\mathbf{R}}$ we get $(x', y') \in F(\mathbf{R})$. \checkmark

Contraction Lemma

Claim: If $G \subseteq \mathcal{N}$ with $|G| \geq 2$ is a coalition that is decisive on all pairs of alternatives, then so is some nonempty coalition $G' \subset G$.

Proof: Take any nonempty G_1, G_2 with $G = G_1 \cup G_2$ and $G_1 \cap G_2 = \emptyset$.

Recall that there are ≥ 3 alternatives. Consider this profile:

- Members of G_1 : $x \succ y \succ z \succ \text{rest}$
- Members of G_2 : $y \succ z \succ x \succ \text{rest}$
- Others: $z \succ x \succ y \succ \text{rest}$

As $G = G_1 \cup G_2$ is decisive, society ranks $y \succ z$. Two cases:

- (1) Society ranks $x \succ z$: Exactly G_1 ranks $x \succ z \Rightarrow$ By IIA, in any profile where exactly G_1 ranks $x \succ z$, society will rank $x \succ z \Rightarrow G_1$ is weakly decisive on (x, z) . Hence (previous slide): G_1 is decisive on all pairs.
- (2) Society ranks $z \succ x$, i.e., $y \succ x$: Exactly G_2 ranks $y \succ x \Rightarrow \dots \Rightarrow G_2$ is decisive on all pairs.

Hence, one of G_1 and G_2 will always be decisive. \checkmark

This concludes the proof of Arrow's Theorem.

Example

We now switch to voting rules (SCFs). Under the *plurality rule* the candidate receiving the highest number of votes wins.

Assume the preferences of the people in, say, Florida are as follows:

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

So even if nobody is cheating, Bush will win this election. **But:**

- It would have been in the interest of the Nader supporters to *manipulate*, i.e., to misrepresent their preferences.

Is there a better voting rule that avoids this problem?

The Gibbard-Satterthwaite Theorem

We are interested in this property:

- A resolute SCF F is *strategy-proof* if there exists no profile where some voter can obtain a preferred outcome by changing her ballot.

We again obtain a surprising negative result:

Theorem 2 (Gibbard-Satterthwaite, 1973/75) Any *resolute SCF* for ≥ 3 alternatives that is *surjective* and *strategy-proof* must be a *dictatorship*.

See my review paper for full definitions as well as a proof using once more the "decisive coalition" technique.

A. Gibbard. Manipulation of Voting Schemes: A General Result. *Econometrica*, 41(4):587–601, 1973.

M.A. Satterthwaite. Strategy-proofness and Arrow's Conditions. *Journal of Economic Theory*, 10:187–217, 1975.

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

Complexity as a Barrier against Manipulation

The Gibbard-Satterthwaite Theorem shows that manipulation is *possible* for any rule. But how *hard* is it to find a manipulating ballot?

Bartholdi et al. (1989) were the first to suggest looking for voting rules for which strategic manipulation is NP-hard.

- For most standard rules this does not work: it's clearly easy for *plurality* and (less obviously so) also for the *Borda rule*.
- Bartholdi and Orlin (1991) showed that the manipulation problem for *Single Transferable Vote* is *NP-hard*.

See Faliszewski et al. (2010) for a review of the state of the art.

J.J. Bartholdi III, C.A. Tovey, and M.A. Trick. The Computational Difficulty of Manipulating an Election. *Soc. Choice and Welfare*, 6(3):227–241, 1989.

J.J. Bartholdi III and J.B. Orlin. Single Transferable Vote Resists Strategic Voting. *Social Choice and Welfare*, 8(4):341–354, 1991.

P. Faliszewski, E. Hemaspaandra, and L.A. Hemaspaandra. Using Complexity to Protect Elections. *Communications of the ACM*, 553(11):74–82, 2010.

Summary: Axiomatic Method

This has been a short introduction to the axiomatic method in social choice theory, which makes use of basic principles from logic, albeit informally. We have seen:

- A formal model for preference aggregation and voting (SWF/SCF)
- Axioms: anonymity, neutrality, Pareto, IIA, strategy-proofness
- Arrow: Pareto + IIA \Rightarrow dictatorial
- Gibbard-Satterthwaite: strategy-proof + surjective \Rightarrow dictatorial

We have also seen one example where modern computational techniques open up a new perspective on old problems of social choice:

- Complexity theory is relevant to the analysis of the problem of strategic manipulation in voting.

Social Choice in Combinatorial Domains

The Paradox of Multiple Elections

13 voters are asked to each vote *yes* or *no* on three issues:

- 3 voters each vote for YNN, NYN, NNY.
- 1 voter each votes for YYY, YYN, YNY, NYY.
- No voter votes for NNN.

If we use the *simple majority* rule *issue-by-issue*, then NNN wins, because on each issue 7 out of 13 vote *no*.

This is an instance of the *paradox of multiple elections*: the winning combination received not a single vote!

S.J. Brams, D.M. Kilgour, and W.S. Zwicker. The Paradox of Multiple Elections. *Social Choice and Welfare*, 15(2):211–236, 1998.

Outline

Elections often have a *combinatorial structure*:

- Electing a committee of k members from amongst n candidates.
- During a referendum (in Switzerland, California, places like that), voters may be asked to vote on several propositions.

Clearly, the number of alternatives can quickly become *very large*.

So we face both a *choice-theoretic* and a *computational challenge*.

Things to be discussed today:

- Definition of the problem: voting in combinatorial domains
- Different approaches to voting in combinatorial domains
- Compact preference representation languages

More details are in the expository paper by Chevaleyre et al. (2008).

Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. Preference Handling in Combinatorial Domains: From AI to Social Choice. *AI Magazine*, 29(4):37–46, 2008.

Voting in Combinatorial Domains

The problem of *voting in combinatorial domains*:

- Domain: variables X_1, \dots, X_p with finite domains D_1, \dots, D_p
- Voters have preferences over set of combinations $D_1 \times \dots \times D_p$.
- What should be the winning combination in $D_1 \times \dots \times D_p$?

Today we only consider *binary* variables: $D_k = \{x_k, \bar{x}_k\}$.

- We have seen that voting issue-by-issue can lead to paradoxical outcomes. What other approaches are there?

Approach 1: Plurality on Combinations

Idea: Vote for combinations directly: ask each voter for her most preferred combination and apply the plurality rule.

This avoids the paradox we have seen and is computationally light.

Problem: This may lead to almost random decisions, unless domains are fairly small and there are many voters.

Example: Suppose there are 10 binary issues and 20 voters. Then there are $2^{10} = 1024$ combinations to vote for. Under the plurality rule, chances are very highly ($\sim 83\%$) that no combination receives more than one vote (so the tie-breaking rule decides everything).

Approach 2: Other Rules on Combinations

Idea: Vote for combinations directly, using your favourite voting rule with the full set of combinations as the set of alternatives.

If we use a voting rule that elicits more information than the plurality rule, then we can avoid the arbitrariness problem noted before.

Problem: This will only be possible in very small domains, certainly when the voting rule requires a complete ranking of all the candidates (such as the Borda rule).

Example: Suppose there are six binary issues. This makes $2^6 = 64$ possible combinations. Hence, under the Borda rule, each voter has to choose from amongst $64! \approx 1.27 \cdot 10^{89}$ possible ballots.

Approach 3: Distance-based Aggregation

Idea: Elicit preferred choices issue-by-issue (as in the paradox), but find a better way to aggregate this information.

Distance-based approaches are promising candidates:

- Define a *distance* metric on ballots (0-1 vectors).
- Extend it to measure distance of a ballot/outcome to a profile.
- Choose the outcome that *minimises* the distance to the profile.

Example: Brams et al. (2007) propose to elect the combination that minimises the maximal Hamming distance to any of the voter ballots.

S.J. Brams, D.M. Kilgour, and M.R. Sanver. A Minimax Procedure for Electing Committees. *Public Choice*, 132:401–420, 2007.

Approach 4: Sequential Voting

Idea: Vote separately on each issue, but do so sequentially to give voters the opportunity to make their vote for one issue dependent on other issues already decided.

We will discuss two basic results for this approach.

Sequential Voting and Condorcet Losers

A *Condorcet loser* is a candidate that loses against any other candidate in a pairwise contest. Electing a CL is very bad.

Lacy and Niou (2000) show that sequential voting can avoid this:

Proposition 3 *Sequential voting (with plurality) over binary issues never results in a winning combination that is a Condorcet loser.*

Proof: Just think what happens during the election for the final issue. The winning combination cannot be a Condorcet loser, because it does, at least, win against the other combination that was still possible after the penultimate election. ✓

A stronger requirement is *Condorcet consistency*: elect the *Condorcet winner* whenever it exists. Sequential voting *cannot* guarantee this.

D. Lacy and E.M.S. Niou. A Problem with Referendums. *Journal of Theoretical Politics*, 12(1):5–31, 2000.

Sequential Voting and Condorcet Winners

A preference order induces a *preferential dependency graph* on issues: issue X depends on Y if there exist situations where you need to know the value of Y before you can decide on your preference regarding X .

Now suppose that there exists a linear order on the issues that is compatible with each voter's preferential dependency graph.

Idea: Vote sequentially in this order!

Lang and Xia (2009) have shown (proof omitted but easy):

Proposition 4 *Under above assumptions, sequential voting is Condorcet-consistent whenever all of the local voting rules are.*

This is useful, particularly when the issues are binary (as then any reasonable local procedure will be Condorcet-consistent).

J. Lang and L. Xia. Sequential Composition of Voting Rules in Multi-issue Domains. *Mathematical Social Sciences*, 57(3):304–324, 2009.

Approach 5: Combinatorial Vote

Idea: Ask voters to report their ballots using a compact preference *representation language* and apply your favourite voting rule to the succinctly encoded ballots received.

Lang (2004) calls this approach *combinatorial vote*.

Discussion: A promising approach, but not too much is known to date about what would be good choices for preference representation languages or voting rules, or what algorithms to use to compute the winners. Also, complexity can be expected to be very high.

J. Lang. Logical Preference Representation and Combinatorial Vote. *Annals of Mathematics and Artificial Intelligence*, 42(1–3):37–71, 2004.

Logic-based Preference Representation

Think of $\{X_1, \dots, X_p\}$ as propositional variables and use propositional formulas to express *goals*. Use a *number* to indicate importance of a goal.

- **Weighted goals:** A set G of weighted goals induces a *utility function* $u_G : D_1 \times \dots \times D_p \rightarrow \mathbb{R}$, mapping each combination/model M to

$$u_G(M) = \sum_{(\varphi, w) \in G[M]} w \quad \text{where } G[M] = \{(\varphi, w) \in G \mid M \models \varphi\}$$

- **Prioritised goals:** (φ, k_1) has higher priority than (ψ, k_2) if $k_1 > k_2$. Under the *lexicographic* form of aggregation, we prefer M to M' if there exists a k such that for all $j > k$ both M and M' satisfy the same number of goals of priority j , and M satisfies more goals of priority k .

Other forms of aggregation are possible (in both settings).

J. Lang. Logical Preference Representation and Combinatorial Vote. *Annals of Mathematics and Artificial Intelligence*, 42(1–3):37–71, 2004.

J. Uckelman. *More than the Sum of its Parts: Compact Preference Representation over Combinatorial Domains*. PhD thesis, ILLC, University of Amsterdam, 2009.

Combinatorial Vote: Example

Use the language of *prioritised goals* (1 has higher priority than 0) with *lexicographic aggregation* together with the *Borda rule*:

- Voter 1: $\{X:1, Y:0\}$ induces order $xy \succ_1 x\bar{y} \succ_1 \bar{x}y \succ_1 \bar{x}\bar{y}$
- Voter 2: $\{X \vee \neg Y:0\}$ induces order $x\bar{y} \sim_2 xy \sim_2 \bar{x}\bar{y} \succ_2 \bar{x}y$
- Voter 3: $\{\neg X:0, Y:0\}$ induces order $\bar{x}y \succ_3 \bar{x}\bar{y} \sim_3 xy \succ_3 x\bar{y}$

As the induced orders need not be strict linear orders, we use a *generalisation of the Borda rule*: an alternative gets as many points as she dominates other alternatives. So we get these Borda scores:

$$\begin{aligned} xy : 3 + 1 + 1 = 5 & & \bar{x}y : 1 + 0 + 3 = 4 \\ x\bar{y} : 2 + 1 + 0 = 3 & & \bar{x}\bar{y} : 0 + 1 + 1 = 2 \end{aligned}$$

So combinatorial alternative xy wins.

Combinatorial vote *proper* would be to compute the winner *directly* from the goalbases, without the detour via the induced orders.

Single Goals and Generalised Plurality

Next a complexity result exemplifying the limitations of the approach.

We will work with the following language and voting rule:

- Using the *language of single goals*, each voter specifies just one goal (an arbitrary propositional formula) with priority 1.
- Under the *generalised plurality rule*, a voter gives 1 point to each undominated alternative.

Here are two examples, for the set of variables $\{X, Y\}$:

- The goal $\neg X \wedge Y$ induces the order $\bar{x}y \succ xy \sim x\bar{y} \sim \bar{x}\bar{y}$, so only combination $\bar{x}y$ receives 1 point.
- The goal $X \vee Y$ induces the order $xy \sim \bar{x}y \sim x\bar{y} \succ \bar{x}\bar{y}$, so combinations $xy, \bar{x}y, x\bar{y}$ receive 1 point each.

Winner Verification under Plurality

Define the following decision problem, for a preference representation language \mathcal{L} and a voting rule F :

AMONG-WINNERS(\mathcal{L}, F)

Instance: Profile \mathbf{R} expressed in \mathcal{L} ; combination x^* .

Question: Is $x^* \in F(\mathbf{R})$?

The following result is due to Lang (2004):

Proposition 5 AMONG-WINNERS is *coNP-complete* for the language of *single goals* and the generalised *plurality* rule.

Proof: Omitted (but easy).

J. Lang. Logical Preference Representation and Combinatorial Vote. *Annals of Mathematics and Artificial Intelligence*, 42(1–3):37–71, 2004.

Compact Preference Representation

The most important language for COMSOC are *CP-nets*. Other important languages include *weighted* and *prioritised goals*.

The study of these languages is an interesting topic in its own right. Questions to investigate (and typical results) include:

- *Expressivity*: with sum aggregation, positive goals with positive weights can express all monotonic functions, and only those
- *Succinctness*: with sum aggregation, conjunctions of literals can express anything general formulas can, but do so less succinctly
- *Complexity*: with max aggregation, social welfare maximisation is NP-hard, even if all weighted goals have the form $(p \wedge q, 1)$

J. Uckelman, Y. Chevaleyre, U. Endriss, and J. Lang. Representing Utility Functions via Weighted Goals. *Mathematical Logic Quarterly*, 55(4):341–361, 2009.

J. Uckelman and U. Endriss. Compactly Representing Utility Functions Using Weighted Goals and the Max Aggregator. *Artif. Intell.*, 174(15):1222–1246, 2010.

Summary: Combinatorial Domains

We have seen several approaches for tackling the problem of voting in *combinatorial domains* (i.e., voting in multi-issue elections).

To date, no clear solution has emerged. Good candidates:

- Distance-based approaches
- Sequential voting
- Voting with compactly expressed preferences

Any approach has to balance a *choice-theoretic challenge* (eliciting too little information from voters leads to paradoxes) and a *computational challenge* (eliciting too much information may be intractable).

Great research area in its own right: *compact preference representation*

Judgment Aggregation

Judgment Aggregation

Preferences are not the only structures we may wish to aggregate.
JA studies the aggregation of judgments on inter-related propositions.

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

Paradox: each *individual* judgment set is *consistent*, but the *collective* judgment arrived at by using the *majority rule* is not

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

C. List and C. Puppe. Judgment Aggregation: A Survey. In *Handbook of Rational and Social Choice*. Oxford University Press, 2009.

Outline

This final part of the tutorial will be an introduction to the field of judgement aggregation. Topics to be covered:

- Formal framework, possible aggregation procedures, axioms
- An impossibility theorem
- A way around the impossibility: domain restrictions
- Complexity of judgment aggregation
- Links between preference aggregation and judgment aggregation

For a more thorough introduction to JA, see the papers cited below.

C. List. The Theory of Judgment Aggregation: An Introductory Review. *Synthese*. In press (2011).

C. List and C. Puppe. Judgment Aggregation: A Survey. In *Handbook of Rational and Social Choice*. Oxford University Press, 2009.

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 $\varphi \in \Phi$
- *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 of n individuals is a function mapping a profile of complete and consistent individual judgment sets to a single collective judgment set: $F : \mathcal{J}(\Phi)^n \rightarrow 2^\Phi$.

Aggregation Procedures

Examples for procedures considered in the literature:

- *Majority rule*: not always consistent (as we have seen)
- *Premise-based procedure*: might be appropriate in practice, but notion of “premise” not easy to define
- *Conclusion-based procedure*: similar issues
- *Distance-based procedure*: choose consistent judgment set that is “closest” to the profile (consistent by design)

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 $\mathbf{J} \in \mathcal{J}(\Phi)$, 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 φ in the agenda Φ and profiles \mathbf{J} and \mathbf{J}' in $\mathcal{J}(\Phi)$, if $\varphi \in J_i \Leftrightarrow \varphi \in J'_i$ for all i , then $\varphi \in F(\mathbf{J}) \Leftrightarrow \varphi \in F(\mathbf{J}')$.

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

Impossibility Theorem

It turns out that our axioms are too demanding:

Theorem 6 (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: Similar impossibilities arise for other agendas with some minimal structural complexity. More recent results fully characterise agendas where consistent aggregation is (im)possible.

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

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 7 *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. ✓

Complexity of Judgment Aggregation

So, under certain domain restrictions, we can safely use the majority rule and never encounter a paradox. But what in general?

Call an agenda Φ *safe* (for the majority rule) if applying the majority rule to *any* profile in $\mathcal{J}(\Phi)^{\mathcal{N}}$ will yield a consistent judgment set.

Unfortunately, not only is it rare that an agenda will guarantee safety, but recognising those agendas that are safe is also very difficult:

Proposition 8 *Deciding the safety of the agenda problem for the majority rule is Π_2^P -complete.*

Π_2^P is also known as “coNP with an NP oracle”, i.e., this is really hard.

Proof: Omitted.

U. Endriss, U. Grandi, and D. Porello. Complexity of Judgment Aggregation: Safety of the Agenda. Proc. AAMAS-2010.

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.

Summary: Judgment Aggregation

This has been a brief introduction to judgment aggregation:

- Impossibility: anonymity, neutrality, independence \Rightarrow inconsistency
- Possibilities via domain restrictions: unidimensional alignment
- Safety of the agenda and complexity questions
- Relation to preference aggregation

Conclusion

Computational Social Choice

COMSOC research can be broadly classified along two dimensions —

The kind of *social choice problem* studied, e.g.:

- electing a winner given individual preferences over candidates
- aggregating individual judgements into a collective verdict
- fairly dividing a cake given individual tastes
- finding a stable matching of students to schools

The kind of *computational technique* employed, e.g.:

- algorithm design to implement complex mechanisms
- complexity theory to understand limitations
- logical modelling to fully formalise intuitions
- knowledge representation techniques to compactly model problems
- deployment in a multiagent system

Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. A Short Introduction to Computational Social Choice. Proc. SOFSEM-2007.

Last Slide

- Tried to give an overview of COMSOC, concentrating on aspects related to logic: axiomatic method, combinatorial domains, judgment aggregation. Of course, there's much more out there.
- Currently a very active area of research, with many opportunities.
- A website where you can find out more about Computational Social Choice (workshops, mailing list, PhD theses, etc.):

<http://www.illc.uva.nl/COMSOC/>

- These slides and my review paper will remain available on the tutorial website, and more extensive materials can be found on the website of my Amsterdam course on COMSOC:

– <http://www.illc.uva.nl/~ulle/teaching/kutaisi-2011/>

– <http://www.illc.uva.nl/~ulle/teaching/comsoc/>

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