

Applications of Logic in Social Choice Theory

Ulle Endriss

Institute for Logic, Language and Computation (ILLC)
University of Amsterdam

Abstract. Social choice theory studies of how groups of people should and do make collective decisions. In this talk I will argue that modern logic can contribute to the study of social choice theory in many different ways, and I will substantiate this claim with examples from recent work by members of my group at the University of Amsterdam.

1 Social Choice Theory

Social choice theory (SCT) is the formal study of mechanisms for collective decision making. As a scientific discipline, it is usually considered to be part of Economic Theory, although it also plays an important role in Political Science and Philosophy. In recent years, furthermore its fundamental significance for work in Multiagent Systems has become to be widely recognised.

The archetypal problem in the field is *preference aggregation*: given the preferences of a number of agents over a set of alternatives, how should we aggregate these individual preferences so as to arrive at a single collective preference order? To see that this is not a trivial question, consider the following example. There are three alternatives, called A , B and C , and five agents. The preferences of each agent are modelled as a linear order over the set of alternatives:

Agent 1:	$A \succ B \succ C$
Agent 2:	$B \succ C \succ A$
Agent 3:	$C \succ A \succ B$
Agent 4:	$C \succ A \succ B$
Agent 5:	$B \succ C \succ A$

The most obvious approach for obtaining a collective preference order is to use the *majority rule*: rank X above Y if and only if a majority of the agents do. If we follow this rule, then we must adopt $A \succ B$ (as three of the agents do), $B \succ C$ (as again three of the agents do), and $C \succ A$ (as four of the agents do). But this means that we get a collective preference order that is cyclic! This surprising outcome is an instance of the *Condorcet Paradox*, named after the 18th century political scientist and mathematician who first discussed it at length.

The question now arises whether there is a better aggregation rule than the majority rule, one that does not suffer from this paradox. Social choice theorists have approached this question using the so-called “axiomatic method”: by formulating desirable properties of aggregation rules as “axioms” in a mathematically rigorous manner, they have been able to obtain results that show that it

is impossible to find a rule that satisfies a certain combination of desirable properties or that a certain combination of such properties uniquely characterises a particular rule. Famous examples include Arrow’s Theorem (showing that there exists no preference aggregation rule for three or more alternatives that respects unanimous choices made by the individuals, that ranks pairs of alternatives independently from how other alternatives are ranked, and that is not dictatorial); the Gibbard-Satterthwaite Theorem (showing that there exists no voting rule for three or more alternatives that does not exclude an alternative from winning *a priori*, that does not allow for situations in which a voter can benefit from submitting a ballot that does not truthfully reflect her actual preferences, and that is not dictatorial); and May’s Theorem (showing that for two alternatives the majority rule is the only aggregation rule that treats all agents and alternatives symmetrically and that respects a basic monotonicity condition).

2 Applications of Logic

Logic has long played an important role in SCT: for instance, some properties of aggregation rules *entail* other properties, and impossibility theorems establish the *inconsistency* of certain sets of properties.

However, this use of “logic” is rather informal in nature. While it does refer to logical concepts such as “consistency”, it does not make use of a formal language. In the sequel I will argue that logic, including formal symbolic logic, has many more applications in SCT, and I will substantiate this claim with examples from recent work by members of my group at the University of Amsterdam.¹

2.1 Representation of Preferences

Before we can tackle the problem of aggregating preferences, we need to decide how to model the preferences themselves. In classical SCT, preferences are taken to be linear (or weak) orders over the set of alternatives, but other types of preference structures are also of interest (see e.g. [4]).

How to actually *represent* preferences, using a formal language, becomes critical when alternatives have a combinatorial structure, e.g., when agents are asked to express their preferences over all combinations of assigning truth values to, say, ten variables. (For an introduction to the field of preference modelling for *social choice in combinatorial domains*, see the expository article authored jointly with Chevaleyre et al. [1].) One family of languages proposed is based on *weighted goals*: agents describe their preferences in terms of propositional formulas they would like to see satisfied, together with numerical weights indicating their relative importance. Besides their application in SCT, also the study of the

¹ This is not intended to be a comprehensive review of the field; in particular, references are restricted to my own work. For extensive references to the use of logic in SCT, as well as to work in classical SCT and modern computational social choice more generally, please refer to the bibliographies of the cited papers.

properties of such languages themselves, e.g., their expressive power, has lead to interesting research questions [11, 12].

Much of this work requires only classical propositional logic, but nonclassical logic also has a role to play. For instance, for resource allocation problems where there may be multiple copies of the same type of resource available, *linear logic* turns out to be the right kind of formalism to represent preferences [9].

2.2 Characterisation and Impossibility Results

A natural application of logic in SCT is to attempt to fully formalise parts of the field. For instance, we have been able to give a full formalisation, in classical first-order logic, of the framework of preference aggregation introduced by Arrow and we have shown that Arrow’s Theorem corresponds to the claim that a particular set of first-order formulas does not have a finite model [6].

In another line of work we have shown that paradoxes of social choice, such as the Condorcet Paradox, can be explained in terms of the violation of an integrity constraint, expressed in propositional logic, that characterises the domain of aggregation [7, 8]. This approach provides a new way of characterising aggregation rules, namely in terms of the types of integrity constraints it respects, as well as a new proof technique for deriving old and new impossibility theorems in SCT.

2.3 Automated Reasoning in Social Choice Theory

One motivation for seeking to formalise SCT is that it opens up the possibility of using automated theorem provers to verify known results. For instance, a by-product of our work on the formalisation of the Arrovian framework of preference aggregation is a specification of that framework in the language of `Prover9`, the automated theorem prover formerly known as `Otter` [6]. Further optimisation may one day lead to a fully automated proof of Arrow’s Theorem.

In another area of SCT, known as *ranking sets of objects*, we already have been able to obtain results in a fully automated manner. Here the choice-theoretic problem is how to extend an agent’s preferences over individual objects to a preference order over sets of such objects. Building on a model-theoretic result, we have been able to use a satisfiability checker to not only verify known impossibility theorems but also to discover nontrivial new results [5].

2.4 Judgment Aggregation

Logic is not only useful for the analysis of aggregation problems, but information expressed in terms of logic may itself be subject to aggregation. This kind of problem is studied in the field of *judgment aggregation* (JA). Suppose three agents each have to judge which of the formulas φ , ψ and $\varphi \wedge \psi$ are true:

	φ	ψ	$\varphi \wedge \psi$
Agent 1:	true	true	true
Agent 2:	true	false	false
Agent 3:	false	true	false

Observe that each agent provides a logically consistent set of judgments. How should we aggregate this information? If we use the *majority rule* for each proposition, then φ should be collectively accepted, ψ should also be collectively accepted, and $\varphi \wedge \psi$ should be collectively rejected—that is, we obtain an inconsistent judgment set! Over the past decade or so, this paradox of JA has given rise to a fast moving area of research, spanning Legal Theory, Philosophy, Economic Theory and AI. For instance, we have recently begun to analyse the *computational complexity* of a number of problems that naturally arise in JA [2, 3].

While originally associated with problems in legal reasoning, it is not hard to see that JA can have a range of significant applications in other fields as well. One of them I believe to be the Semantic Web, and more specifically the aggregation of knowledge distributed over a number of different ontologies [10].

References

1. Chevaleyre, Y., Endriss, U., Lang, J., Maudet, N.: Preference handling in combinatorial domains: From AI to social choice. *AI Magazine* 29(4), 37–46 (2008)
2. Endriss, U., Grandi, U., Porello, D.: Complexity of judgment aggregation: Safety of the agenda. In: Proc. 9th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2010) (2010)
3. Endriss, U., Grandi, U., Porello, D.: Complexity of winner determination and strategic manipulation in judgment aggregation. In: Proc. 3rd International Workshop on Computational Social Choice (COMSOC-2010) (2010)
4. Endriss, U., Pini, M.S., Rossi, F., Venable, K.B.: Preference aggregation over restricted ballot languages: Sincerity and strategy-proofness. In: Proc. 21st International Joint Conference on Artificial Intelligence (IJCAI-2009) (2009)
5. Geist, C., Endriss, U.: Automated search for impossibility theorems in social choice theory: Ranking sets of objects. *Journal of Artificial Intelligence Research (JAIR)* 40, 143–174 (2011)
6. Grandi, U., Endriss, U.: First-order logic formalisation of Arrow’s theorem. In: Proc. 2nd International Workshop on Logic, Rationality and Interaction (LORI-2009). LNAI, vol. 5834, pp. 133–146. Springer-Verlag (2009)
7. Grandi, U., Endriss, U.: Lifting rationality assumptions in binary aggregation. In: Proc. 24th AAAI Conference on Artificial Intelligence (AAAI-2010) (2010)
8. Grandi, U., Endriss, U.: Binary aggregation with integrity constraints. In: Proc. 22nd International Joint Conference on Artificial Intelligence (IJCAI-2011) (2011)
9. Porello, D., Endriss, U.: Modelling combinatorial auctions in linear logic. In: Proc. 12th International Conference on the Principles of Knowledge Representation and Reasoning (KR-2010) (2010)
10. Porello, D., Endriss, U.: Ontology merging as social choice. In: Proc. 12th International Workshop on Computational Logic in Multiagent Systems (CLIMA-2011). Springer-Verlag (2011)
11. Uckelman, J., Chevaleyre, Y., Endriss, U., Lang, J.: Representing utility functions via weighted goals. *Mathematical Logic Quarterly* 55(4), 341–361 (2009)
12. Uckelman, J., Endriss, U.: Compactly representing utility functions using weighted goals and the max aggregator. *Artificial Intelligence* 174(15), 1222–1246 (2010)