UNIVERSITY OF AMSTERDAM

UvA-DARE (Digital Academic Repository)

Binary aggregation with integrity constraints

Grandi, U.

Publication date
2012

Link to publication

Citation for published version (APA):
Grandi, U. (2012). *Binary aggregation with integrity constraints*. [Thesis, fully internal, Universiteit van Amsterdam]. Institute for Logic, Language and Computation.

General rights
It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations
If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: https://uba.uva.nl/en/contact, or a letter to: Library of the University of Amsterdam, Secretariat, P.O. Box 19185, 1000 GD Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (https://dare.uva.nl)

Download date:15 Nov 2026

Chapter 8

Conclusions and Perspectives

If we have succeeded in our purpose, then the reader will agree that collective rationality should take centre stage in the study of aggregation problems. Inspired by the wide spectrum of potential applications in Artificial Intelligence, we have put forward a systematic study of collective rationality in binary aggregation, focusing on the syntactic structure of the integrity constraint that defines an aggregation problem. Moreover, we have shown that a failure in collective rationality is at the basis of most of the classical paradoxes in aggregation theory, and that the source of many (im)possibility results in the literature on preference and judgment aggregation lies in a clash between a set of axiomatic properties and requirements of collective rationality. By providing a unifying view on aggregation problems, our framework of binary aggregation with integrity constraints proved to be a useful and flexible tool for both the analysis of theoretical questions as well as for the development of solutions for application-oriented problems.

Let us now look back at what has been achieved in this dissertation. **Chapter 2** defined the framework of binary aggregation with integrity constraints, in which individuals make yes/no choices over a finite set of issues and an aggregation procedure merges them into a collective choice. The choices of individuals are bound by an integrity constraint or rationality assumption, which we represented as a formula in a simple propositional language. We provided a definition of collective rationality which features integrity constraints as a parameter: an aggregation procedure is collectively rational with respect to an integrity constraint if the collective outcome satisfies the constraint whenever all individual ballots do. We called a counterexample to collective rationality a paradox.

In **Chapter 3** we explored the generality of our definition of paradox by showing that many classical paradoxes from the literature on Social Choice Theory can be seen as instances of our definition. In particular, we focused on the Condorcet paradox, the discursive dilemma, the Ostrogorski paradox, and, to a lesser extent, the paradox of multiple elections.

In **Chapter 4** we defined the class $\mathcal{CR}[\mathcal{L}]$ as the class of collectively rational

procedures with respect to all integrity constraints in a given language $\mathcal{L}$. By studying those classes, we discovered that classical axiomatic properties from the literature on Social Choice Theory correspond to requirements of collective rationality with respect to some natural syntactically defined languages. For instance, the class of procedures that are collectively rational with respect to the language of literals corresponds to the class of unanimous procedures; the language of equivalences $\mathcal{L}_{\leftrightarrow}$ is associated with the requirement of issue-neutrality; and the XOR-language $\mathcal{L}_{XOR}$ is associated with the axiom of domain-neutrality. Similar correspondences cannot be proven for certain other axiomatic requirements. For the axioms of anonymity, independence, and two versions of monotonicity, we proved negative results which rule out possible characterisations in terms of collective rationality. We then moved to the study of procedures defined by means of acceptance quotas, a class that includes the majority rule. We characterised the set of integrity constraints lifted by the majority rule as the language of *2-clauses*, i.e., disjunctions of size at most 2, and we provided conditions about the size of a clause for other quota rules.

Classical frameworks like preference and judgment aggregation can be embedded into binary aggregation by devising suitable integrity constraints, and in Chapters 5 and 6 we compared theoretical results in these settings with our findings in binary aggregation.

Chapter 5 focused on preference aggregation. We obtained a possibility and an impossibility theorem for different representations of preferences by making use of our characterisation results from Chapter 4. We also put forward a new proof of Arrow's Theorem which aimed at reducing the impossibility to a clash between the Arrovian axiomatic requirements on the one hand, and collective rationality with respect to the preferential integrity constraint on the other.

In **Chapter 6** we presented in detail the framework of judgment aggregation. We focused on the problem of safety of the agenda, i.e., characterising the set of agendas on which a given class of procedures is guaranteed to output a consistent outcome. For several classes of procedures defined axiomatically, we identified the class of safe agendas by giving conditions on the structure of the inconsistent subsets that can be created from formulas in the agenda. For instance, we proved that a safe agenda for independent procedures only contains inconsistent subsets composed of a formula and its negation (we called this condition the syntactically simplified median property). We confronted these results with our findings in binary aggregation and we investigated the computational complexity of recognising safe agendas. For all the classes under consideration, we proved that the problem is Π_2^p-complete, i.e., highly intractable.

Chapter 7 is devoted to studying aggregation procedures that are designed to be collectively rational, and that are good candidates to be used in practice. We focus on three rules that can be employed in both binary and judgment aggregation: the average voter rule, the premise-based procedure and the distance-based procedure. For each rule, we investigated their axiomatic properties as well as the

computational complexity of the problems of winner determination and strategic manipulation (except for the distance-based rule, for which determining a winner is already too hard a problem to justify a study of manipulation). We concluded the chapter by comparing the computational complexity of some simple problems in binary and judgment aggregation.

There are numerous directions in which this work can be extended. **Voting theory** represents a closely connected topic in which many of the results presented in this dissertation may find application. First, as shown by our Example 2.1.6, approval voting, k-approval voting and the plurality rule can be seen as binary aggregation procedures over suitably defined domains. The use of constraints is not common in the study of voting procedures, but constitutes an interesting direction for future investigations, as testified by recent work on the topic (Lu and Boutilier, 2011). Second, our framework can easily generalise to account for the case of full (rather than binary) combinatorial domains, i.e., product spaces $\mathcal{D} = D_1 \times \cdots \times D_m$ with $|D_i| \geqslant 2$. In this more general setting, integrity constraints could be expressed in the propositional language $\mathcal{L}'_{PS}$ built on a set of atomic formulas $PS = \{x_j{=}a_j \mid j = 1, \ldots, m \text{ and } a_j \in D_j\}$ (cf. Example 2.1.7). Axioms for aggregation procedures on general combinatorial domains can be adapted from the literature (see, e.g., Lang, 2007). Preliminary results linking collective rationality with axiomatic properties can be obtained for the simpler case of voting for committees, in which the domains D_j are equal to a set of available candidates D and the combinatorial domain is the set $\mathcal{D} = D^m$. Moreover, a language to compactly express preferences over a combinatorial domain could be devised from $\mathcal{L}'_{PS}$ by adding a binary predicate $<$. A study of the expressivity and succintness of this language constitutes another interesting idea for future work, comparing it to existing languages for compact representation of preferences (Uckelman, 2009; Bienvenu et al., 2010).

Binary issues can also be used to represent **preferential dependencies**, for instance by encoding the graph underlying a given CP-net (Boutilier et al., 2004). Every edge in the graph could be associated with a binary issue, and integrity constraints could then be devised to encode rationality assumptions *about* preferential dependencies. Moreover, interesting aggregation procedures may be devised by first aggregating individual dependency structures, and then devising a sequential order of local elections based on the collective dependency graph obtained. A preliminary study along these lines has been carried out by Airiau et al. (2011). A similar approach has also been followed by Xia et al. (2008) and Conitzer et al. (2011).

The generality of the framework of binary aggregation with integrity constraints may also suggest that the **aggregation of logical structures** constitute an interesting theoretical setting for the study of collective agents. As models describing autonomous agents get more refined, our simple propositional language may not be sufficiently expressive or compact to encode more complex constraints.

The general framework for judgment aggregation developed by Dietrich (2007) goes in this direction. However, in this line of research the domain of aggregation is still composed of a sequence of yes/no choices over a set of (non-propositional) formulas, a framework that does not deviate significantly from binary aggregation. Instead, it would be interesting to study the aggregation of non-propositional logical *models* (e.g., first-order models, Kripke structures, etc.) using formulas in more complex logical languages to express properties of such structures, such as rationality assumptions. We have recently made an initial step in this direction by providing a first study of graph aggregation in which integrity constraints may be expressed in modal logic (Endriss and Grandi, 2012).

Our study of the **computational complexity** of aggregation procedures is based on a worst-case analysis. Despite constituting a solid first step in preparation for a possible implementation of a procedure, this approach has been criticised as failing to capture the actual distribution of the hard instances of a problem (Faliszewski and Procaccia, 2010; Walsh, 2011a). Thus, an important step to be taken in future work is a more refined analysis of those problems studied in the dissertation, by for instance using tools from average complexity (Faliszewski and Procaccia, 2010) or parametrized complexity (Betzler, 2010). Another interesting possibility is performing an evaluation of the distribution of the hard instances of some of our problems, an approach that may be useful to shed light on very hard questions such as our findings on the computational complexity of recognising safe agendas, as recently done in the domain of the manipulation of elections (Walsh, 2011b).

Perhaps the most intriguing direction for future research is to move **from theory to applications**. The flexibility of our framework has the potential of bringing interesting insights to the study and development of a diverse range of aggregation problems, like ranking and recommender systems, but it is in the design of mechanisms for complex collective decisions that lies its biggest potential. A recent position paper by Boutilier and Lu (2011) advocates a move from the study of high-stake and low-frequency situations like elections to low-stake and high-frequency situations such as those that are encountered in electronic commerce and similar applications. While the theoretical framework developed in this dissertation does not straightforwardly point in this direction, it is in its applicability and in its potential of suggesting novel solutions to real-world problems that it will find its main testing ground.