Binary Aggregation with Integrity Constraints

Umberto Grandi and Ulle Endriss

Institute for Logic, Language and Computation, University of Amsterdam
`u.grandi@uva.nl, ulle.endriss@uva.nl`

Abstract

We consider problems where several individuals each need to make a yes/no choice regarding a number of issues and these choices then need to be aggregated into a collective choice. We describe rationality assumptions in terms of a propositional formula and we explore the question of whether or not a given aggregation procedure will *lift* the rationality assumptions from the individual to the collective level. For various fragments of propositional logic, we provide an axiomatic characterisation of the class of aggregation procedures that will lift all rationality assumptions expressible in this fragment. We also show how several classical frameworks of Social Choice Theory, particularly preference and judgment aggregation, can be viewed as binary aggregation problems by designing suitable integrity constraints.

1 Introduction

In recent years, the AI community has dedicated more and more attention to the study of methods coming from Social Choice Theory (SCT). The reasons for this focus are clear: SCT provides tools for the analysis of collective choices of groups of agents, and as such is of immediate relevance to the study of multiagent systems. At the same time, studies in AI have led to a new and broadened perspective on classical results in SCT, e.g., via the use of knowledge representation languages for modelling preferences in social choice problems or via the complexity-theoretic analysis of the implementation of social choice rules. Particularly close to the interests of AI is the problem of social choice in combinatorial domains (Chevaleyre et al. 2008), where the space of choices the individuals have to make has a combinatorial structure.

Many of the questions studied in SCT arise from the observation of paradoxes, such as the Condorcet Paradox in preference aggregation (Gaertner 2006) or the Doctrinal Paradox in judgment aggregation (List and Puppe 2009). One of the scopes of this paper is to show how these can all be viewed as instances of a general definition of *paradox*, and to do so we translate classical frameworks for SCT into a canonical (and more easily implementable) one. This framework is *binary aggregation with integrity constraints*, introduced and studied in our previous works (Grandi and Endriss 2010; 2011), building on research initiated by Wilson (1975) and more recently developed by Dokow and Holzman (2010).

Dokow and Holzman (2010) characterise domains of aggregation over which every independent and unanimous procedure is dictatorial. This is a good example for the use of the axiomatic method in economic theory: the aim is to identify *the* appropriate set of axioms (e.g., to model real-word economies, specific moral ideals, etc.) and then to prove a characterisation (or impossibility) result for those axioms. AI suggests an alternative approach: with every new application the principles underlying a system may change; so we may be more interested in devising languages for expressing a range of different axioms rather than identifying the "right" set of axioms; and we may be more interested in developing methods that will help us to understand the dynamics of a range of different social choice scenarios rather than in technical results for a specific such scenario.

For this purpose we separate two parameters in the framework of binary aggregation. On the one hand, we introduce a propositional language to define the domain of aggregation by expressing a rationality assumption common to all individuals. On the other, we state a list of axioms to classify aggregation procedures over these domains. We call an aggregation procedure *collectively rational* with respect to a language if whenever all individuals submit ballots satisfying a formula in the language, so does the outcome of aggregation. We characterise, for several simple fragments of the language of propositional logic, the associated class of collectively rational procedures as the set of procedures satisfying a certain set of axioms.

We then turn to the study of classical frameworks of SCT as instances of binary aggregation with integrity constraints. We show how characterisation results proved in binary aggregation can be used to derive a new impossibility theorem in preference aggregation, a variant of Arrow's Theorem, by identifying a clash between the syntactic shape of the integrity constraints defining the framework of preference aggregation and a number of axiomatic postulates. In a similar fashion, we are able to translate problems in judgment aggregation into binary aggregation problems with a specific integrity constraint, and we identify a syntactic analogue of classical agenda properties guaranteeing consistent aggregation.

The paper is organised as follows: Section 2 presents the framework of binary aggregation with integrity constraints. In Section 3 we prove several characterisation results relating axiomatic requirements and collective rationality. Section 4 and Section 5, respectively, deal with the translation of preference and judgment aggregation to binary aggregation, and Section 6 concludes.

2 Binary Aggregation with Integrity Constraints

In this section we introduce the framework of binary aggregation with integrity constraint, based on work by Wilson (1975) and Dokow and Holzman (2010). We introduce two crucial definitions: a new definition of the notion of *paradox* and the definition of *collective rationality*. We conclude by stating classical axioms for aggregation procedures adapted to the framework of binary aggregation.

2.1 Terminology and Notation

Let $\mathcal{I} = \{1, \ldots, m\}$ be a finite set of *issues*, and let $\mathcal{D} = D_1 \times \cdots \times D_m$ be a boolean combinatorial *domain*, i.e., $|D_i| = 2$ for all $i \in \mathcal{I}$ (we assume $D_i = \{0, 1\}$). Let $PS = \{p_1, \ldots, p_m\}$ be a set of propositional symbols, one for each issue, and let $\mathcal{L}_{PS}$ be the corresponding propositional language. For any $\varphi \in \mathcal{L}_{PS}$, let $\mathrm{Mod}(\varphi)$ be the set of *models* that satisfy φ. For example, $\mathrm{Mod}(p_1 \wedge \neg p_2) = \{(1,0,0), (1,0,1)\}$ if $PS = \{p_1, p_2, p_3\}$. We call *integrity constraint* any formula $\mathrm{IC} \in \mathcal{L}_{PS}$. Any such formula defines a *domain of aggregation* $X := \mathrm{Mod}(\mathrm{IC})$.

Integrity constraints can be used to define what tuples in $\mathcal{D}$ we consider *rational* choices. For example, as we shall see in Section 4, $\mathcal{D}$ might be used to encode a binary relation representing preferences, in which case we may want to declare only those elements of $\mathcal{D}$ rational that correspond to relations that are transitive. We shall therefore use the terms "integrity constraints" and "rationality assumptions" interchangeably.

Let $\mathcal{N} = \{1, \ldots, n\}$ be a finite set of *individuals*. A *ballot* B is an element of $\mathcal{D}$ (i.e., an assignment to the variables $p_1, \ldots, p_m$); and a *rational ballot* B is an element of $\mathcal{D}$ that satisfies the integrity constraint, i.e., an element of Mod(IC). A *profile* $\mathbf{B}$ is a vector of (rational) ballots, one for each individual in $\mathcal{N}$. We write b_j for the jth element of a ballot B, and $b_{i,j}$ for the jth element of ballot B_i within a profile $\mathbf{B} = (B_1, \ldots, B_n)$. An *aggregation procedure* is a function $F : \mathcal{D}^{\mathcal{N}} \to \mathcal{D}$, mapping each profile to an element of $\mathcal{D}$. $F(\mathbf{B})_j$ denotes the result of the aggregation on issue j.

2.2 Paradoxes and Collective Rationality

Consider the following example: Let IC $= \neg(p_1 \wedge p_2 \wedge p_3)$ and suppose there are three individuals, choosing $(1, 1, 0)$, $(1, 0, 1)$ and $(0, 1, 1)$, respectively, i.e., their choices are rational (they all satisfy IC). If we use issue-wise majority (accepting p_i if a majority of individuals do) to aggregate their choices, however, we obtain $(1, 1, 1)$, which fails to be rational. This kind of observation is often referred to as a paradox.

We now give a general definition of paradoxical behaviour of an aggregation procedure in terms of the violation of certain rationality assumptions:

Definition 2.1. A paradox is a triple $(F, \mathbf{B}, \text{IC})$, where $F : \mathcal{D}^{\mathcal{N}} \to \mathcal{D}$ is an aggregation procedure, $\mathbf{B}$ is a profile in $\mathcal{D}^{\mathcal{N}}$, IC $\in \mathcal{L}_{PS}$, and $B_i \models$ IC for all $i \in \mathcal{N}$ but $F(\mathbf{B}) \not\models$ IC.

As we shall see in the following sections, various classical paradoxes in SCT are instances of this definition. A closely related notion is that of collective rationality:

Definition 2.2. Given an integrity constraint IC $\in \mathcal{L}_{PS}$, an aggregation procedure $F : \mathcal{D}^{\mathcal{N}} \to \mathcal{D}$ is called collectively rational (CR) for IC, if for all rational profiles $\mathbf{B} \in \text{Mod(IC)}^{\mathcal{N}}$ we have that $F(\mathbf{B}) \in$ Mod(IC).

Thus, F is CR if it can *lift* the rationality assumptions given by IC from the individual to the collective level. An aggregation procedure that is CR with respect to IC cannot generate a paradox with IC as integrity constraint.

2.3 Axiomatic Method

In social choice theory, aggregation procedures are studied using the axiomatic method. Axioms are used to express desirable properties of a procedure. In this section, we adapt the most important axioms familiar from standard social choice theory, and more specifically from judgment aggregation (List and Puppe 2009) and binary aggregation theory (Dokow and Holzman 2010), to our setting. We start with four common axioms:

Unanimity (U): For any profile $\mathbf{B} \in X^N$ and any $x \in \{0, 1\}$, if $b_{i,j} = x$ for all $i \in N$, then $F(\mathbf{B})_j = x$.

Anonymity (A): For any profile $\mathbf{B} \in X^{\mathcal{N}}$ and any permutation $\sigma : \mathcal{N} \to \mathcal{N}$, we have that $F(B_1, \ldots, B_n) = F(B_{\sigma(1)}, \ldots, B_{\sigma(n)})$.

Issue-Neutrality ($N^{\mathcal{I}}$): For any two issues $j, j' \in \mathcal{I}$ and any profile $\mathbf{B} \in X^{\mathcal{N}}$, if for all $i \in \mathcal{N}$ we have that $b_{i,j} = b_{i,j'}$, then $F(\mathbf{B})_j = F(\mathbf{B})_{j'}$.

Independence (I): For any issue $j \in \mathcal{I}$ and profiles $\mathbf{B}, \mathbf{B}' \in X^{\mathcal{N}}$, if $b_{i,j} = b'_{i,j}$ for all $i \in \mathcal{N}$, then $F(\mathbf{B})_j = F(\mathbf{B}')_j$.

Unanimity postulates that, if all individuals agree on issue j, then the aggregation procedure should implement that choice for j. Anonymity requires the procedure to be symmetric with respect to individuals. Issue-neutrality (a variant of the standard axiom of neutrality introduced in judgment aggregation) asks that the procedure be symmetric with respect to issues. Finally, independence requires the outcome of aggregation on a certain issue j to depend only on the individual choices regarding that issue. Combining independence with issue-neutrality, we get the axiom of systematicity (S) = (I) + ($N^{\mathcal{I}}$).

It is important to remark that all axioms are domain-dependent. For instance, many aggregation procedures, such as the majority rule, are independent over the full combinatorial domain $\mathcal{D}$, while others, such as the one presented in the next example, are not. With two issues, let IC = $(p_2 \to p_1)$ and let F be equal to the majority rule on the first issue, and accept the second issue only if the first one was accepted and the second one has the support of a majority of the individuals. This procedure is not independent on the full domain, but it is easy to see that it satisfies independence when restricted to $X^N = \text{Mod}(\text{IC})^N$.

As a generalisation of the axiom of neutrality introduced by May (1952), we introduce the following:

Domain-Neutrality ($N^{\mathcal{D}}$): For any two issues $j, j' \in \mathcal{I}$ and any profile $\mathbf{B} \in X^{\mathcal{N}}$, if $b_{i,j} = 1 - b_{i,j'}$ for all $i \in \mathcal{N}$, then $F(\mathbf{B})_j = 1 - F(\mathbf{B})_{j'}$.

The two notions of neutrality are uncorrelated but dual: issue-neutrality requires the outcome on two issues to be the same if all individuals agree on these issues; domain-neutrality requires it to be reversed if all the individuals make opposed choices on the two issues.

The following axiom of monotonicity is often called *positive responsiveness*, and is formulated as an (inter-profile) axiom for independent aggregation procedures:[1]

I-Monotonicity (M): For any issue $j \in \mathcal{I}$ and profiles $\mathbf{B}=(B_1..B_i..B_n)$ and $\mathbf{B}'=(B_1..B'_i..B_n)$ in $X^{\mathcal{N}}$, if $b_{i,j}=0$ and $b'_{i,j}=1$, then $F(\mathbf{B})_j = 1$ entails $F(\mathbf{B}')_j = 1$.

Every set of axioms identifies a class of aggregation procedures that satisfy these properties. A characterisation in mathematical terms can be obtained for some classes. One example is the class of *quota rules* $\mathcal{QR}$ introduced by Dietrich and List (2007): an aggregation procedure F for n individuals is a quota rule if for every issue j there exists a quota $0 \leqslant q_j \leqslant n+1$ such that, if we denote by $N^{\mathbf{B}}_j = |\{i \mid b_{i,j}=1\}|$, then $F(\mathbf{B})_j=1$ if and only if $N^{\mathbf{B}}_j \geqslant q_j$. The following representation result holds:

Proposition 1 (Dietrich and List, 2007)**.** *An aggregation procedure F satisfies* A, I, *and* M^{I} *if and only if it is a quota rule.*

A quota rule is called *uniform* if the quota is the same for all issues. By adding the axiom of issue-neutrality to Proposition 1 we get an axiomatisation of this class. The uniform quota rule with $q_j = \lceil \frac{n}{2} \rceil$ for all issues j is the *majority rule*. If n is odd, then the majority rule satisfies all of the axioms listed above—but, as we have seen, it is not CR even for simple integrity constraints such as $\neg(p_1 \wedge p_2 \wedge p_3)$. It is interesting to link these results with May's Theorem (1952) on the axiomatic characterisation of the majority rule in voting. We can obtain a more general version of his result (which deals with the case of a single issue) by adding the axiom of domain-neutrality: this forces the quota to treat $N^{\mathbf{B}}_j$ and $n - N^{\mathbf{B}}_j$ symmetrically, and thus the only possibility is to fix the quota as the majority of the individuals.

[1] A variant of this axiom for issue-neutral aggregators has been defined in previous work (Endriss et al. 2010).

3 Lifting Individual Rationality

We now want to establish connections between aggregation procedures characterised in terms of axioms and aggregation procedures characterised in terms of languages for integrity constraints for which they are collectively rational. To this end, we first define the class of procedures that can lift the integrity constraints belonging to a given language $\mathcal{L}$ (recall Definition 2.2).

Definition 3.1. For any language $\mathcal{L} \subseteq \mathcal{L}_{PS}$, define the class $\mathcal{CR}[\mathcal{L}]$ of aggregation procedures that lift $\mathcal{L}$:

$$\mathcal{CR}[\mathcal{L}] \quad = \quad \{F : \mathcal{D}^N \to \mathcal{D} \mid F \text{ is CR for all IC} \in \mathcal{L}\}$$

Next, we establish some basic properties of $\mathcal{CR}[\mathcal{L}]$. In our framework, we have made the assumption of IC being a single formula (rather than a set of formulas); we now provide a formal underpinning for this choice. For any $\mathcal{L} \subseteq \mathcal{L}_{PS}$, let $\mathcal{L}^{\wedge}$ be the language of conjunctions of formulas in $\mathcal{L}$.

Lemma 1. *$\mathcal{CR}[\mathcal{L}^{\wedge}] = \mathcal{CR}[\mathcal{L}]$ for all $\mathcal{L} \subseteq \mathcal{L}_{PS}$.*

Proof. $\mathcal{CR}[\mathcal{L}^{\wedge}]$ is clearly included in $\mathcal{CR}[\mathcal{L}]$, since $\mathcal{L} \subseteq \mathcal{L}^{\wedge}$. It remains to be shown that, if an aggregation procedure F lifts every constraint in $\mathcal{L}$, then it lifts any conjunction of formulas in $\mathcal{L}$. Let $\bigwedge_k \mathrm{IC}_k$ be a conjunction of formulas in $\mathcal{L}$, and let $\mathbf{B} \in \mathrm{Mod}(\bigwedge_k \mathrm{IC}_k)^N$ be a profile satisfying this integrity constraint. Since $\mathrm{Mod}(\bigwedge_k \mathrm{IC}_k) = \bigcap_k \mathrm{Mod}(\mathrm{IC}_k)$, we have that $\mathbf{B} \in \mathrm{Mod}(\mathrm{IC}_k)$ for every k. Thus, if $F \in \mathcal{CR}[\mathcal{L}]$, then $F(\mathbf{B}) \in \mathrm{Mod}(\mathrm{IC}_k)$ for every k. Therefore, F will also be in $\mathrm{Mod}(\bigwedge_k \mathrm{IC}_k)$, and this concludes the proof. □

In particular, we have that $\mathcal{CR}[cubes] = \mathcal{CR}[literals]$ and $\mathcal{CR}[clauses] = \mathcal{CR}[\mathcal{L}_{PS}]$. The latter holds, because for every propositional formula there is an equivalent formula in conjunctive normal form (CNF).

The following lemma is an immediate consequence of our definitions:

Lemma 2. *$\mathcal{CR}[\mathcal{L}_1 \cup \mathcal{L}_2] = \mathcal{CR}[\mathcal{L}_1] \cap \mathcal{CR}[\mathcal{L}_2]$ for all $\mathcal{L}_1, \mathcal{L}_2 \subseteq \mathcal{L}_{PS}$.*

Next we introduce notation for defining classes of aggregation procedures in terms of axioms. As mentioned earlier, a particular axiom may be satisfied on a subdomain of interest, but not on the full domain. Here, our subdomains of interests are subdomains that correspond to rational ballots for a given constraint. We therefore need to be able to speak about the procedures that satisfy an axiom on the subdomain $\mathrm{Mod}(\mathrm{IC})^N$ induced by a given integrity constraint IC.

Let $F_{\restriction \mathrm{Mod(IC)}^N}$ denote the restriction of the aggregation procedure F to the subdomain $\mathrm{Mod(IC)}^N$.

Definition 3.2. An aggregation procedure F satisfies a set of axioms AX wrt. a language $\mathcal{L} \subseteq \mathcal{L}_{PS}$, if for all constraints IC $\in \mathcal{L}$ the restriction $F_{\restriction \mathrm{Mod(IC)}^N}$ satisfies the axioms in AX. This defines the following class:

$$\mathcal{F}_{\mathcal{L}}[\mathrm{AX}] = \{F : \mathcal{D}^N \to \mathcal{D} \mid F_{\restriction \mathrm{Mod(IC)}^N} \text{ sat. AX for all IC} \in \mathcal{L}\}$$

We write $\mathcal{F}[\mathrm{AX}]$ as a shorthand for $\mathcal{F}_{\{\top\}}[\mathrm{AX}]$, the class of procedures that satisfy AX over the *full* domain $\mathcal{D}$. It is easy to see that the following lemma holds:

Lemma 3. $\mathcal{F}[\mathrm{AX}] \subseteq \mathcal{F}_{\mathcal{L}}[\mathrm{AX}]$ *for all* $\mathcal{L} \subseteq \mathcal{L}_{PS}$.

We shall now seek to obtain results that link the two kinds of classes defined, i.e., results of the form

$$\mathcal{CR}[\mathcal{L}] = \mathcal{F}_{\mathcal{L}}[\mathrm{AX}],$$

for certain languages $\mathcal{L}$ and certain axioms AX.

3.1 Characterisation Results

Our first characterisation result shows that the aggregation procedures that can lift all rationality constraints expressible in terms of a conjunction of literals (a cube) is precisely the class of unanimous procedures:

Proposition 2. $\mathcal{CR}[cubes] = \mathcal{F}_{cubes}[\mathrm{U}]$.

Proof. One direction is easy: If X is a domain defined by a cube, then every individual must agree on every literal in the conjunction, and, by unanimity, so will the collective. For the other direction, suppose that $F \in \mathcal{CR}[cubes]$. Fix $j \in \mathcal{I}$. Pick a profile $\mathbf{B} \in \mathcal{D}^n$ such that $b_{i,j} = 1$ (or 0) for all $i \in N$. That is, $\mathbf{B} \in \mathrm{Mod}(p_j)^N$ (or $\neg p_j$, respectively). Since F is collectively rational on every domain defined by a cube (and this includes literals), it must be the case that $F(\mathbf{B})_j = 1$ (or 0, respectively), proving the unanimity of the aggregator. □

Observe that, as $\mathcal{F}_{cubes}[\mathrm{U}] = \mathcal{F}[\mathrm{U}]$, the explicit mentioning of cubes on the righthand side of Proposition 2 is not needed; we chose this form of presentation for uniformity with later results on other axioms.[2] By Lemma 1, we also get $\mathcal{CR}[literals] = \mathcal{F}_{literals}[\mathrm{U}]$ (it is easy to see that $\mathcal{F}_{literals}[\mathrm{U}] = \mathcal{F}_{cubes}[\mathrm{U}]$).

[2] The same remark applies to Propositions 3 and 4 below.

Let $\mathcal{L}_{\leftrightarrow}$ be the language of bi-implications of positive literals: $\mathcal{L}_{\leftrightarrow} = \{p_j \leftrightarrow p_k \mid p_j, p_k \in PS\}$. This language allows us to characterise issue-neutral aggregators:

Proposition 3. $\mathcal{CR}[\mathcal{L}_{\leftrightarrow}] = \mathcal{F}_{\mathcal{L}_{\leftrightarrow}}[\mathrm{N}^{\mathcal{I}}]$.

Proof. To prove the first inclusion ($\supseteq$), pick a positive bi-implication $p_j \leftrightarrow p_k$: issues j and k share the same pattern of acceptances/rejections and since the procedure is neutral over issues, we get $F(\mathbf{B})_j = F(\mathbf{B})_k$. The constraint is therefore lifted. For the other direction ($\subseteq$), suppose that a profile $\mathbf{B}$ is such that $b_{i,j} = \mathbf{B}_{i,k}$ for every $i \in N$. Then $\mathbf{B} \in \mathrm{Mod}(p_j \leftrightarrow p_k)^N$, and if F is in $\mathcal{CR}[\mathcal{L}_{\leftrightarrow}]$, then $F(\mathbf{B})_j$ must be equal to $F(\mathbf{B})_k$. Since this holds for every such $\mathbf{B}$, this proves that F is neutral over issues. □

Let $\mathcal{L}_{\not\leftrightarrow}$ be the language of bi-implications of one negative and one positive literal: $\mathcal{L}_{\not\leftrightarrow} = \{p_j \leftrightarrow \neg p_k \mid p_j, p_k \in PS\}$. That is, $\mathcal{L}_{\not\leftrightarrow}$ is the language of XOR-formulas over pairs of positive literals. With a proof analogous to the one above we can characterise domain-neutrality:

Proposition 4. $\mathcal{CR}[\mathcal{L}_{\not\leftrightarrow}] = \mathcal{F}_{\mathcal{L}_{\not\leftrightarrow}}[\mathrm{N}^{\mathcal{D}}]$.

Let $\mathcal{F} = \{F : \mathcal{D}^N \to \mathcal{D}\}$ be the class of *all* aggregation procedures (for fixed $\mathcal{D}$ and N). The next result is an immediate consequence of our definitions:

Proposition 5. $\mathcal{CR}[\{\bot\}] = \mathcal{CR}[\{\top\}] = \mathcal{F}$.

Hence, by Lemma 2, $\mathcal{CR}[\mathcal{L} \cup \{\bot\}] = \mathcal{CR}[\mathcal{L}]$, which shows that unsatisfiable formulas can be omitted from languages for integrity constraints.

We now move on to characterising more specific classes of procedures. A *dictatorship* is an aggregation procedure that copies in every profile the ballot of a certain fixed individual, the dictator. The class $\mathcal{F}_{\mathcal{L}}[\mathrm{DIC}]$ is composed by all functions that are dictatorships when restricted to $\mathrm{Mod}(\mathrm{IC})^N$ for all $\mathrm{IC} \in \mathcal{L}$. Now, let us call a language $\mathcal{L} \subseteq \mathcal{L}_{PS}$ *trivial*, if it is composed only of formulas having a single model each. Clearly:

Proposition 6. *If $\mathcal{L}$ is trivial, then* $\mathcal{CR}[\mathcal{L}] = \mathcal{F}_{\mathcal{L}}[\mathrm{DIC}]$.

We propose the following definition of a class of aggregators that generalises the notion of dictatorship:

Definition 3.3. An aggregation procedure $F : \mathcal{D}^N \to \mathcal{D}$ is a generalised dictatorship, if there exists a map $g : \mathcal{D}^N \to N$ such that $F(\mathbf{B}) = \mathbf{B}_{g(\mathbf{B})}$ for every $\mathbf{B} \in \mathcal{D}^N$.

That is, a generalised dictatorship copies the ballot of a (possibly different) individual in every profile. Call this class $\mathcal{F}$[GDIC]. This class fully characterises the aggregators that can lift *any* integrity constraint:

Proposition 7. $\mathcal{CR}[\mathcal{L}_{PS}] = \mathcal{F}[\text{GDIC}]$.

Proof. Clearly, every generalised dictatorship lifts any arbitrary integrity constraint IC $\in \mathcal{L}_{PS}$. To prove the other direction, suppose that $F \notin \mathcal{F}[\text{GDIC}]$. Then there exists a profile $\mathbf{B} \in \mathcal{D}^N$ such that $F(\mathbf{B}) \neq \mathbf{B}_i$ for all $i \in N$. This means that for every i there exists an issue j_i such that $F(\mathbf{B})_{j_i} \neq \mathbf{B}_{i,j_i}$. Define now ℓ_{j_i} to be equal to p_{j_i} if $\mathbf{B}_{i,j_i} = 1$, to $\neg p_{j_i}$ otherwise. Call IC the following formula: $\bigvee_i \ell_{j_i}$. Clearly, $\mathbf{B}_i \models$ IC for every $i \in N$, i.e., $\mathbf{B}$ is a rational profile for the integrity constraint IC. Since $F(\mathbf{B}) \not\models$ IC by construction, F is not in $\mathcal{CR}[\{\text{IC}\}]$ and therefore also not in $\mathcal{CR}[\mathcal{L}_{PS}]$. □

All of the characterisation results presented thus far characterise a class of procedures determined by a *single* axiom (or apply to a very specific class of procedures) and by a *uniform* description of the language. So we might ask to what extent such results can be combined to allow us to make predictions regarding the collective rationality of procedures satisfying several such axioms, or in the case where the integrity constraints can be chosen from a more complex language. To illustrate the application of our results to such cases, suppose $\mathcal{CR}[\mathcal{L}_1] = \mathcal{F}_{\mathcal{L}_1}[\text{AX}_1]$ and $\mathcal{CR}[\mathcal{L}_2] = \mathcal{F}_{\mathcal{L}_2}[\text{AX}_2]$. Then Lemma 2 and the fact that $\mathcal{F}_{\mathcal{L}_1\cup\mathcal{L}_2}[\text{AX}_1,\text{AX}_2] \subseteq \mathcal{F}_{\mathcal{L}_1}[\text{AX}_2] \cap \mathcal{F}_{\mathcal{L}_2}[\text{AX}_2]$ entail $\mathcal{F}_{\mathcal{L}_1\cup\mathcal{L}_2}[\text{AX}_1,\text{AX}_2] \subseteq \mathcal{CR}[\mathcal{L}_1\cup\mathcal{L}_2]$. (But note that the other inclusion is not always true.) Now, if we start from the language $\mathcal{L}_1 \cup \mathcal{L}_2$ or any of its sublanguages, then this shows that picking procedures from $\mathcal{F}_{\mathcal{L}_1\cup\mathcal{L}_2}[\text{AX}_1,\text{AX}_2]$ is a sufficient condition for collective rationality. If, instead, we start from the axioms in AX_1 and AX_2, then we can infer that the procedures we obtain will lift any language $\mathcal{L} \subseteq \mathcal{L}_1 \cup \mathcal{L}_2$, since $\mathcal{F}[\text{AX}_1,\text{AX}_2] \subseteq \mathcal{F}_{\mathcal{L}_1\cup\mathcal{L}_2}[\text{AX}_1,\text{AX}_2] \subseteq \mathcal{CR}[\mathcal{L}_1 \cup \mathcal{L}_2] \subseteq \mathcal{CR}[\mathcal{L}]$. (The first inclusion follows from Lemma 3.)

3.2 Negative Results

For two important classes of aggregators, it is not possible to obtain a characterisation result:

Proposition 8. *There is no language* $\mathcal{L} \subseteq \mathcal{L}_{PS}$ *such that* $\mathcal{CR}[\mathcal{L}] = \mathcal{F}_{\mathcal{L}}[\text{I}]$.

Proof. We prove this proposition by constructing, for any choice of a language $\mathcal{L}$, an independent function that is not collectively rational for a certain IC $\in \mathcal{L}$.

Fix a language $\mathcal{L}$. W.l.o.g., this language will contain a falsifiable formula φ (otherwise $\mathcal{CR}[\mathcal{L}] = \mathcal{F}$ by Proposition 5 and we are done, as $\mathcal{F} \neq \mathcal{F}_{\mathcal{L}}[\text{I}]$). Choose a ballot/model $B^\star \in \mathcal{D}$ such that $B^\star \not\models \varphi$. Then the constant function $F \equiv B^\star$ is an independent function (on the full domain) that is not collectively rational. □

Proposition 9. *There is no language $\mathcal{L} \subseteq \mathcal{L}_{PS}$ such that $\mathcal{CR}[\mathcal{L}] = \mathcal{F}_{\mathcal{L}}[\text{A}]$.*

Proof. Employing a different technique than in the previous proof, we show that for every language $\mathcal{L}$ there exists a procedure that is collectively rational but not anonymous. First, in case $\mathcal{L}$ is trivial, by Proposition 6, $\mathcal{CR}[\mathcal{L}] = \mathcal{F}_{\mathcal{L}}[\text{DIC}]$, which is strictly included in the class of all anonymous functions. Second, if $\mathcal{L}$ is not trivial, then a dictatorship is always collectively rational (cf. Proposition 7), and it is not anonymous since due to nontriviality there is an IC $\in \mathcal{L}$ that allows for at least two different rational ballots. □

These results are coherent with the intuition that any assumption of collective rationality of an aggregator can only condition the outcome in view of a single profile at a time, without being able to express inter-profile requirements such as anonymity and independence. Similar remarks apply to the axiom of monotonicity (note that M^{I} is meaningful only in connection with I).

3.3 Quota Rules and Languages of Clauses

In view of the negative results proved above, we now focus on procedures satisfying anonymity, independence and monotonicity, and analyse the ability of procedures to lift rationality assumptions *within* that class. In previous work (Grandi and Endriss 2010) we proved several preliminary results for collective rationality of quota rules with respect to language of clauses. Here instead we present a recent result (Grandi and Endriss 2011) that gives precise bounds on quotas for languages of positive clauses.

Recall from Proposition 1 that the independent, anonymous and monotone procedures are exactly the quota rules, i.e., procedures that assign a quota q_j to every issue j such that $F(\mathbf{B})_j = 1 \Leftrightarrow |\{i \mid \mathbf{B}_{i,j} = 1\}| \geqslant q_j$. That is, in our notation, $\mathcal{QR} = \mathcal{F}[\text{A,I,M}^{\text{I}}]$. By Proposition 7 and Lemma 1, we know that $\mathcal{CR}[\mathit{clauses}]$ is the collection of generalised dictatorships. Therefore, to obtain results for more attractive classes of procedures, we restrict attention to clauses of limited length. For $k \geqslant 1$, let *k-clauses* be the set of clauses of length $\leqslant k$, *k-pclauses* the set of positive k-clauses, i.e., disjunctions where all literals are positive.

Proposition 10. *A quota rule is CR for a k-pclause* IC *if and only if* $\sum_j q_j < n + k$, *with j ranging over all issues that occur in* IC *and n being the number of individuals, or* $q_j = 0$ *for at least one issue j that occurs in* IC.

Proof. Suppose IC $= p_1 \vee \cdots \vee p_k$ and call $i_1, \ldots, i_k$ the corresponding issues. Given that IC is a positive clause, the only way to generate a paradox is by rejecting all issues $i_1, \ldots, i_k$. Suppose that we can create a paradoxical profile **B**. Suppose moreover that all quotas are > 0 (for otherwise one issue is always accepted and the IC trivially lifted). Every individual ballot B_i must accept at least one issue to satisfy the integrity constraint; therefore the profile **B** contains at least n acceptances. Since $F(\mathbf{B})_j = 0$ for all $j = 1, \ldots k$, we have that the number of individuals accepting an issue j is strictly lower than q_j. As previously remarked, there are at least n acceptances on the profile **B**; hence, this is possible if and only if $n \leqslant \sum_j (q_j - 1)$. Therefore, we can construct a paradox with this IC if and only if $n + k \leqslant \sum_j q_j$, and by taking the contrapositive we obtain the statement of Proposition 10. □

4 Preference Aggregation

In this section we give a translation of the framework of preference aggregation for linear orders into binary aggregation for a particular language of integrity constraints.

The framework of *preference aggregation* (see e.g. Gaertner 2006) considers a finite set of individuals $\mathcal{N}$ expressing preferences over a finite set of alternatives $\mathcal{X}$. A preference relation is represented by a binary relation P over $\mathcal{X}$. Here, we shall assume that P is a linear order, i.e., an antisymmetric, transitive and complete binary relation, thus reading aPb as "alternative a is strictly preferred to b". Let $\mathcal{L}(\mathcal{X})$ denote the set of all linear orders on $\mathcal{X}$. Aggregation procedures in this framework are functions $F : \mathcal{L}(\mathcal{X})^{\mathcal{N}} \to \mathcal{L}(X)$ and are called *social welfare functions* (SWFs).

4.1 Translation

Let us now consider the following setting for binary aggregation: define a set of issues $\mathcal{I}_{\mathcal{X}}$ as the set of all pairs (a, b) in $\mathcal{X}$. The domain $\mathcal{D}_{\mathcal{X}}$ of aggregation is therefore $\{0,1\}^{|\mathcal{X}|^2}$. In this setting a binary ballot corresponds to a binary relation P over $\mathcal{X}$: $B_{(a,b)} = 1$ iff a is in relation to b (aPb). Given this representation, we can associate with every SWF for $\mathcal{X}$ and $\mathcal{N}$ an aggregation procedure on a subdomain of $\mathcal{D}_{\mathcal{X}}^{\mathcal{N}}$.

Using the propositional language $\mathcal{L}_{PS}$, we can express properties of binary ballots in $\mathcal{D}_{\mathcal{X}}$. In this case the language consists of $|\mathcal{X}|^2$ propositional symbols, which we shall call p_{ab} for every issue (a, b). The properties of linear orders can be enforced on binary ballots using the following set of integrity constraints, which we shall call $\text{IC}_<$:[3]

Completeness and antisymmetry:
$p_{ab} \leftrightarrow \neg p_{ba}$ for $a \neq b \in \mathcal{X}$ $\quad \neg p_{aa}$ for all $a \in \mathcal{X}$

Transitivity: $p_{ab} \wedge p_{bc} \rightarrow p_{ac}$ for $a, b, c \in \mathcal{X}$ pairwise distinct

Note that the size of this set of integrity constraints is polynomial in the number of alternatives in $\mathcal{X}$.

It is now straightforward to see that every SWF corresponds to an aggregation procedure that is collectively rational wrt. $\text{IC}_<$, and *vice versa*. Moreover, if the SWF satisfies the unanimity axiom of preference aggregation (Gaertner 2006), then the associated binary aggregation procedure satisfies unanimity as defined in Section 2.3. The same is true for the axioms of anonymity, independence, and monotonicity (but note that for the two axioms of neutrality the correspondence is not straightforward).

4.2 Condorcet Paradox and Impossibilities

The translation presented above enables us to express the famous Condorcet Paradox in terms of Definition 2.1. Let $\mathcal{X} = \{a, b, c\}$ and let $\mathcal{N}$ contain three individuals. Consider the following profile **B**, where we have omitted the values of the reflexive issues aa (always 0 by $\text{IC}_<$), and specified the value of only one of ab and ba (the other can be obtained by taking the opposite of the value of the first):

	ab	*bc*	*ac*
Agent 1	1	1	1
Agent 2	0	1	0
Agent 3	1	0	0
Majority	1	1	0

[3]We will use the notation IC both for a single integrity constraint and for a set of formulas—in the latter case considering as the actual constraint the conjunction of all the formulas in IC.

Clearly, every individual ballot satisfies $\mathrm{IC}_<$, but the outcome obtained by taking majorities violates one formula, namely $p_{ab} \wedge p_{bc} \rightarrow p_{ac}$. Therefore, $(F_{\mathrm{maj}}, \mathbf{B}, \mathrm{IC}_<)$ is a paradox by Definition 2.1, where F_{maj} is the majority rule.

Now, by a syntactic analysis of the transitivity constraints introduced before, we can observe that they are in fact equivalent to just two positive clauses: The first one, $p_{ab} \vee p_{bc} \vee p_{ca}$, rules out the cycle $a{<}b{<}c{<}a$, and the second one, $p_{ba} \vee p_{cb} \vee p_{ac}$, rules out the opposite cycle $c{<}b{<}a{<}c$. That is, these constraints correspond exactly to the two Condorcet cycles that can be created from three alternatives.

We will now show how characterisation results of CR procedures for specific propositional languages, such as those given in Section 3, can be used to prove impossibility theorems in preference aggregation, similar to Arrow's Theorem (Arrow 1963). Call an SWF *imposed* if for some pair of distinct alternatives a and b we have that a is always collectively preferred to b in every profile.

Proposition 11. *If $|\mathcal{X}| \geqslant 3$ and $|\mathcal{N}| \geqslant 2$, then any anonymous, independent and monotonic SWF for $\mathcal{X}$ and $\mathcal{N}$ is imposed.*

Proof. In the first part of Section 4 we have seen that every anonymous, independent and monotonic SWF corresponds to a binary aggregation procedure that is collectively rational for $\mathrm{IC}_<$ and that satisfies A, I and M. By Proposition 1, every A, I, M aggregation procedure is a quota rule. We will now prove that, if a quota rule is collectively rational for $\mathrm{IC}_<$, then it is imposed, i.e., at least one of the quotas q_{ab} is equal to 0.

Suppose, for the sake of contradiction, that every quota $q_{ab} > 0$. As remarked before, for any three alternatives $a, b, c \in \mathcal{X}$ the integrity constraints corresponding to transitivity are $p_{ba} \vee p_{ca} \vee p_{bc}$ and $p_{ab} \vee p_{ac} \vee p_{cb}$. These are positive clauses of size 3; thus, by Proposition 10 we obtain:

$$q_{ab} + q_{bc} + q_{ca} < n + 3$$
$$q_{ba} + q_{cb} + q_{ac} < n + 3$$

Furthermore, it is easy to see that the $\mathrm{IC}_<$ for completeness and antisymmetry force the quotas to satisfy the following: $q_{ab}{+}q_{ba}{=}n{+}1$, $q_{bc}{+}q_{cb}{=}n{+}1$, and $q_{ac}{+}q_{ca}{=}n{+}1$.

Now, adding the two inequalities we obtain that $\sum_{a,b \in \mathcal{X}} q_{ab} < 2n + 6$ and adding the three equalities we obtain $\sum_{a,b \in \mathcal{X}} q_{ab} = 3n + 3$. The two constraints together admit a solution only if $n < 3$. Thus, it remains to analyse the case of 2 individuals; but it is easy to see that our constraints do not admit a solution in positive integers for $n = 2$. This shows that there must be a quota $q_{ab} = 0$ for certain distinct a and b as soon as $n \geqslant 2$; hence, the SWF is imposed. □

Arrow's Theorem states that every SWF satisfying U and I is dictatorial, and, although intuitively stronger, it does not imply Proposition 11. The importance of our result lies in the structure of its proof: most proofs of Arrow's Theorem and similar results concentrate on so-called "decisive coalitions". Here instead we point out a clash between axiomatic requirements and the syntactic shape of integrity constraints.

5 Judgment Aggregation

In this section we review the framework of *judgment aggregation* (List and Puppe 2009), and we provide a characterisation of judgment aggregation procedures as collectively rational procedures wrt. a particular set of integrity constraints.

Judgement aggregation (JA) considers problems where a finite set of individuals $\mathcal{N}$ has to generate a collective judgment over a set of interconnected propositional formulas Φ. Formally, we call *agenda* a finite nonempty set Φ of propositional formulas, not containing any doubly-negated formulas, that is closed under complementation (i.e, $\alpha \in \Phi$ whenever $\neg\alpha \in \Phi$, and $\neg\alpha \in \Phi$ for every positive $\alpha \in \Phi$). Each individual in $\mathcal{N}$ expresses a *judgment set* $J \subseteq \Phi$, as the set of those formulas in the agenda that she judges to be true. Every individual judgment set J is assumed to be *complete* (i.e., for each $\alpha \in \Phi$ either α or its complement are in J) and *consistent* (i.e., there exists an assignment that makes all formulas in J true). If we denote by $\mathcal{J}(\Phi)$ the set of all complete and consistent subsets of Φ, we can define a *JA procedure* for Φ and $\mathcal{N}$ as a function $F : \mathcal{J}(\Phi)^{\mathcal{N}} \to 2^{\Phi}$. A JA procedure is called *complete* (resp. *consistent*) if the judgment set it returns is complete (resp. consistent) on every profile.

5.1 Translation

Let us now consider the following binary aggregation framework. Let the set of issues $\mathcal{I}_\Phi$ be equal to the set of formulas in Φ. The domain $\mathcal{D}_\Phi$ of aggregation is therefore $\{0, 1\}^{|\Phi|}$. In this setting, a binary ballot corresponds to a judgment set: $B_\alpha = 1$ iff $\alpha \in J$. Given this representation, we can associate with every JA procedure for Φ and $\mathcal{N}$ a binary aggregation procedure on a subdomain of $\mathcal{D}_\Phi^{\mathcal{N}}$. Note that this translation is different from the one given by Dokow and Holzman (2010), which deals with models of judgment sets (rather than judgment sets) as input of the aggregation.

As before, we now define a set of integrity constraints for $\mathcal{D}_\Phi$ to enforce the properties of consistency and completeness. The propositional language in this case consists of $|\Phi|$ propositional symbols p_α, one for every $\alpha \in \Phi$. Recall that a *minimally inconsistent set* (mi-set) of propositional formulas is an inconsistent set each proper subset of which is consistent. Let IC_Φ be the following set of integrity constraints:

Completeness: $p_\alpha \vee p_{\neg\alpha}$ for all $\alpha \in \Phi$

Consistency: $\neg(\bigwedge_{\alpha \in S} p_\alpha)$ for every mi-set $S \subseteq \Phi$

Note that the size of IC_Φ might be exponential in the size of the agenda. This is in agreement with considerations of computational complexity: Since checking the consistency of a judgment set is an NP-hard problem, while model checking on binary ballots is in P, the translation from JA to binary aggregation must contain an exponential step.

The same kind of correspondence we have shown for SWFs holds between complete and consistent JA procedures and binary aggregation procedures that are collectively rational with respect to IC_Φ. We also obtain a perfect correspondence between the axioms, as every unanimous (resp. anonymous, independent, neutral, monotonic) JA procedure corresponds to a unanimous (resp. anonymous, independent, issue-neutral, monotonic) binary aggregation procedure.

5.2 Doctrinal Paradox and Agenda Properties

The paradox of JA was first studied in the literature discussing legal doctrines and then formalised in JA under the name of Doctrinal Paradox (List and Puppe 2009). Let Φ be the agenda $\{\alpha, \beta, \alpha \wedge \beta\}$[4] and let **B** be the following profile:

	α	β	$\alpha \wedge \beta$
Agent 1	1	1	1
Agent 2	0	1	0
Agent 3	1	0	0
Majority	1	1	0

Every individual ballot satisfies IC_Φ, while the outcome contradicts the constraint $\neg(p_\alpha \wedge p_\beta \wedge p_{\neg(\alpha\wedge\beta)}) \in IC_\Phi$. Hence, $(F_{\text{maj}}, \mathbf{B}, IC_\Phi)$ constitutes a paradox by Definition 2.1.

The notion of *safety of the agenda* introduced in previous work (Endriss et al. 2010) is related to our definition of paradox. An agenda Φ is *safe* wrt. a class of

[4] We omit negated formulas; for any $J \in \mathcal{J}(\Phi)$ their acceptance can be inferred from the acceptance of the positive counterparts.

JA procedures if any procedure in the class will return consistent outcomes for any profile over Φ. Several characterisation results have been proved that links agenda properties ensuring safety and classes of procedures defined axiomatically. As we shall see next, the translation of the JA framework into binary aggregation enables us to obtain a syntactic analogue of these properties. To simplify presentation, we shall assume that agendas do not include tautologies (or contradictions).

We say that an agenda Φ satisfies the *syntactic simplified median property* (SSMP) if every mi-subset of Φ is of the form $\{\alpha, \neg\alpha\}$. This corresponds to IC_Φ being equivalent to the conjunction of $p_\alpha \leftrightarrow \neg p_{\neg\alpha}$ for all positive $\alpha \in \Phi$. A weaker condition is the *simplified median property* (SMP), which holds if every mi-subset of Φ is of the form $\{\alpha, \neg\beta\}$ for α logically equivalent to β. Equivalences between formulas are expressed using bi-implications; thus, the SMP corresponds to adding to the previous set of constraints a set of positive bi-implications $p_\alpha \leftrightarrow p_\beta$ for any equivalent α and β in Φ. These considerations enable us to give a new proof for and strengthen a result that was proved in previous work (Endriss et al. 2010, Theorem 8). Call a procedure *complement-free* if the outcome never includes two formulas that are (syntactic) complements, for any profile in $\mathcal{J}(\Phi)^{\mathcal{N}}$.

Proposition 12. *An agenda Φ is safe for the class of complete, complement-free, and neutral JA procedures if and only if Φ satisfies the SMP.*

Proof. By translating JA into binary aggregation we have that Φ is safe wrt. complete, complement-free and neutral JA procedures if and only if IC_Φ does not generate a paradox with any issue-neutral procedure. It is easy to see that complete and complement-free procedures are characterised by procedures that are CR wrt. to constraints of the form $p_\alpha \leftrightarrow \neg p_{\neg\alpha}$. Therefore, we can concentrate on the remaining condition. We know by Proposition 3 that an issue-neutral procedure is collectively rational for IC_Φ iff IC_Φ is expressible in $\mathcal{L}_\leftrightarrow$, and using our earlier syntactic characterisation we conclude that this is the case iff Φ satisfies the SMP. □

The statement of Proposition 12 drops the axiom of anonymity, which was assumed in the previous statement of the theorem, and it does not require a representation result for its proof.

5.3 Another Characterisation Result: the Majority Rule

We will now glance back at the lifting results of Section 3, obtaining a characterisation of the set of integrity constraints that are lifted by the majority rule by exploiting the link between binary aggregation and JA. A result proved by Nehring and Puppe (2007) in the framework of JA shows that the majority rule is consistent if and only if the agenda Φ satisfies the *median property*, i.e., if there exists no mi-subset of Φ of size greater than 2. Binary aggregation problems with integrity constraints can be viewed as JA over atomic agendas: a ballot over issues $i_1, \dots, i_m$ can be viewed as a complete judgment set over a set of propositional symbols $p_1, \dots, p_m$, the consistency of a judgment set being defined as consistency *with respect to* the constraint IC. Ballots are assignments that may satisfy or falsify IC. Therefore, a mi-subset of the agenda corresponds to a *minimally falsifying partial assignment* (mifap-assignment) for IC: an assignment to some of the propositional variables that cannot be extended to a satisfying assignment, although each of its proper subsets can. Therefore, we obtain the following characterisation:

Lemma 4. *The majority rule is CR wrt. to* IC *if and only if there is no mifap-assignment for* IC *of size greater than 2.*

Let us now prove a crucial lemma about mifap-assignments. Associate with each mifap-assignment ρ a conjunction $C_\rho = \ell_1 \wedge \cdots \wedge \ell_k$, where ℓ_i=p_i if $\rho(p_i)$=1 and ℓ_i=$\neg p_i$ if $\rho(p_i)$=0, for all propositional symbols p_i on which ρ is defined.

Lemma 5. *Every non-tautological formula φ is equivalent to $(\bigwedge_\rho \neg C_\rho)$ with ρ ranging over all mifap-assignments of φ.*

Proof. Let A be a total assignment for φ. Suppose $A \not\models \varphi$, i.e., A is a falsifying assignment for φ. Since φ is not a tautology there exists at least one such A. By sequentially deleting propositional symbols from the domain of A we find a mifap-assignment ρ_A included in A. Hence, A falsifies the conjunct associated with ρ_A, and thus the whole formula $(\bigwedge_\rho \neg C_\rho)$.

Assume now $A \models \varphi$ but $A \not\models (\bigwedge_\rho \neg C_\rho)$. Then there is a ρ such that $A \models C_\rho$. This implies $\rho \subseteq A$, and since ρ is a mifap-assignment for φ this contradicts the assumption $A \models \varphi$. □

Proposition 13. *The majority rule is CR wrt.* IC *if and only if* IC *is expressible as a conjunction of clauses of size $\leqslant 2$.*

Proof. The proof for one direction can be found in previous work (Grandi and Endriss 2010, Proposition 18): the majority rule is CR wrt. conjunctions of 2-clauses. The other direction is entailed by the two lemmas above: Suppose that the majority rule is CR wrt. IC, then, by Lemma 4, IC does not have any mifap-assignment of size > 2. Therefore, by Lemma 5, we can construct a conjunction of 2-clauses that is equivalent to IC, as every conjunct C_ρ in the statement of Lemma 5 has size $\leqslant 2$. The case of IC being a tautology is straightforward, as every tautology is equivalent to a 2-clause, namely $p \vee \neg p$. □

6 Conclusions and Future Work

We introduced a simple propositional language to express individual rationality constraints in the framework of binary aggregation, and we defined an aggregation procedure to be collectively rational if the collective outcome satisfies a certain constraint whenever all individuals do. We proved several results to characterise, for various subsets of the language, a set of axioms that guarantees the collective rationality of a procedure for all constraints in this subset, and we have outlined an approach for how to apply these results in more complex situations. We have explored the potential of the framework of binary aggregation with integrity constraints as a general framework for the analysis of collective choice problems, by showing how two of the main frameworks of Social Choice Theory, preference and judgment aggregation, can be embedded into binary aggregation by defining suitable integrity constraints. We were able to give new and simpler proofs of theoretical results in both frameworks, and to characterise seemingly unrelated paradoxes as instances of the same general definition.

This work can be extended in a number of ways. The first step towards a generalisation to the case of full (rather than boolean) combinatorial domains (Lang 2007; 2004) is a study of the case of *voting for committees*, where the domain is a product space of domains D of equal cardinality. By defining a language from propositional symbols $\{p_{x_j=a} \mid a \in D, j \in \mathcal{I}\}$ it is possible to generate integrity constraints to model various voting procedures, such as approval voting, and prove preliminary results linking axioms with syntactic requirements on additional integrity constraints. Another direction is to allow for sequential aggregation procedures: by analysing the integrity constraints we might be able to devise a meaningful agenda for the decision process. Finally, by using more powerful languages to express rationality assumptions we can move towards more complex logical models of artificial agents.

Acknowledgements The authors would like to thank the anonymous referees of AAAI-2010 and IJCAI-2011 for their precious comments, as well as the audience of SCW-2010, the COMSOC seminar and the LIRa seminar at ILLC, and the Doctoral School on Computational Social Choice in Estoril, where different versions of this paper have been presented.

References

K. J. Arrow. *Social Choice and Individual Values*. John Wiley & Sons, 2nd edition, 1963.

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.

F. Dietrich and C. List. Judgment aggregation by quota rules: Majority voting generalized. *J. Theoret. Politics*, 19(4):391–424, 2007.

E. Dokow and R. Holzman. Aggregation of binary evaluations. *Journal of Economic Theory*, 145(2):495–511, 2010.

U. Endriss, U. Grandi, and D. Porello. Complexity of judgment aggregation: Safety of the agenda. In *Proceedings of AAMAS-2010*, 2010.

W. Gaertner. *A Primer in Social Choice Theory*. Oxford University Press, 2006.

U. Grandi and U. Endriss. Lifting rationality assumptions in binary aggregation. In *Proceedings of AAAI-2010*, 2010.

U. Grandi and U. Endriss. Binary aggregation with integrity constraints. In *Proceedings of IJCAI-2011*, 2011. To appear.

J. Lang. Logical preference representation and combinatorial vote. *Annals of Mathematics and Artificial Intelligence*, 42(1–3):37–71, 2004.

J. Lang. Vote and aggregation in combinatorial domains with structured preferences. In *Proceedings of IJCAI-2007*, 2007.

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

K. O. May. A set of independent necessary and sufficient conditions for simple majority decision. *Econometrica*, 20(4):680–684, 1952.

K. Nehring and C. Puppe. The structure of strategy-proof social choice. Part I: General characterization and possibility results on median spaces. *Journal of Economic Theory*, 135(1):269–305, 2007.

R. B. Wilson. On the theory of aggregation. *Journal of Economic Theory*, 10(1): 89–99, 1975.