

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

Binary aggregation studies problems in which individuals express yes/no choices over a number of possibly correlated issues, and these individual choices need to be aggregated into a collective choice. We show how several classical frameworks of Social Choice Theory, particularly preference and judgment aggregation, can be viewed as binary aggregation problems by designing an appropriate set of integrity constraints for each specific setting. We explore the generality of this framework, showing that it makes available useful techniques both to prove theoretical results, such as a new impossibility theorem in preference aggregation, and to analyse practical problems, such as the characterisation of safe agendas in judgment aggregation in a syntactic way. The framework also allows us to formulate a general definition of paradox that is independent of the domain under consideration, which gives rise to the study of the class of aggregation procedures of generalised dictatorships.

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 (Chevalleyre 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). In this paper, we 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*, which we introduced in earlier work (Grandi and Endriss, 2010), building on work initiated by Wilson (1975) and more recently developed by Dokow and Holzman (2010). In this setting a group of individuals have to aggregate their choices over a set of yes/no questions, and the range of answers that is considered *rational* is constrained by means of a propositional formula. In our 2010 paper we have shown how to characterise the class of collectively rational (i.e., paradox-free) aggregation procedures for several fragments of propositional logic in classical axiomatic terms.

In this paper, we prove two further such results and we employ them 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. Thus, we demonstrate that binary aggregation with integrity constraints constitutes a powerful general framework for the study of aggregation problems, including in particular preference and judgment aggregation. We also show how our approach can help us to identify attractive aggregation procedures of practical interest by studying the class of binary aggregation procedures that respect *all* integrity constraints expressible in the language of propositional logic. Specifically, we define a procedure we call *distance-based generalised dictatorship* and show that it enjoys good axiomatic properties.

The paper is organised as follows: Section 2 presents the framework of binary aggregation with integrity constraints, and some general results about collective rationality. Section 3 and Section 4, respectively, deal with the translation of preference and judgment aggregation to binary aggregation. Section 5 studies the class of generalised dictatorships, respecting all integrity constraints, and Section 6 concludes.

2 Binary Aggregation

In this section we review the framework of binary aggregation we defined in previous work (Grandi and Endriss, 2010),

which in turn is based on work by Wilson (1975) and Dokow and Holzman (2010). We also propose a new definition of the notion of *paradox* and we prove a new result characterising the class of aggregation procedures that are collectively rational wrt. any integrity constraint expressible in terms of positive clauses of bounded length.

2.1 Terminology and Notation

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

Integrity constraints can be used to define what tuples in \mathcal{D} we consider *rational* choices. For example, as we shall see in Section 3, \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 $\mathcal{N} = \{1, \dots, n\}$ be a finite set of *individuals*. A *ballot* B is an element of \mathcal{D} (i.e., an assignment to the variables p_1, \dots, p_m); and a *rational ballot* B is an element of \mathcal{D} that satisfies the integrity constraint, i.e., an element of $\text{Mod}(IC)$. A *profile* \mathbf{B} is a vector of (rational) ballots, one for each individual in \mathcal{N} . We write b_j for the j th element of a ballot B , and $b_{i,j}$ for the j th element of ballot B_i within a profile $\mathbf{B} = (B_1, \dots, B_n)$. An *aggregation procedure* is a function $F : \mathcal{D}^{\mathcal{N}} \rightarrow \mathcal{D}$, 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 1. A *paradox* is a triple (F, \mathbf{B}, IC) , where $F : \mathcal{D}^{\mathcal{N}} \rightarrow \mathcal{D}$ 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 sequel, various classical paradoxes in SCT are instances of this definition. A closely related notion is that of collective rationality (Grandi and Endriss, 2010):

Definition 2. Given an integrity constraint $IC \in \mathcal{L}_{PS}$, an aggregation procedure $F : \mathcal{D}^{\mathcal{N}} \rightarrow \mathcal{D}$ is called *collectively rational (CR)* for IC , if for all rational profiles $\mathbf{B} \in \text{Mod}(IC)^{\mathcal{N}}$ we have that $F(\mathbf{B}) \in \text{Mod}(IC)$.

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

Aggregation procedures are traditionally studied using the axiomatic method. Axioms are used to express desirable properties of a procedure. We now quickly review the most important axioms for binary aggregation procedures.

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

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

Issue-Neutrality ($\mathcal{N}^{\mathcal{I}}$): 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 $F(\mathbf{B})_j = F(\mathbf{B})_{j'}$.

Domain-Neutrality ($\mathcal{N}^{\mathcal{D}}$): 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 $F(\mathbf{B})_j = 1 - F(\mathbf{B})_{j'}$.

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 $F(\mathbf{B})_j = F(\mathbf{B}')_j$.

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

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 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. Independence requires the collective outcome on a certain issue j to depend only on the individual choices regarding j . I-monotonicity is a monotonicity axiom (designed for independent aggregators) asking that any collectively accepted issue receiving additional support will still be accepted.

Representation results can be proved for several sets of axioms, characterising in mathematical terms the class of aggregators satisfying those axioms. One example is the class of *quota rules* 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 $0 \leq q_j \leq n + 1$ such that, if we denote by $\mathcal{N}_j^{\mathbf{B}} = \{i | B_{i,j} = 1\}$, then $F(\mathbf{B})_j = 1$ if and only if $|\mathcal{N}_j^{\mathbf{B}}| \geq q_j$. A straightforward adaptation of a result by Dietrich and List (2007) gives the following representation result:

Proposition 1. An aggregation procedure F satisfies A, I, and M 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 $\neg(p_1 \wedge p_2 \wedge p_3)$.

2.4 Lifting Rationality Assumptions

Call an aggregation procedure collectively rational wrt. a certain sublanguage $\mathcal{L} \subseteq \mathcal{L}_{PS}$ if it is CR for every IC $\in \mathcal{L}$, and write $\mathcal{CR}[\mathcal{L}]$ for the class of all such procedures. In earlier work we have analysed the classes of CR procedures for different fragments of the language of propositional logic, characterising several of them in axiomatic terms (Grandi and Endriss, 2010). Here, we report one such result that we will use later on, and we prove a new result related to quota rules.

Let $\mathcal{F}_{\mathcal{L}_{\leftrightarrow}}[\mathbb{N}^{\mathcal{I}}]$ be the class of procedures that satisfy the axiom of issue-neutrality over domains defined by integrity constraints in $\mathcal{L}_{\leftrightarrow} := \{p_j \leftrightarrow p_k \mid p_j, p_k \in PS\}$. Then the following characterisation holds (Grandi and Endriss, 2010):

Proposition 2. $\mathcal{CR}[\mathcal{L}_{\leftrightarrow}] = \mathcal{F}_{\mathcal{L}_{\leftrightarrow}}[\mathbb{N}^{\mathcal{I}}]$.

A problem that was left open is the characterisation of CR procedures for languages of clauses. We provide here a precise result for the class of positive clauses. Let k -*pclauses* be the set of positive clauses of size $\leq k$.

Proposition 3. *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 $IC = p_1 \vee \dots \vee p_k$ and call i_1, \dots, 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, \dots, i_k . Suppose that we can create a paradoxical profile \mathbf{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 \mathbf{B} contains at least n acceptances. Since $F(\mathbf{B})_j = 0$ for all $j = 1, \dots, k$, 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 \mathbf{B} ; hence, this is possible if and only if $n \leq \sum_j (q_j - 1)$. Therefore, we can construct a paradox with this IC if and only if $n + k \leq \sum_j q_j$, and by taking the contrapositive we obtain the statement of Proposition 3. \square

We will prove another characterisation result in Section 4.3.

3 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 $F : \mathcal{L}(\mathcal{X})^{\mathcal{N}} \rightarrow \mathcal{L}(\mathcal{X})$ and are called *social welfare functions* (SWFs).

3.1 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 $IC_{<}$:¹

Completeness and antisymmetry:

$$p_{ab} \leftrightarrow \neg p_{ba} \text{ for } a \neq b \in \mathcal{X} \quad \neg p_{aa} \text{ for all } a \in \mathcal{X}$$

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. $IC_{<}$, 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).

3.2 Condorcet Paradox and Impossibilities

The translation presented above enables us to express the famous Condorcet Paradox in terms of Definition 1. Let $\mathcal{X} = \{a, b, c\}$ and let \mathcal{N} contain three individuals. Consider the following profile \mathbf{B} , where we have omitted the values of the reflexive issues aa (always 0 by $IC_{<}$), 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

Clearly, every individual ballot satisfies $IC_{<}$, but the outcome obtained by taking majorities violates one formula, namely $p_{ab} \wedge p_{bc} \rightarrow p_{ac}$. Therefore, $(F_{\text{maj}}, \mathbf{B}, IC_{<})$ is a paradox by Definition 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, $p_{ba} \vee p_{cb} \vee p_{ac}$, rules out the cycle $a < b < c < a$, and the second one, $p_{ab} \vee p_{bc} \vee p_{ca}$, 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.

¹We 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 .

We will now show how characterisation results of CR procedures for specific propositional languages, such as those given in Section 2.4, 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 4. *If $|\mathcal{X}| \geq 3$ and $|\mathcal{N}| \geq 2$, then any anonymous, independent and monotonic SWF for \mathcal{X} and \mathcal{N} is imposed.*

Proof. In the first part of Section 3 we have seen that every anonymous, independent and monotonic SWF corresponds to a binary aggregation procedure that is collectively rational for $IC_{<}$ 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 $IC_{<}$, 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 $p_{ba} \vee p_{cb} \vee p_{ac}$ and $p_{ab} \vee p_{bc} \vee p_{ca}$. These are positive clauses of size 3; thus, by Proposition 3 we obtain:

$$\begin{aligned} q_{ba} + q_{cb} + q_{ac} &< n + 3 \\ q_{ab} + q_{bc} + q_{ca} &< n + 3 \end{aligned}$$

Furthermore, it is easy to see that the $IC_{<}$ 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 $n \geq 2$; hence, the SWF is imposed. \square

Arrow’s Theorem states that every SWF satisfying U and I is dictatorial, and, although intuitively stronger, it does not imply Proposition 4. 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.

4 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 $F : \mathcal{J}(\Phi)^{\mathcal{N}} \rightarrow 2^{\Phi}$. A JA procedure is called *complete* (resp. *consistent*) if the judgment set it returns is complete (resp. consistent) on every profile.

4.1 Translation

Let us now consider the following binary aggregation framework. Let the set of issues \mathcal{I}_{Φ} be equal to the set of formulas in Φ . The domain \mathcal{D}_{Φ} 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}_{Φ} 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 NP-hard, while model checking on binary ballots is in P, the translation from JA to binary aggregation must contain a superpolynomial step (unless P=NP).

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.

4.2 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 $\{\alpha, \beta, \alpha \wedge \beta\}$ ² and let \mathbf{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

²We omit negated formulas; for any $J \in \mathcal{J}(\Phi)$ their acceptance can be inferred from the acceptance of the positive counterparts.

Every individual ballot satisfies IC_Φ , while the outcome contradicts the constraint $\neg(p_\alpha \wedge p_\beta \wedge p_{\neg(\alpha \wedge \beta)}) \in \text{IC}_\Phi$. Hence, $(F_{\text{maj}}, \mathbf{B}, \text{IC}_\Phi)$ constitutes a paradox by Definition 1.

The notion of *safety of the agenda* introduced by Endriss et al. (2010) is related to our definition of paradox. An agenda Φ is *safe* wrt. a class of JA procedures if any procedure in the class will return consistent outcomes for any profile over Φ . Endriss et al. (2010) prove several characterisation results linking 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).

Following Endriss et al. (2010), 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 of 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 $\mathcal{J}(\Phi)^N$.

Proposition 5. *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 2 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. \square

The statement of Proposition 5 drops the axiom of anonymity, which was assumed by Endriss et al. (2010, Theorem 8).

4.3 Median Property and Majority Rule

A problem that was left open in our previous work is the characterisation of the set of integrity constraints that is lifted by the majority rule (Grandi and Endriss, 2010). We will now settle this issue, exploiting the link to 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 6. *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 $C_\rho = \ell_1 \wedge \dots \wedge \ell_k$, where $\ell_i = p_i$ if $\rho(p_i) = 1$ and $\ell_i = \neg p_i$ if $\rho(p_i) = 0$, for all propositional symbols p_i on which ρ is defined.

Lemma 7. *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 $A \models \varphi$. \square

Proposition 8. *The majority rule is CR wrt. IC if and only if IC is expressible as a conjunction of clauses of size ≤ 2 .*

Proof. We proved one direction in earlier 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 6, IC does not have any mifap-assignment of size > 2 . Therefore, by Lemma 7, we can construct a conjunction of 2-clauses that is equivalent to IC, as every conjunct C_ρ in the statement of Lemma 7 has size ≤ 2 . The case of IC being a tautology is straightforward, as every tautology is equivalent to a 2-clause, namely $p \vee \neg p$. \square

5 Generalised Dictatorships

Any decision process that involves several binary issues (such as referenda, board decisions, ...) generates a binary aggregation problem. In the previous sections we have studied situations where an integrity constraint that is common to all individuals can be settled before the aggregation process takes place. In other situations this might not be possible, or the constraint might be subject to individual opinion. Therefore, we call the result of a referendum using aggregator R *contestable* if there exists a formula IC such that $(R, \mathbf{B}, \text{IC})$ is a paradox according to Definition 1. We shall now study the class of procedures that lift *any* integrity constraint, making the result of the aggregation incontestable.

Definition 3. *An aggregation procedure $F : \mathcal{D}^N \rightarrow \mathcal{D}$ is a generalised dictatorship, if there exists a map $g : \mathcal{D}^N \rightarrow \mathcal{N}$ 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. In our previous work we proved the following characterisation result (Grandi and Endriss, 2010):

Proposition 9. *An aggregator F is collectively rational wrt. every IC $\in \mathcal{L}_{PS}$ if and only if it is a generalised dictatorship.*

This class contains undesirable procedures such as (proper) dictatorships, where g is a constant function choosing the same individual in all profiles. But in general a generalised dictatorship is not an unreasonable rule: for instance, in large electorates it is often the case that the result of aggregation coincides with the ballot of a “median” voter, different in every situation. Moreover, it is easy to see that any generalised dictatorship satisfies the axiom of unanimity and both axioms of neutrality, $N^{\mathcal{I}}$ and $N^{\mathcal{D}}$.

We propose the following definition for a particular generalised dictatorship that outputs the ballot submitted by the “median” voter. It is a *non-resolute* procedure, i.e., an aggregation procedure that returns a set of ballots rather than a single one. Any non-resolute aggregation procedure can be made resolute by choosing a tie-breaking rule.

Definition 4. *We call distance-based generalised dictatorship the following (non-resolute) aggregation procedure:*

$$\text{DBGD}(\mathbf{B}) = \operatorname{argmin}_{\{B_i | i \in \mathcal{N}\}} \sum_{i' \in \mathcal{N}} H(B_i, B_{i'}),$$

where $H(B, B') = \sum_{j \in \mathcal{I}} |b_j - b'_j|$ is the Hamming distance.

This rule chooses as a representative individual one whose ballot is closest (in terms of the Hamming distance, i.e., the number of issues on which two individuals disagree) to the ballots of the others and adopts that individual’s choice as the collective choice. The DBGD rule has attractive properties:

Proposition 10. *The DBGD procedure satisfies U, $N^{\mathcal{I}}$, $N^{\mathcal{D}}$, and A for every choice of an anonymous tie-breaking rule.*

Proof. As a generalised dictatorship, DBGD satisfies U, $N^{\mathcal{I}}$ and $N^{\mathcal{D}}$ for any choice of a tie-breaking rule. Moreover, if tie-breaking is anonymous, then it clearly satisfies A. \square

The DBGD satisfies also a form of monotonicity adapted to non-resolute voting rules: if there exists a ballot B in the collective outcome such that $b_j = 1$ and we increase acceptance for issue j , then there still exists a winning ballot B' such that $b'_j = 1$. Indeed, if B is the ballot of one of the median voters in profile \mathbf{B} and $b_j = 1$, then increasing acceptance of j will only make B closer to the other individual ballots, thus B remains in the collective outcome.

Let us turn our analysis to the relation between the DBGD and the majority rule. Independence is clearly only satisfied by the latter. However, it is easy to see that the DBGD coincides with the majority rule whenever the result of the latter coincides with one of the individual ballots expressed in the profile. In the other cases it picks the individual ballot that is closest to the outcome of the majority rule. Therefore, the DBGD can be seen as a computationally tractable (winner determination is clearly polynomial) compromise between the good axiomatic properties of the majority rule and the full collective rationality of generalised dictatorships.

6 Conclusions and Future Work

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. We have shown 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. With similar techniques, we characterised the set of integrity constraints that are lifted by the majority rule. We have also analysed the more practical problem of designing a CR procedure for arbitrary integrity constraints, defining a distance-based procedure with very attractive properties.

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, 2004, 2007) 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.

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. *J. Economic Theory*, 145(2):495–511, 2010.
- U. Endriss, U. Grandi, and D. Porello. Complexity of judgment aggregation: Safety of the agenda. In *Proc. 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 *Proc. AAI-2010*, 2010.
- 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 *Proc. IJCAI-2007*, 2007.
- C. List and C. Puppe. Judgment aggregation: A survey. In *Handbook of Rational and Social Choice*. Oxford University Press, 2009.
- K. Nehring and C. Puppe. The structure of strategy-proof social choice. Part I: General characterization and possibility results on median spaces. *J. Economic Theory*, 135(1):269–305, 2007.
- R. B. Wilson. On the theory of aggregation. *J. Economic Theory*, 10(1):89–99, 1975.