

The Common Structure of Paradoxes in Aggregation Theory

Umberto Grandi

Abstract

In this paper we analyse some of the classical paradoxes in Social Choice Theory (namely, the Condorcet paradox, the discursive dilemma, the Ostrogorski paradox and the multiple election paradox) using a general framework for the study of aggregation problems called binary aggregation with integrity constraints. We provide a definition of paradox that is general enough to account for the four cases mentioned, and identify a common structure in the syntactic properties of the rationality assumptions that lie behind such paradoxes. We generalise this observation by providing a full characterisation of the set of rationality assumptions on which the majority rule *does not* generate a paradox.

1 Introduction

Most work in Social Choice Theory (SCT) started with the observation of paradoxical situations. From the Marquis de Condorcet (1785) to more recent American court cases (Kornhauser and Sager, 1986), a wide collection of paradoxes have been analysed and studied in the literature on Social Choice Theory (see, e.g., Nurmi, 1999). More recently, researchers in Artificial Intelligence (AI) have become interested in the study of collective choice problems in which the set of alternatives has a combinatorial structure (Lang, 2007; Xia et al., 2011). Novel paradoxical situations emerged from the study of these situations, and the combinatorial structure of the domains gave rise to interesting computational challenges.

This paper concentrates on the use of the majority rule on binary combinatorial domains, and investigates the question of what constitutes a paradox in such a setting. We identify a common structure behind the most classical paradoxes in SCT, putting forward a general definition of paradox in aggregation theory. By characterising paradoxical situations by means of computationally recognisable properties, we aim at providing more domain-specific research with new tools for the development of safe procedures for collective decision making.

We base the analysis on our previous work on binary aggregation with integrity constraints (Grandi and Endriss, 2011), which constitutes a general framework for the study of aggregation problems. In this setting, a set of individuals needs to take a decision over a set of binary issues, and these choices are then aggregated into a collective one. Given a rationality assumption that binds the choices of the individuals, we define a paradox as a situation in which all individuals are rational but the collective outcome is not. We present some of the most well-known paradoxes that arise from the use of the majority rule in different contexts, and we show how they can be expressed in binary aggregation as instances of this general definition. Our analysis focuses on the Condorcet paradox (1785), the discursive dilemma in judgment aggregation (List and Pettit, 2002), the Ostrogorski paradox (1902) and the more recent work of Brams et al. (1998) on multiple election paradoxes.

Such a uniform representation of the most important paradoxes in SCT enables us to make a crucial observation concerning the syntactic structure of the rationality assumptions that lie behind these paradoxes. We represent rationality assumptions by means of propositional formulas, and we observe that all formulas formalising a number of classical paradoxes feature a disjunction of literals of size at least 3. This observation can be generalised to a full characterisation of the rationality assumptions on which the majority rule does not generate

a paradox, and in Theorem 4 we identify them as those formulas that are equivalent to a conjunction of clauses of size at most 2.

The paper is organised as follows. In Section 2 we give the basic definitions of the framework of binary aggregation with integrity constraints, and we provide a general definition of paradox. In Section 3 we show how a number of paradoxical situations in SCT can be seen as instances of our general definition of paradox, and we identify a syntactic property that is common to all paradoxical rationality assumptions. Section 4 provides a characterisation of the paradoxical situations for the majority rule and Section 5 concludes the paper.

2 Binary Aggregation with Integrity Constraints

In this section we provide the basic definitions of the framework of binary aggregation with integrity constraints (Grandi and Endriss, 2011), based on work by Wilson (1975) and Dokow and Holzman (2010). In this setting, a number 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. Paradoxical situations may occur when a set of individual choices that is considered rational leads to a collective outcome which fails to satisfy the same rationality assumption of the individuals.

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 $i \in \mathcal{I}$. Without loss of generality we assume that $D_j = \{0, 1\}$ for all j . Thus, given a set of issues \mathcal{I} , the domain associated with it is $\mathcal{D} = \{0, 1\}^{\mathcal{I}}$. A *ballot* B is an element of \mathcal{D} .

In many applications it is necessary to specify which elements of the domain are rational and which should not be taken into consideration. Propositional logic provides a suitable formal language to express possible restrictions of rationality on binary combinatorial domains. If \mathcal{I} is a set of m issues, let $PS = \{p_1, \dots, p_m\}$ be a set of propositional symbols, one for each issue, and let \mathcal{L}_{PS} be the propositional language constructed by closing PS under propositional connectives. For any formula $\varphi \in \mathcal{L}_{PS}$, let $\text{Mod}(\varphi)$ be the set of assignments that satisfy φ . For example, $\text{Mod}(p_1 \wedge \neg p_2) = \{(1, 0, 0), (1, 0, 1)\}$ when $PS = \{p_1, p_2, p_3\}$. An *integrity constraint* is any formula $IC \in \mathcal{L}_{PS}$.

Integrity constraints can be used to define what tuples in \mathcal{D} we consider *rational* choices. Any ballot $B \in \mathcal{D}$ is an assignment to the variables p_1, \dots, p_m , and we call B a *rational ballot* if it satisfies the integrity constraint IC , i.e., if B is an element of $\text{Mod}(IC)$. In the sequel we shall use the terms “integrity constraints” and “rationality assumptions” interchangeably.

Let $\mathcal{N} = \{1, \dots, n\}$ be a finite set of *individuals*. We make the assumption that there are at least 2 individuals. Each individual submits a ballot $B_i \in \mathcal{D}$ to form a *profile* $\mathbf{B} = (B_1, \dots, B_n)$. 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 $\mathbf{B} = (B_1, \dots, B_n)$. Given a finite set of issues \mathcal{I} and a finite set of individuals \mathcal{N} , an *aggregation procedure* is a function $F : \mathcal{D}^{\mathcal{N}} \rightarrow \mathcal{D}$, mapping each profile of binary ballots to an element of \mathcal{D} . Let $F(\mathbf{B})_j$ denote the result of the aggregation of profile \mathbf{B} on issue j .

2.2 A General Definition of Paradox

Consider the following example: Let $IC = p_1 \wedge p_2 \rightarrow p_3$ and suppose there are three individuals, choosing ballots $(0, 1, 0)$, $(1, 0, 0)$ and $(1, 1, 1)$. Their choices are rational (they all satisfy IC). However, if we employ the *majority rule*, i.e., we accept an issue j if and only if a majority of individuals do, we obtain the ballot $(1, 1, 0)$ as collective outcome, 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 $(F, \mathbf{B}, \text{IC})$, where F is an aggregation procedure, \mathbf{B} is a profile in $\mathcal{D}^{\mathcal{N}}$, IC is an integrity constraint in \mathcal{L}_{PS} , and $B_i \in \text{Mod}(\text{IC})$ for all $i \in \mathcal{N}$ but $F(\mathbf{B}) \notin \text{Mod}(\text{IC})$.

A closely related notion is that of collective rationality:

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

Thus, F is CR with respect to IC if it *lifts* the rationality assumption given by IC from the individual to the collective level, i.e., if $F(\mathbf{B}) \in \text{Mod}(\text{IC})$ whenever $B_i \in \text{Mod}(\text{IC})$ for all $i \in \mathcal{N}$. An aggregation procedure that is CR with respect to IC cannot generate a paradoxical situation with IC as integrity constraint. Given an aggregation procedure F , let $\mathcal{LF}[F] = \{\varphi \in \mathcal{L}_{PS} \mid F \text{ is CR with respect to } \varphi\}$ be the set of integrity constraints that are lifted by F .

3 Unifying Paradoxes in Binary Aggregation

In this section we present a number of classical paradoxes from SCT, and we show how they can be seen as instances of our Definition 1. In Section 3.1 we introduce the Condorcet paradox, and we show how settings of preference aggregation can be seen as instances of binary aggregation by devising a suitable integrity constraint. Section 3.2 repeats this construction for the framework of judgment aggregation and for the discursive dilemma. In Section 3.3 we then deal with the Ostrogorski paradox, in which a paradoxical feature of representative majoritarian systems is analysed, and in Section 3.4 we focus on the paradox of multiple elections. In Section 3.5 we conclude by identifying a common structure in the integrity constraints that lie behind those paradoxes.

3.1 The Condorcet Paradox and Preference Aggregation

One of the earliest observation of paradoxical behaviour of the majority rule was made by the Marquis de Condorcet in 1785. A simple version of the paradox he discovered is explained in the following paragraphs:

Condorcet Paradox. Three individuals need to decide on the ranking of three alternatives $\{\Delta, \circ, \square\}$. Each individual expresses her own ranking and the collective outcome is aggregated by pairwise majority: an alternative is preferred to a second one if and only if a majority of the individuals prefer the first alternative to the second. Consider the following situation:

$$\begin{array}{c} \Delta <_1 \circ <_1 \square \\ \square <_2 \Delta <_2 \circ \\ \circ <_3 \square <_3 \Delta \\ \hline \Delta < \circ < \square < \Delta \end{array}$$

When computing the outcome of the pairwise majority rule, we notice that there is a majority of individuals preferring the circle to the triangle ($\Delta < \circ$); that there is a majority of individuals preferring the square to the circle ($\circ < \square$); and, finally, that there is a majority of individuals preferring the triangle to the square ($\square < \Delta$). The resulting outcome fails to be a linear order, giving rise to a circular collective preference between the alternatives.

3.1.1 Preference Aggregation

Condorcet’s paradox was rediscovered in the second half of the XXth century while a whole theory of *preference aggregation* was being developed (see, e.g., Gaertner, 2006). This framework 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 over \mathcal{X} . Preference relations are traditionally assumed to be *weak orders*, i.e., reflexive, transitive and complete binary relations. Another common assumption is representing preferences as *linear orders*, i.e., irreflexive, transitive and complete binary relations. In the sequel we shall assume that preferences are represented as linear orders, writing aPb for “alternative a is strictly preferred to b ”. Each individual in \mathcal{N} submits a linear order P_i , forming a profile $\mathbf{P} = (P_1, \dots, P_{|\mathcal{N}|})$. Let $\mathcal{L}(\mathcal{X})$ denote the set of all linear orders on \mathcal{X} . Given a finite set of individuals \mathcal{N} and a finite set of alternatives \mathcal{X} , a *social welfare function* is a function $F : \mathcal{L}(\mathcal{X})^{\mathcal{N}} \rightarrow \mathcal{L}(\mathcal{X})$.

3.1.2 Translation

Given a preference aggregation problem defined by a set of individuals \mathcal{N} and a set of alternatives \mathcal{X} , let us 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 $\{0, 1\}^{|\mathcal{I}_{\mathcal{X}}|}$. In this setting, a binary ballot B corresponds to a binary relation P over \mathcal{X} : $B_{(a,b)} = 1$ if and only if 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 that is defined on a subdomain of $\mathcal{D}_{\mathcal{X}}^{\mathcal{N}}$. We now characterise this domain as the set of models of a suitable integrity constraint.

Using the propositional language \mathcal{L}_{PS} constructed over the set $\mathcal{I}_{\mathcal{X}}$, we can express properties of binary ballots in $\mathcal{D}_{\mathcal{X}}$. In this case \mathcal{L}_{PS} consists of $|\mathcal{X}|^2$ propositional symbols, which we 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 $\text{IC}_{<}$:¹

- Irreflexivity:** $\neg p_{aa}$ for all $a \in \mathcal{X}$
- Completeness:** $p_{ab} \vee p_{ba}$ for all $a \neq b \in \mathcal{X}$
- 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} . In case preferences are expressed using weak orders rather than linear orders, it is sufficient to replace the integrity constraints of irreflexivity in $\text{IC}_{<}$ with their negation to obtain a similar correspondence between SWFs and aggregation procedures.

3.1.3 The Condorcet Paradox in Binary Aggregation

The translation presented in the previous section enables us to express the Condorcet paradox in terms of Definition 1. Let $\mathcal{X} = \{\Delta, \circ, \square\}$ and let \mathcal{N} contain three individuals. Consider the profile \mathbf{B} for $\mathcal{I}_{\mathcal{X}}$ in the following table, where we have omitted the values of the reflexive issues (Δ, Δ) (always 0 by $\text{IC}_{<}$), and specified the value of only one of (Δ, \circ) and (\circ, Δ) (the other can be obtained by taking the opposite of the value of the first), and accordingly for the other alternatives.

	$\Delta\circ$	$\circ\square$	$\Delta\square$
Agent 1	1	1	1
Agent 2	1	0	0
Agent 3	0	1	0
<i>Maj</i>	1	1	0

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 .

Every individual ballot satisfies $IC_{<}$, but the outcome obtained using the majority rule Maj (which corresponds to pairwise majority in preference aggregation) does not satisfy $IC_{<}$: the formula $p_{\Delta\circ} \wedge p_{\circ\Delta} \rightarrow p_{\Delta\Delta}$ is falsified by the outcome. Therefore, $(Maj, \mathcal{B}, IC_{<})$ is a paradox by Definition 1.

3.2 The Discursive Dilemma and Judgment Aggregation

The discursive dilemma emerged from the formal study of court cases that was carried out in recent years in the literature on law and economics, generalising the observation of a paradoxical situation known as the “doctrinal paradox” (Kornhauser and Sager, 1986). Such a setting was first given mathematical treatment by List and Pettit (2002), giving rise to an entirely new research area in SCT known as *judgment aggregation*.

Discursive Dilemma. A court of three judges has to decide on the liability of a defendant under the charge of breach of contract. An individual is considered liable if there was a valid contract and her behaviour was such as to be considered a breach of the contract. The court takes three majority decisions on the following issues: there was a valid contract (α), the individual broke the contract (β), the defendant is liable ($\alpha \wedge \beta$). Consider the following situation:

	α	β	$\alpha \wedge \beta$
Judge 1	yes	yes	yes
Judge 2	no	yes	no
Judge 3	yes	no	no
Majority	yes	yes	no

All judges express consistent judgments: they accept the third proposition if and only if the first two are accepted. However, even if there is a majority of judges who believe that there was a valid contract, and even if there is a majority of judges who believe that the individual broke the contract, the individual is considered *not liable* by a majority of the individuals.

3.2.1 Judgment Aggregation

Judgement aggregation (JA) considers problems in which a finite set of individuals \mathcal{N} has to generate a collective judgment over a set of interconnected propositional formulas (see, e.g., List and Puppe, 2009). Formally, given a propositional language \mathcal{L} , an *agenda* is a finite nonempty subset $\Phi \subseteq \mathcal{L}$ that does not contain doubly-negated formulas and is closed under complementation (i.e., $\alpha \in \Phi$ whenever $\neg\alpha \in \Phi$, and $\neg\alpha \in \Phi$ for non-negated $\alpha \in \Phi$).

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). Denote by $\mathcal{J}(\Phi)$ the set of all complete and consistent subsets of Φ . Given a finite agenda Φ and a finite set of individuals \mathcal{N} , a *JA procedure* for Φ and \mathcal{N} is a function $F : \mathcal{J}(\Phi)^{\mathcal{N}} \rightarrow 2^{\Phi}$.

3.2.2 Translation

Given a JA framework defined by an agenda Φ and a set of individuals \mathcal{N} , let us now construct a setting for binary aggregation with integrity constraints that interprets it. 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 B corresponds to a judgment set: $B_{\alpha} = 1$ if and only if $\alpha \in J$. Given this representation, we can associate with every JA procedure for Φ and \mathcal{N} a binary aggregation procedure on a subdomain of $\mathcal{D}_{\Phi}^{\mathcal{N}}$.

As we did for the case of preference aggregation, we now define a set of integrity constraints for \mathcal{D}_Φ to enforce the properties of consistency and completeness of individual judgment sets. Recall that the propositional language is constructed in this case on $|\Phi|$ propositional symbols p_α , one for every $\alpha \in \Phi$. Call an inconsistent set of formulas each proper subset of which is consistent *minimally inconsistent set* (mi-set). 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 $S \subseteq \Phi$

While the interpretation of the first formula is straightforward, we provide some further explanation for the second one. If a judgment set J is inconsistent, then it contains a minimally inconsistent set, obtained by sequentially deleting one formula at the time from J until it becomes consistent. This implies that the constraint previously introduced is falsified by the binary ballot that represents J , as all issues associated with formulas in a mi-set are accepted. *Vice versa*, if all formulas in a mi-set are accepted by a given binary ballot, then clearly the judgment set associated with it is inconsistent.

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 polynomial, the translation from JA to binary aggregation must contain a superpolynomial step (unless $\text{P}=\text{NP}$).

3.2.3 The Discursive Dilemma in Binary Aggregation

The same procedure that we have used to show that the Condorcet paradox is an instance of our general definition of paradox applies here for the case of the discursive dilemma. Let Φ be the agenda $\{\alpha, \beta, \alpha \wedge \beta\}$, in which we have omitted negated formulas, as for any $J \in \mathcal{J}(\Phi)$ their acceptance can be inferred from the acceptance of their positive counterparts. Consider the profile \mathbf{B} for \mathcal{L}_Φ described in the following profile:

	α	β	$\alpha \wedge \beta$
Judge 1	1	1	1
Judge 2	0	1	0
Judge 3	1	0	0
<i>Maj</i>	1	1	0

Every individual ballot satisfies IC_Φ , while the outcome obtained by using the majority rule contradicts one of the constraints of consistency, namely $\neg(p_\alpha \wedge p_\beta \wedge p_{\neg(\alpha \wedge \beta)})$. Hence, $(\text{Maj}, \mathbf{B}, \text{IC}_\Phi)$ constitutes a paradox by Definition 1.

3.3 The Ostrogorski Paradox

A less well-known paradox concerning the use of the majority rule on multiple issues is the Ostrogorski paradox (Ostrogorski, 1902).

Ostrogorski Paradox. Consider the following situation: there is a two party contest between the Mountain Party (MP) and the Plain Party (PP); three individuals (or, equivalently, three equally big groups in an electorate) will vote for one of the two parties if their view agrees with that party on a majority of the three following issues: economic policy (E), social policy (S), and foreign affairs policy (F). Consider the situation described in Table 1. The result of the two party contest, assuming that the party that has the support of a majority of the voters wins, declares the Plain Party the winner. However, a majority of individuals support the Mountain Party both on the economic policy E and on

	<i>E</i>	<i>S</i>	<i>F</i>	Party supported
Voter 1	MP	PP	PP	PP
Voter 2	PP	PP	MP	PP
Voter 3	MP	PP	MP	MP
<i>Maj</i>	MP	PP	MP	PP

Table 1: The Ostrogorski paradox.

the foreign policy *F*. Thus, the elected party (the PP) is in disagreement with a majority of the individuals on a majority of the issues.

3.3.1 The Ostrogorski Paradox in Binary Aggregation

In this section, we provide a binary aggregation setting that represents the Ostrogorski paradox as a failure of collective rationality with respect to a suitable integrity constraint.

Let $\{E, S, F\}$ be the set of issues at stake, and let the set of issues $\mathcal{I}_O = \{E, S, F, A\}$ consist of the same issues plus an extra issue *A* to encode the support for the first party (MP). A binary ballot over these issues represents the individual view on the three issues *E*, *S* and *F*: if, for instance, $b_E = 1$, then the individual supports the first party MP on the first issue *E*. Moreover, it also represents the overall support for party MP (in case issue *A* is accepted) or PP (in case *A* is rejected). In the Ostrogorski paradox, an individual votes for a party if and only if she agrees with that party on a majority of the issues. This rule can be represented as a rationality assumption by means of the following integrity constraint IC_O :

$$p_A \leftrightarrow [(p_E \wedge p_S) \vee (p_E \wedge p_F) \vee (p_S \wedge p_F)]$$

An instance of the Ostrogorski paradox can therefore be represented in the following profile:

	<i>E</i>	<i>S</i>	<i>F</i>	<i>A</i>
Voter 1	1	0	0	0
Voter 2	0	0	1	0
Voter 3	1	0	1	1
<i>Maj</i>	1	0	1	0

Each individual accepts issue *A* if and only if she accepts a majority of the other issues. However, the outcome of the majority rule is a rejection of issue *A*, even if a majority of the issues gets accepted by the same rule. Therefore, the triple (Maj, \mathbf{B}, IC_O) constitutes a paradox by Definition 1.

3.4 The Paradox of Multiple Elections

Whilst the Ostrogorski paradox was devised to stage an attack against representative systems of collective choice based on parties, the paradox of multiple elections (MEP) is based on the observation that when voting directly on multiple issues, a combination that was not supported nor liked by any of the voters can be the winner of the election (Brams et al., 1998; Lacy and Niou, 2000). While the original model takes into account the full preferences of individuals over combinations of issues, if we focus on only those ballots that are submitted by the individuals, then an instance of the MEP can be represented as a paradox of collective rationality. Let us consider the following simple example:

Multiple election paradox. Suppose three voters need to take a decision over three binary issues *A*, *B* and *C*. Their ballots are described as follows.

	A	B	C
Voter 1	1	0	1
Voter 2	0	1	1
Voter 3	1	1	0
<i>Maj</i>	1	1	1

The outcome of the majority rule is the acceptance of all three issues, even if this combination was not voted for by any of the individuals.

While there seems to be no integrity constraint directly causing this paradox, we may represent the profile in the example above as a situation in which the three individual ballots are bound by a constraint, e.g., $\neg(p_A \wedge p_B \wedge p_C)$. Even if each individual accepts at most two issues, the result of the aggregation is the unfeasible acceptance of all three issues.

As can be deduced from our previous discussion, every instance of the MEP gives rise to several instances of a binary aggregation paradox for Definition 1. To see this, it is sufficient to find an integrity constraint that is satisfied by all individuals and not by the outcome of the aggregation.² On the other hand, every instance of Definition 1 in binary aggregation may represent an instance of the MEP, as the irrational outcome cannot have been voted for by any of the individuals.

3.5 The Common Structure of Paradoxical Integrity Constraints

We can now make a crucial observation concerning the syntactic structure of the integrity constraints that formalise the paradoxes we have presented so far.

First, for the case of the Condorcet paradox, we observe that the formula encoding the transitivity of a preference relation is the implication $p_{ab} \wedge p_{bc} \rightarrow p_{ac}$. This formula is equivalent to $\neg p_{ab} \vee \neg p_{bc} \vee p_{ac}$, which is a clause of size 3, i.e., it is a disjunction of three different literals. Second, the formula which appears in the translation of the discursive dilemma is also equivalent to a clause of size 3, namely $\neg p_\alpha \vee \neg p_\beta \vee \neg p_{\neg(\alpha \wedge \beta)}$. Third, the formula which formalises the majoritarian constraint underlying the Ostrogorski paradox, is equivalent to the following conjunction of clauses of size 3:

$$(p_A \vee \neg p_E \vee \neg p_F) \wedge (p_A \vee \neg p_E \vee \neg p_S) \wedge (p_A \vee \neg p_S \vee \neg p_F) \wedge \\ \wedge (\neg p_A \vee p_E \vee p_F) \wedge (\neg p_A \vee p_E \vee p_S) \wedge (\neg p_A \vee p_S \vee p_F)$$

Finally, the formula which exemplifies the MEP is equivalent to a negative clause of size 3.

Thus, we observe that the **integrity constraints formalising the most classical paradoxes in aggregation theory all feature a clause of size at least 3.**³

4 The Majority Rule: Characterisation of Paradoxes

In this section we generalise the observation made in the previous section, and we provide a full characterisation of the class of integrity constraints that are lifted by the majority rule as those formulas that can be expressed as a conjunction of clauses of maximal size 2.

Under the majority rule, an issue is accepted if and only if a majority of the individuals accept it. Let $N_j^{\mathbf{B}}$ be the set of individuals that accept issue j in profile \mathbf{B} . In case the number of individuals is odd, the *majority rule* (*Maj*) has a unique definition by accepting issue j if and only if $|N_j^{\mathbf{B}}| \geq \frac{n+1}{2}$. The case of an even number of individuals is more problematic, to account for profiles in which exactly half of the individuals accept an issue

²Such a formula always exists: consider the disjunction of the formulas specifying the individual ballots.

³This observation is strongly related to a result proven by Nehring and Puppe (2007) in the framework of judgment aggregation, which characterises the set of paradoxical agendas for the majority rule as those agendas containing a minimal inconsistent subset of size at least 3.

and exactly half reject it. We give two different definitions. The *weak majority rule* (*W-Maj*) accepts an issue if and only if $|N_j^B| \geq \frac{n}{2}$, favouring acceptance. The *strict majority rule* (*S-Maj*) accepts an issue if and only if $|N_j^B| \geq \frac{n+2}{2}$, favouring rejection.

4.1 Odd Number of Individuals: The Majority Rule

We begin with a base-line result that proves collective rationality of the majority rule in case the integrity constraint is equivalent to a conjunction of 2-clauses. Let *2-clauses* denote the set of propositional formulas in \mathcal{L}_{PS} that are equivalent to a conjunction of clauses of maximal size 2.⁴

Proposition 1. *The majority rule is collectively rational with respect to 2-clauses.*

Proof. Let us first consider the case of a single 2-clause $IC = \ell_j \vee \ell_k$, where ℓ_j and ℓ_k are two literals, i.e., atoms or negated atoms. A paradoxical profile for the majority rule with respect to this integrity constraint features a first majority of individuals not satisfying literal ℓ_j , and a second majority of individuals not satisfying literal ℓ_k . By the pigeonhole principle these two majorities must have a non-empty intersection, i.e., there exists one individual that does not satisfy both literals ℓ_j and ℓ_k , but this is incompatible with the requirement that all individual ballots satisfy IC. To conclude the proof, it is sufficient to observe that if IC is equivalent to a conjunction of two clauses, then all individuals satisfy each of these clauses, and by the previous discussion all these clauses will also be satisfied by the outcome of the majority rule. \square

An easy corollary of this proposition covers the case of just 2 issues:

Corollary 2. *If $|\mathcal{I}| \leq 2$, then the majority rule is collectively rational with respect to all integrity constraints $IC \in \mathcal{L}_{PS}$.*

Proof. This follows immediately from Proposition 1 and from the observation that every formula built with two propositional symbols is equivalent to a conjunction of clauses of size at most 2 (e.g., its conjunctive normal form). \square

As we have remarked in Section 3.5, all classical paradoxes involving the majority rule can be formalised in our framework by means of an integrity constraint that consists of (or is equivalent to) one or more clauses with size bigger than two. We now generalise this observation to a theorem that completes the characterisation of the integrity constraints lifted by the majority rule. We need some preliminary definitions and a lemma.

Call a *minimally falsifying partial assignment* (mifap-assignment) for an integrity constraint IC an assignment to some of the propositional variables that cannot be extended to a satisfying assignment, although each of its proper subsets can. We first prove a crucial lemma about mifap-assignments. Given a propositional formula φ , associate with each mifap-assignment ρ for φ 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. The conjunction C_ρ represents the mifap-assignment ρ and it is clearly inconsistent with φ .

Lemma 3. *Every non-tautological formula φ is equivalent to $(\bigwedge_\rho \neg C_\rho)$ with ρ ranging over all mifap-assignments of φ .⁵*

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 eventually find a mifap-assignment ρ_A for

⁴The set of *2-clauses* can be equivalently defined by closing the set of 2-CNF under logical equivalence.

⁵Formulas $\neg C_\rho$ associated to mifap-assignments ρ for IC are also known as the *prime implicates* of IC. Lemma 3 is a reformulation of the fact that a formula is equivalent to the conjunction of its prime implicates.

φ 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 exists a ρ such that $A \models C_{\rho}$. This implies that $\rho \subseteq A$, as C_{ρ} is a conjunction. Since ρ is a mifap-assignment for φ , i.e., it is a falsifying assignment for φ , this contradicts the assumption that $A \models \varphi$. \square

We are now able to provide a full characterisation of the set of integrity constraints that are lifted by the majority rule in case the set of individuals is odd. Recall from Section 2 that $\mathcal{LF}[F]$ is the set of integrity constraints that are lifted by F .

Theorem 4. $\mathcal{LF}[Maj] = 2\text{-clauses}$.⁶

Proof. One direction is entailed by Proposition 1: the majority rule is CR with respect to formulas in 2-clauses . For the opposite direction assume that $IC \notin 2\text{-clauses}$, i.e., IC is not equivalent to a conjunction of 2-clauses. We now build a paradoxical profile for the majority rule. By Lemma 3 we know that IC is equivalent to the conjunction $\bigwedge_{\rho} \neg C_{\rho}$ of all mifap-assignments ρ for IC. We can therefore infer that at least one mifap-assignment ρ^* has size > 2 , for otherwise IC would be equivalent to a conjunction of 2-clauses.

Consider the following profile. Let y_1, y_2, y_3 be three propositional variables that are fixed by ρ^* . Let the first individual i_1 accept the issue associated with y_1 if $\rho(y_1) = 0$, and reject it otherwise, i.e., let $b_{1,1} = 1 - \rho^*(y_1)$. Furthermore, let i_1 agree with ρ^* on the remaining propositional variables. By minimality of ρ^* , this partial assignment can be extended to a satisfying assignment for IC, and let B_{i_1} be such an assignment. Repeat the same construction for individual i_2 , this time changing the value of ρ^* on y_2 and extending it to a satisfying assignment to obtain B_{i_2} . The same construction for i_3 , changing the value of ρ^* on issue y_3 and extending it to a satisfying assignment B_{i_3} . Recall that there are at least 3 individuals in \mathcal{N} . If there are other individuals, let individuals i_{3s+1} have the same ballot B_{i_1} , individuals i_{3s+2} ballot B_{i_2} and individuals i_{3s+3} ballot B_{i_3} . The basic profile for 3 issues and 3 individuals is shown in the following table:

	y_1	y_2	y_3
i_1	$1 - \rho^*(y_1)$	$\rho^*(y_2)$	$\rho^*(y_3)$
i_2	$\rho^*(y_1)$	$1 - \rho^*(y_2)$	$\rho^*(y_3)$
i_3	$\rho^*(y_1)$	$\rho^*(y_2)$	$1 - \rho^*(y_3)$
<i>Maj</i>	$\rho^*(y_1)$	$\rho^*(y_2)$	$\rho^*(y_3)$

As can be seen in the previous table, and easily generalised to the case of more than 3 individuals, there is a majority supporting ρ^* on every variable on which ρ^* is defined. Since ρ^* is a mifap-assignment and therefore cannot be extended to an assignment satisfying IC, the majority rule in this profile is not collectively rational with respect to IC. \square

4.2 Even Number of Individuals: Weak and Strict Majority

While a result analogous to Theorem 4 for the case of an even number of individuals cannot be proven, we provide the following result (proof is omitted for lack of space).

Proposition 5. *W-Maj and S-Maj are CR with respect to 2-clauses in which one literal is negative and one is positive. W-Maj is CR with respect to positive 2-clauses, in which all literals occur positively. S-Maj is CR with respect to negative 2-clauses, in which all literals occur negatively.*

⁶This result may be considered a “syntactic counterpart” of a result by Nehring and Puppe (2007) in the framework of judgment aggregation, characterising profiles on which the majority rule outputs a consistent outcome. In the interest of space, we refer to our previous work (Grandi and Endriss, 2011) for a more detailed discussion of the relation between the two results.

5 Conclusions

The first conclusion that can be drawn from this paper dedicated to paradoxes of aggregation is that the majority rule is to be avoided when dealing with collective choices over multiple issues. This fact stands out as a counterpart to May's Theorem (1952), which proves that the majority rule is the only aggregation rule *for a single binary issue* that satisfies a set of highly desirable conditions. The sequence of paradoxes we have analysed in this paper shows that this is not the case when multiple issues are involved. While this fact may not add anything substantially new to the existing literature, the wide variety of paradoxical situations encountered in this paper stresses even further the negative features of the majority rule on multi-issue domains.

A second conclusion is that most paradoxes of SCT share a common structure, and that this structure is formalised by our Definition 1, which stands out as a truly general definition of paradox in aggregation theory. Moreover, by analysing the integrity constraints that underlie some of the most classical paradoxes, we were able to identify a common syntactic feature of paradoxical constraints. Starting from this observation, we have provided a full characterisation of the integrity constraints that are lifted by the majority rule, as those formulas that are equivalent to a conjunction of clauses of size at most 2.

The paradoxical situations presented in this paper constitute a fragment of the problems that can be encountered in the formalisation of collective choice problems. For instance, paradoxical situations concerning voting procedures (Nurmi, 1999), which take as input a set of preferences and output a set of winning candidates, are not included in our analysis.

Recent work on paradoxes of aggregation also pointed at similarities within different frameworks, e.g., comparing the Ostrogorski paradox with the discursive dilemma (Pigozzi, 2005), or proposing a geometric approach for the study of paradoxical situations (Eckert and Klamler, 2009). The MEP gives rise to a different problem than that of collective rationality, not being directly linked to an integrity constraint established in advance. The problem formalised by the MEP is rather the *compatibility* of the outcome of aggregation with the individual ballots. Individuals in such a situation may be forced to adhere to a collective choice which, despite it being rational, they do not perceive as representing their views (Grandi and Pigozzi, 2012). Some answers to the problem raised by the multiple election paradox have already been proposed in the literature on AI, by for instance devising a suitable sequence of local elections (Xia et al., 2011), or by approximating the collective outcome (Conitzer and Xia, 2012).

Elections over multi-issue domains cannot be escaped: not only do they represent a model for the aggregation of more complex objects like preferences and judgments, but they also stand out as one of the biggest challenges to the design of more complex automated systems for collective decision making. A crucial problem in the modelling of real-world situations of collective choice is that of identifying the set of issues that best represent a given domain of aggregation, and devising an integrity constraint that models correctly the correlations between those issues. This problem obviously represents a serious obstacle to a mechanism designer, and is moreover open to manipulation. However, a promising direction for future work consists in structuring collective decision problems with more detailed models *before* the aggregation takes place, e.g., by discovering a shared order of preferential dependencies between issues (Lang and Xia, 2009; Airiau et al., 2011), facilitating the definition of collective choice procedures on complex domains without having to elicit the full preferences of individuals. Such models can be employed in the design and the implementation of automated decision systems, in which a safe aggregation, i.e., one that avoids paradoxical situations, is of the utmost necessity.

References

- S. Airiau, U. Endriss, U. Grandi, D. Porello, and J. Uckelman. Aggregating dependency graphs into voting agendas in multi-issue elections. In *Proceedings of IJCAI-2011*, 2011.
- S. J. Brams, D. M. Kilgour, and W. S. Zwicker. The paradox of multiple elections. *Social Choice and Welfare*, 15(2):211–236, 1998.
- Condorcet, Marie Jean Antoine Nicolas de Caritat, Marquis de. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Paris, 1785.
- V. Conitzer and L. Xia. Paradoxes of multiple elections: An approximation approach. In *Proceedings of KR-2012*, 2012.
- E. Dokow and R. Holzman. Aggregation of binary evaluations. *Journal of Economic Theory*, 145(2):495–511, 2010.
- D. Eckert and C. Klamler. A geometric approach to paradoxes of majority voting in abstract aggregation theory. In *Proceedings of ADT-2009*, 2009.
- W. Gaertner. *A Primer in Social Choice Theory*. Oxford University Press, 2006.
- U. Grandi and U. Endriss. Binary aggregation with integrity constraints. In *Proceedings IJCAI-2011*, 2011.
- U. Grandi and G. Pigozzi. On compatible multi-issue group decisions. In *Proceedings of LOFT-2012*, 2012.
- L. A. Kornhauser and L. G. Sager. Unpacking the court. *Yale Law Journal*, 96(1):82–117, 1986.
- D. Lacy and E. M. S. Niou. A problem with referendums. *Journal of Theoretical Politics*, 12(1):5–31, 2000.
- J. Lang. Vote and aggregation in combinatorial domains with structured preferences. In *Proceedings of IJCAI-2007*, 2007.
- J. Lang and L. Xia. Sequential composition of voting rules in multi-issue domains. *Mathematical Social Sciences*, 57:304–324, 2009.
- C. List and P. Pettit. Aggregating sets of judgments: An impossibility result. *Economics and Philosophy*, 18(1):89–110, 2002.
- 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. Pt. I: General characterization and possibility results on median spaces. *J. Economic Theory*, 135(1):269–305, 2007.
- H. Nurmi. *Voting paradoxes and how to deal with them*. Springer, 1999.
- M. Ostrogorski. *La démocratie et l'organisation de partis politiques*. Calmann-Levy, Paris, 1902.
- G. Pigozzi. Two aggregation paradoxes in social decision making: the Ostrogorski paradox and the discursive dilemma. *Episteme: A Journal of Social Epistemology*, 2(2):33–42, 2005.
- R. B. Wilson. On the theory of aggregation. *Journal of Economic Theory*, 10(1):89–99, 1975.
- L. Xia, V. Conitzer, and J. Lang. Strategic sequential voting in multi-issue domains and multiple-election paradoxes. In *Proceedings of ACM-2011*, 2011.

Umberto Grandi

Institute for Logic, Language and Computation

University of Amsterdam

Postbus 94242, 1090 GE Amsterdam, The Netherlands

Email: umberto.uni@gmail.com