

Goal-Based Collective Decisions: Axiomatics and Computational Complexity

Arianna Novaro¹, Umberto Grandi¹, Dominique Longin², Emiliano Lorini²

¹ IRIT, University of Toulouse

² IRIT, CNRS, Toulouse

{arianna.novaro, umberto.grandi, dominique.longin, emiliano.lorini}@irit.fr

Abstract

We study agents expressing propositional goals over a set of binary issues to reach a collective decision. We adapt properties and rules from the literature on Social Choice Theory to our setting, providing an axiomatic characterisation of a majority rule for goal-based voting. We study the computational complexity of finding the outcome of our rules (i.e., winner determination), showing that it ranges from Nondeterministic Polynomial Time (NP) to Probabilistic Polynomial Time (PP).

1 Introduction

Social choice and voting have become part of the standard computational toolbox for the design of rational agents that need to act in situations of collective choice [Shoham and Leyton-Brown, 2009; Brandt *et al.*, 2016]. In a variety of applications ranging from product configuration to multiple sensor control, the space of alternatives from which a collective choice needs to be taken is often combinatorial. This has brought many researchers to introduce compact languages for preference representation, and to design collective procedures that act directly on a compactly represented preference input (see, e.g., the survey by Lang and Xia [2016]).

When facing collective decisions with multiple binary issues, the framework of reference is judgment aggregation (see, e.g., List [2012], Lang and Slavkovik [2014], and Endriss [2016]). A vast literature explores the computational complexity of this framework [Endriss *et al.*, 2012; Baumeister *et al.*, 2015; de Haan and Slavkovik, 2017] and applications range from multiagent argumentation [Awad *et al.*, 2017] to the collective annotation of linguistic corpora [Qing *et al.*, 2014]. However, when considering collective decision-making in practice, the rigidity of representing individual views as complete judgments over issues poses serious obstacles, as becomes evident in the following example inspired by the *traveling group problem* [Klamer and Pferschy, 2007]:

Example 1. *An automated travel planner is organising a city trip for a group of friends, Ann, Barbara, and Camille, deciding whether to include a visit to a Church, a Museum, and a Park. Ann wants to see all the points of interest, Barbara prefers to have a walk in the Park, and Camille would like*

to visit a single point of interest but she does not care about which one. A judgment-based automated planner would require agents to specify a full valuation for each of the issues at stake, obtaining the following:

	Church	Museum	Park
Ann	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Barbara	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Camille	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>

The result by majority is a plan to visit both the Museum and the Park. However, Camille voted for the Museum only because asked for a complete judgment, and she was unable to express her truthful goal to “visit a single place, no matter which one” that the result does not satisfy.

The option of allowing individuals to abstain on issues, as proposed by Dietrich and List [2008] and Dokow and Holzman [2010b], is easily seen to be insufficient in modelling Camille’s preference in Example 1. Moreover, the obvious candidate for aggregating propositional goals, logic-based belief merging [Konieczny and Pérez, 2002], is quickly ruled out as its rules are not decisive, i.e., they often output a number of equally preferred plans. Building on an original idea of Lang [2004], we use a simple language of propositional goals to model individual preferences, defining and studying several rules to find the most preferred common alternative directly on such input.

A general tension exists in current models of collective decision making in combinatorial domains: on one side is the decisiveness or resoluteness of the rule—i.e., its ability to take a unique decision in most situations—and on the other side are fairness requirements, with respect to issues and individuals. Resoluteness is the primary concern in the development of decision-aid tools such as automated travel planners, or collective product configurators, to avoid returning to the users an excessive number of final options to choose from. Therefore, our purpose is to define rules that are as decisive as possible, whilst keeping high standards of fairness as defined by classical work in social choice and economic theory.

Related work. Judgment aggregation can be seen as goal-based voting in which individuals express single-model propositional goals. This is particularly evident in the binary aggregation model [Dokow and Holzman, 2010a; Grandi and Endriss, 2011], and is also true of judgments with abstentions,

which correspond to goals specified as partial conjunctions of (possibly negated) variables. Propositional goals have been proposed as compact representations of dichotomous preferences over combinatorial alternatives described by binary variables. Social choice with dichotomous preferences has been widely studied as a possible solution to the computational barriers affecting classical preference aggregation (see, e.g., the recent survey by Elkind *et al.* [2017]). However, to the best of our knowledge it has not been applied to combinatorial domains such as those studied in this paper. The vast literature on boolean games [Harrenstein *et al.*, 2001] studies similar situations in which agents are endowed with propositional goals; yet, our agents do not strategize and they control *together* the value of the issues at stake. Finally, we acknowledge an attempt at using logic-based belief merging to represent individual goals, using axiomatic properties from belief revision [Dastani and van der Torre, 2002].

Paper structure. In Section 2 we introduce the framework of goal aggregation, presenting our goal-based voting rules. In Section 3 we list desirable properties for such rules, and prove a characterisation result. In Section 4 we analyse the computational complexity of determining the winner of goal-based voting, and in Section 5 we conclude.

2 Goal-Based Voting

We begin with basic definitions and we introduce voting rules for goal-based collective decisions in multi-issue domains.

2.1 Basic Definitions

Let $\mathcal{N} = \{1, \dots, n\}$ be a set of *agents* deciding over a set $\mathcal{I} = \{1, \dots, m\}$ of binary *issues* or *propositions*. Agent i has *individual goal* γ_i , expressed as a consistent propositional formula over variables in \mathcal{I} (using standard connectives $\wedge, \vee, \rightarrow$ and \neg). For instance, in Example 1 Camille’s individual goal is $\gamma_3 = (1 \wedge \neg 2 \wedge \neg 3) \vee (\neg 1 \wedge 2 \wedge \neg 3) \vee (\neg 1 \wedge \neg 2 \wedge 3)$, since she wants to visit a single place.

An *interpretation* is a function $v : \mathcal{I} \rightarrow \{0, 1\}$ associating a binary value to each propositional variable. We often visualise v as the vector $(v(1), \dots, v(m))$. The set $\text{Mod}(\varphi) = \{v \mid v \models \varphi\}$ consists of all the *models* of formula φ . A goal is exponentially more succinct than the set of its models. In voting terminology, interpretations correspond to alternatives, and models of γ_i are the alternatives supported by agent i . We assume that issues in \mathcal{I} are independent, i.e., all interpretations over \mathcal{I} are feasible alternatives.

We indicate by $m_i(j) = (m_{ij}^1, m_{ij}^0)$ the number of 1/0 choices of agent i for issue j in the different models of her goal γ_i , where $m_{ij}^x = |\{v \in \text{Mod}(\gamma_i) \mid v(j) = x\}|$ for $x \in \{0, 1\}$. For example, if $\text{Mod}(\gamma_1) = \{(100), (010), (001)\}$ for issue $j = 3$ we have $m_1(3) = (m_{13}^1, m_{13}^0) = (1, 2)$.

A *goal-profile* $\Gamma = (\gamma_1, \dots, \gamma_n)$ collects the goals of all agents in \mathcal{N} and a *goal-based voting rule* is a function taking a goal-profile and returning a set of interpretations as the collective outcome. Formally, it is a collection of functions $F : (\mathcal{L}_{\mathcal{I}})^n \rightarrow \mathcal{P}(\{0, 1\}^m) \setminus \emptyset$ defined over any n and m whose input are n formulas submitted by the agents, and whose output is a set of models over the m issues in \mathcal{I} .

If a rule always outputs a singleton we call it *resolute*, and *irresolute* otherwise. We let $F(\Gamma)_j = (F(\Gamma)_j^0, F(\Gamma)_j^1)$, where $F(\Gamma)_j^x = |\{v \in F(\Gamma) \mid v_j = x\}|$ for $x \in \{0, 1\}$, indicate the amount of 0/1 choices in the outcome of F for j . We write $F(\Gamma)_j = x$ in case $F(\Gamma)_j^{1-x} = 0$ for $x \in \{0, 1\}$.

2.2 Conjunction and Approval Rules

We begin by introducing the following baseline rule:

$$\text{Conj}_v(\Gamma) = \begin{cases} \text{Mod}(\gamma_1 \wedge \dots \wedge \gamma_n) & \text{if non-empty} \\ \{v\} \text{ for } v \in \{0, 1\}^m & \text{otherwise} \end{cases}$$

The *conjunction rule* is an irresolute rule that outputs those alternatives on which all agents agree, and a default otherwise. While such consensual alternatives are clearly an optimal choice, they rarely exist—the purpose of voting being to find compromises among conflicting individual goals. To avoid default options, we introduce the following rule:

$$\text{Approval}(\Gamma) = \arg \max_{v \in \text{Mod}(\bigvee_{i \in \mathcal{N}} \gamma_i)} |\{i \in \mathcal{N} \mid v \in \text{Mod}(\gamma_i)\}|.$$

This rule expresses simple approval voting [Brams and Fishburn, 2007; Laslier and Sanver, 2010]. It is also studied by Lang [2004] as the plurality rule, and in belief merging as an instance of $\Delta_{\mu}^{\sum, d}$ -rules by Konieczny and Pérez [2002]. Despite its intuitive appeal, approval-based voting is not adapted to combinatorial domains in which a large number of alternatives might be approved by a few agents only.

2.3 Issue-Wise Voting

We first introduce a large class of goal-based rules inspired by the well-known quota rules from judgment aggregation [Dietrich and List, 2007]. Let $\mu_{\varphi} : \text{Mod}(\varphi) \rightarrow \mathbb{R}$ be a function associating to each model v of φ some weight $\mu_{\varphi}(v)$, giving (possibly) different weights to distinct models of the same formula. Let *threshold rules* be defined as follows:

$$\text{TrSh}^{\mu}(\Gamma)_j = 1 \text{ iff } \left(\sum_{i \in \mathcal{N}} (w_i \cdot \sum_{v \in \text{Mod}(\gamma_i)} v(j) \cdot \mu_{\gamma_i}(v)) \right) \geq q_j$$

such that $1 \leq q_j \leq n$ for all $j \in \mathcal{I}$ is the *quota* of issue j , where for each $v \in \text{Mod}(\gamma_i)$ we have $\mu_{\gamma_i}(v) \neq 0$ and $w_i \in [0, 1]$ is the individual weight of agent i . To ease notation we omit the vector $\mathbf{q} = (q_1, \dots, q_m)$ from TrSh^{μ} , specifying the particular choice of thresholds for the issues. Intuitively, threshold rules set a quota to be passed for each issue to be accepted, with the additional flexibility of weights for agents and for models of the individual goals.

From here, we can provide a first adaptation of the classical issue-wise majority voting for goal-based settings. Inspired by *equal and even cumulative voting* [Campbell, 1954] we call *EQuota* rules those TrSh^{μ} procedures having $\mu_{\gamma_i}(v) = \frac{1}{|\text{Mod}(\gamma_i)|}$ and $w_i = 1$ for all $v \in \text{Mod}(\gamma_i)$ and for all $i \in \mathcal{N}$. Thus, the *equal and even majority rule* EMaj is the *EQuota* rule having $q_j = \lceil \frac{n}{2} \rceil$ for all $j \in \mathcal{I}$.

A second (irresolute) version of majority voting simply compares for each issue the number of acceptances with the number of rejections, weighting each goal model as *EQuota*:

$$\text{TrueMaj}(\Gamma) = \Pi_{j \in \mathcal{I}} M(\Gamma)_j$$

	Γ^1	Γ^2	Γ^3
Mod(γ_1)	(111)	(111)	(000)
Mod(γ_2)	(001)	(111) (011) (000)	(111) (110) (011)
Mod(γ_3)	(100) (010) (101)	(111) (011) (000)	(111) (110) (011)
EMaj	(001)	—	(010)
TrueMaj	(101)	(111)	—
2sMaj	—	(011)	(111)

Table 1: Three goal-profiles on which majority-based rules differ.

where for each $j \in \mathcal{I}$:

$$M(\Gamma)_j = \begin{cases} \{x\} & \text{if } \sum_{i \in \mathcal{N}} \frac{m_{ij}^x}{|\text{Mod}(\gamma_i)|} > \sum_{i \in \mathcal{N}} \frac{m_{ij}^{1-x}}{|\text{Mod}(\gamma_i)|} \\ \{0, 1\} & \text{otherwise} \end{cases}$$

Intuitively, *TrueMaj* computes a weighted count of the 1s and the 0s in all models of the individual goals, discounted by the number of models of the formula sent by the agent. In case of ties on an issue, the rule outputs all interpretations with either 0 or 1 for that issue.

We define a third version of the majority rule as $2sMaj(\Gamma) = Maj(Maj(\gamma_1), \dots, Maj(\gamma_n))$, where *Maj* is the classical issue-by-issue strict majority rule, that accepts an issue if and only if a strict majority of the models of γ_i does. This procedure belongs to a wider class of rules that can be obtained by applying a first rule on each individual goal, and a second, possibly different, rule on the results obtained in the first step.

We now prove that the three proposed versions of goal-based majority do not always return the same result.

Proposition 1. *There exists goal-profiles on which the outcomes of EMaj, TrueMaj and 2sMaj differ.*

Proof sketch. In Table 1 we provide a profile for each pair of rules on which their results differ. Consider Γ^1 and *EMaj*. For agents 1 and 2, the weight of the single model satisfying their goal is 1, while for the third agent is $\frac{1}{|\text{Mod}(\gamma_3)|} = \frac{1}{3}$.

If we focus on the first issue, $\frac{m_{11}^1}{|\text{Mod}(\gamma_1)|} + \frac{m_{21}^1}{|\text{Mod}(\gamma_2)|} + \frac{m_{31}^1}{|\text{Mod}(\gamma_3)|} = 1 + \frac{2}{3} < 2$, hence $EMaj(\Gamma^1)_1 = 0$. Take *TrueMaj* instead.

Since $\sum_{i \in \mathcal{N}} \frac{m_{i1}^1}{|\text{Mod}(\gamma_i)|} = 1 + \frac{2}{3} > 1 + \frac{1}{3} = \sum_{i \in \mathcal{N}} \frac{m_{i1}^0}{|\text{Mod}(\gamma_i)|}$ we get $TrueMaj(\Gamma^1)_1 = 1$. The calculations for the other cases can be obtained straightforwardly. \square

3 Axiomatic Analysis

In this section we conduct an axiomatic analysis of the proposed rules and we provide a characterisation of *TrueMaj*.

3.1 Axiom Definitions

A first straightforward generalisation of an axiom from the literature on Social Choice Theory is the following:

Definition 1. *A rule F is anonymous (A) if for any profile Γ and permutation $\sigma : \mathcal{N} \rightarrow \mathcal{N}$, we have that $F(\gamma_1, \dots, \gamma_n) = F(\gamma_{\sigma(1)}, \dots, \gamma_{\sigma(n)})$.*

Observe that all the presented rules are anonymous, except for threshold rules with varying weights for the agents.

Define $\varphi[j \mapsto k]$ for $j, k \in \mathcal{I}$ as the replacement of each occurrence of j by k in φ . The next axiom ensures that issues are treated equally:

Definition 2. *A rule F is neutral (N) if for all Γ and $\sigma : \mathcal{I} \rightarrow \mathcal{I}$, we have $F(\gamma_1^\sigma, \dots, \gamma_n^\sigma) = \{(v(\sigma(1)), \dots, v(\sigma(m))) \mid v \in F(\Gamma)\}$ where $\gamma_i^\sigma = \gamma_i[1 \mapsto \sigma(1), \dots, m \mapsto \sigma(m)]$.*

TrSh $^\mu$ and *EQuota* rules are not neutral when the quotas for two issues differ. Neither is *Conj $_v$* , by permuting issues in a profile of inconsistent goals resulting in a profile of inconsistent goals, so that the same default v is chosen. *Approval* is neutral, since the values for the issues are permuted in the models of the agents' goals. Both *TrueMaj* and *2sMaj* have the same quota for all issues, and hence they are neutral.

We then move to a controversial yet well-known axiom in the literature, used in both characterisation and impossibility results [List, 2012; Brandt *et al.*, 2016].

First, let $\mathcal{D}_m = \{(a, b) \mid a, b \in \mathbb{N} \text{ and } a + b \leq 2^m\}$ and $\mathcal{C} = \{\{0\}, \{1\}, \{0, 1\}\}$. Independence is formally defined as:

Definition 3. *A rule F is independent (I) if there are functions $f : \mathcal{D}_m \rightarrow \mathcal{C}$ for $j \in \mathcal{I}$ such that for all profiles Γ we have $F(\Gamma) = \prod_{j \in \mathcal{I}} f(m_1(j), \dots, m_n(j))$.*

Albeit being often identified as one of the main sources of impossibilities in aggregation theory [List, 2012], we believe that independent (i.e., issue-wise) rules are crucial in solving the tension between fairness and resoluteness in goal-based voting. From the definitions we easily see that *TrSh $^\mu$* , *EQuota* and *TrueMaj* are independent, while *Conj $_v$* and *Approval* are not since they consider the profile globally.

The next axiom holds whenever the unanimous choice of the agents for an issue is respected in the outcome:

Definition 4. *A rule F is unanimous (U) if for all profiles Γ and for all $j \in \mathcal{I}$, if $m_{ij}^x = 0$ for all $i \in \mathcal{N}$ then $F(\Gamma)_j = 1 - x$ for $x \in \{0, 1\}$.*

While if all agents accept or reject an issue the output of *TrueMaj* and *2sMaj* will agree with the profile, interestingly *TrSh $^\mu$* and *EQuota* rules do not satisfy it (by setting a high enough quota) as well as *Conj $_v$* (for a profile where goals are inconsistent and thus the default is chosen).

We say that profiles Γ and Γ' are *comparable* if and only if for all $i \in \mathcal{N}$ we have that $|\text{Mod}(\gamma_i)| = |\text{Mod}(\gamma'_i)|$. Then, a rule is positively responsive if adding (deleting) support for issue j when the result for j is equally irresolute or favouring acceptance (rejection), results in an outcome strictly favouring acceptance (rejection) for j .

Definition 5. *A rule F satisfies positive responsiveness (PR) if for all comparable profiles $\Gamma = (\gamma_1, \dots, \gamma_i, \dots, \gamma_n)$ and $\Gamma^* = (\gamma_1, \dots, \gamma_i^*, \dots, \gamma_n)$, for all $j \in \mathcal{I}$ and $i \in \mathcal{N}$, if $m_{ij}^{x^*} \geq m_{ij}^x$ for $x \in \{0, 1\}$, then $F(\Gamma)_j^{1-x} \geq F(\Gamma)_j^x$ implies $F(\Gamma^*)_j^{1-x} > F(\Gamma^*)_j^x$.*

Observe that all our presented versions of majority are positively responsive, since they have a threshold of acceptance.

We conclude by presenting two important fairness axioms. The first aims at formalising the “one man, one vote” principle, and ensures that a rule is giving *equal* weight to the models of each individual goal for all the agents. It is satisfied by all *EQuota* rules as well as by *TrueMaj*:

Definition 6. A rule F is egalitarian (E) if for all Γ , on the profile Γ' with $|\mathcal{N}'| = |\mathcal{N}| \cdot \text{lcm}(|\text{Mod}(\gamma_1)|, \dots, |\text{Mod}(\gamma_n)|)$, and for all $i \in \mathcal{N}$ and $v \in \text{Mod}(\gamma_i)$ there are $\frac{|\mathcal{N}'|}{|\mathcal{N}| \cdot |\text{Mod}(\gamma_i)|}$ agents in Γ' having goal γ with $\text{Mod}(\gamma) = \{v\}$, it holds that $F(\Gamma) = F(\Gamma')$.

The second axiom instead focuses on possible biases towards acceptance or rejection of the issues.

Definition 7. A rule is dual (D) if for all profiles Γ , $F(\bar{\gamma}_1, \dots, \bar{\gamma}_n) = \{(1 - v(1), \dots, 1 - v(m)) \mid v \in F(\Gamma)\}$ where $\bar{\gamma} = \gamma[-1 \mapsto 1, \dots, -m \mapsto m]$.

A similar requirement is called neutrality by May [1952], while in binary aggregation this is known as domain-neutrality [Grandi and Endriss, 2011].

3.2 Characterising Goal-Based Majority Rules

A seminal result in characterising aggregation rules is May’s Theorem [1952], where an axiomatisation of the majority rule in the context of voting over two alternatives is provided. A natural question to ask after defining three versions of the majority rule is therefore whether one can be axiomatised, building on May’s results. We answer this question in the positive:

Theorem 1. A rule F satisfies (E), (I), (N), (A), (PR), (U) and (D) if and only if it is *TrueMaj*.

Proof. Right-to-left follows from discussion in Section 3.1. For left-to-right, consider a rule F . Let Γ be an arbitrary profile over n voters and m issues. By (E), we can construct a profile Γ' for m issues and n' agents, where n' is as in Definition 6, in which each agent submits a single-model goal and such that $v \in F(\Gamma)$ if and only if $v \in F(\Gamma')$. We therefore consider the restriction of F on profiles over n' agents and m issues where agents submit single-model goals. We denote \mathcal{G}^\square such a set of profiles (hence, in particular, $\Gamma' \in \mathcal{G}^\square$). We now show that $F(\Gamma') = \text{TrueMaj}(\Gamma')$.

By (I), there are functions f_1, \dots, f_m such that $F(\Gamma') = f_1(m_1(1), \dots, m_n(1)) \times \dots \times f_m(m_1(m), \dots, m_n(m))$. Observe that since $\Gamma' \in \mathcal{G}^\square$, we have $m_i(j) \in \{(0, 1), (1, 0)\}$ for all $i \in \mathcal{N}$ and $j \in \mathcal{I}$. Hence, we can equivalently see each f on profiles in \mathcal{G}^\square as a function from $\{0, 1\}^n$ to \mathcal{C} . By (N) and (I) we get that $f_1 = \dots = f_m$, i.e., the *same* function applies to all issues, let us denote it with f .

By (A), any permutation of agents in Γ' gives the same result $F(\Gamma')$. Hence, combining (A) with (I) and (N), we have that only the number of ones (and zeroes) and not their position is necessary to determine the outcome of f . Hence, we can write it as $f : \{0, \dots, n\} \rightarrow \mathcal{C}$.

Consider now a profile $\Gamma^+ \in \mathcal{G}^\square$ such that for all $i \in \mathcal{N}$ we have $m_{ij}^0 = 0$. By (U) we know that $F(\Gamma^+)_j^0 = 0$, i.e., $v(j) = 1$ for all $v \in F(\Gamma)$, and consequently that $f(n) = \{1\}$. Analogously we obtain that $f(0) = \{0\}$.

Let now s be a sequence of \mathcal{G}^\square -profiles $\Gamma^- = \Gamma^0, \Gamma^1, \dots, \Gamma^n = \Gamma^+$ where *exactly* one agent i at a step k changes her

goal γ_i such that $m_{ij}^1 = 0$ in Γ^k and $m_{ij}^1 = 1$ in Γ^{k+1} . By (I) and the definition of cartesian product, for any Γ and j , $F(\Gamma)_j$ is either equal to $(a, 0)$, (b, b) or $(0, c)$ for $a, b, c \in \mathbb{N}$.

By (PR), the outcome of the Γ^k profiles in s can only switch from $(a, 0)$ to (b, b) or $(0, c)$, and from (b, b) to $(0, c)$. In particular, this means that there is some number q such that $f(0) = \{0\}, \dots, f(q-1) = \{0\}, f(q) = \{0, 1\}$ or $f(q) = \{1\}$, and $f(q+1) = \{1\}, \dots, f(n) = \{1\}$.

We now show that for n even, $q = \frac{n}{2}$ and $f(q) = \{0, 1\}$, while for n odd we have $q = \frac{n+1}{2}$ and $f(q) = \{1\}$. For n even, consider profile Γ^ℓ where exactly half of the agents accept j . If $F(\Gamma^\ell)_j = (0, a)$ or $(c, 0)$, by (D) we would need to reverse the outcome in $F(\bar{\Gamma}^\ell)_j$. However, the decision in both profiles is determined by $f(\frac{n}{2})$, which is therefore equal to $\{0, 1\}$. For n odd, suppose that $q < \frac{n+1}{2}$ and consider a profile $\bar{\Gamma}$ where there are exactly q agents accepting j . By (PR) we have $F(\Gamma)_j = (0, c)$. Consider now profile $\bar{\Gamma}$: we have $|\{i \mid m_i(j) = (0, 1) \text{ for } \gamma_i \in \Gamma\}| = |\{i \mid m_i(j) = (1, 0) \text{ for } \gamma_i \in \bar{\Gamma}\}| = q < \frac{n+1}{2}$. Hence, $|\{i \mid m_i(j) = (0, 1) \text{ for } \gamma_i \in \bar{\Gamma}\}| \geq \frac{n+1}{2} > q$. Hence, $F(\bar{\Gamma})_j = (0, c)$, contradicting (D). Suppose $q > \frac{n+1}{2}$ and consider a profile Γ where $\frac{n+1}{2} \leq |\{i \mid m_i(j) = (0, 1) \text{ for } \gamma_i \in \Gamma\}| < q$. Then, $F(\Gamma)_j = (a, 0)$ and $F(\bar{\Gamma})_j = (a, 0)$, again contradicting (D).

To sum up, F is defined as the cartesian product of binary decisions taken by the same function $f : \{0, \dots, n\} \rightarrow \mathcal{C}$. on each issue, with $f(k) = \{0, 1\}$ for n even and $k = \frac{n}{2}$, $f(k) = \{0\}$ for $\sum_{i \in \mathcal{N}'} m_{ij}^x > \sum_{i \in \mathcal{N}'} m_{ij}^{1-x}$ for $x \in \{0, 1\}$, corresponding to the definition of *TrueMaj*. Since Γ is an arbitrary goal profile, and *TrueMaj* satisfies (E), we obtained the desired equivalence. \square

While both *EMaj* and *2sMaj* are based on similar intuitions as *TrueMaj*, *EMaj* has a bias towards the rejection of the issues, while *2sMaj* does not satisfy the equality axiom. *TrueMaj* however remains the only irresolute rule of the three, once more showing the tension between fairness criteria and the decisiveness of a goal-based voting rule.

We conclude with a seemingly negative result. A rule is *grounded* if $v \in F(\Gamma)$ implies $v \in \text{Mod}(\gamma_1 \vee \dots \vee \gamma_n)$.

Proposition 2. *EQuota*, *TrueMaj* and *2sMaj* are not grounded.

Proof. Consider profile Γ for 3 agents and 3 issues where $\text{Mod}(\gamma_1) = \{(111)\}$, $\text{Mod}(\gamma_2) = \{(010)\}$ and $\text{Mod}(\gamma_3) = \{(001)\}$. Both *EQuota* (with uniform quota 2), *TrueMaj* and *2sMaj* return (011) , contradicting groundedness. \square

Hence, the three majority rules do not guarantee that the collective choice will satisfy the goal of at least one agent. However, by considering aggregation as compromising between agents, it becomes less important for a rule to be grounded.

4 Computational Complexity

In this section we study the computational complexity of determining the result of goal-based voting, showing that propositional goals entail a significant increase from standard voting, in some cases from P to Probabilistic Polynomial time.

4.1 Winner Determination

We present two definitions for the winner determination problem, for resolute and irresolute rules, in line with the literature on judgment aggregation [Endriss *et al.*, 2012; Baumeister *et al.*, 2015; de Haan and Slavkovik, 2017].

Note that we provide the existential version of the winner determination problem — a universal definition is also possible [Lang and Slavkovik, 2014]. We start with resolute rules:

WINDET(F)

Input: profile Γ , issue j

Question: Is it the case that $F(\Gamma)_j = 1$?

The outcome for $F(\Gamma)$ can then be computed by repeatedly answering the question in WINDET over all issues $j \in \mathcal{I}$. Next, we introduce the problem for irresolute rules.

WINDET*(F)

Input: profile Γ , set $S \subseteq \mathcal{I}$, partial model $\rho : S \rightarrow \{0, 1\}$

Question: Is there a $v \in F(\Gamma)$ with $v(j) = \rho(j)$ for $j \in S$?

By answering to the question in WINDET* starting from a set S with one issue and filling it with all the issues in \mathcal{I} , and checking possible values for the partial function ρ we can construct a complete model in the outcome of $F(\Gamma)$.

4.2 Conjunction and Approval Rules

Our first complexity result provides a lower bound for the family of conjunction rules $Conj_v(\Gamma)$.

Theorem 2. WINDET*($Conj_v$) is NP-hard.

Proof. We reduce from SAT. Let $\varphi^*[p_1^*, \dots, p_k^*]$ be the formula over k variables whose satisfiability we want to check. Construct an instance of WINDET($Conj_v$) as follows. Let $\mathcal{I} = \{p_1^*, \dots, p_k^*\} \cup \{q\}$. Consider a profile $\Gamma = (\gamma_1)$ for a single agent 1, such that $\gamma_1 = q \oplus \varphi^*$, for \oplus the exclusive or. This formula is true if and only if either q is true or φ^* is true, so that the default model v is not needed. If we set $S = \{q\}$ and $\rho(q) = 0$, we get that φ^* is satisfiable if and only if for this instance of WINDET*($Conj_v$) the answer is yes. \square

Membership in NP is still open. The intuitive algorithm that guesses a model v , then checks whether $v \models \bigwedge_{i \in \mathcal{N}} \gamma_i$, if the answer is negative it checks $v \models \bigwedge_{\substack{j \in S \\ \rho(j)=1}} j \wedge \bigwedge_{\substack{j' \in S \\ \rho(j')=0}} \neg j'$

(i.e., the formula expressing ρ), excludes the case in which the conjunction of the goals is satisfiable, but v is not a model.

The *Approval* rule is significantly harder. We first need some definitions. Let $\Theta_2^p = \text{P}^{\text{NP}[\log]}$ be the class of decision problems solvable in polynomial time by a Turing machine that can make $\mathcal{O}(\log n)$ queries to an NP oracle, for n the size of the input. Consider the following Θ_2^p -complete problem [Chen and Toda, 1995]:

MAX-MODEL

Input: satisfiable propositional formula φ , variable p of φ

Question: Is there a model $v \in \text{Mod}(\varphi)$ that sets a maximal number of variables of φ to true and such that $v(p) = 1$?

We are now ready to prove the following:

Theorem 3. WINDET*(*Approval*) is Θ_2^p -complete.

Proof. Membership in Θ_2^p can be obtained from Proposition 4 by Lang [2004], using the following formula in the definition of the ELECT-SAT problem:

$$\psi = \bigwedge_{\substack{j \in S \\ \rho(j)=1}} j \wedge \bigwedge_{\substack{j' \in S \\ \rho(j')=0}} \neg j'.$$

For completeness, we give a reduction from MAX-MODEL. Consider an instance of MAX-MODEL where $\varphi[p_1, \dots, p_m]$ is a satisfiable formula and p_i for $i \in \{1, \dots, m\}$ is one of its variables. Construct now an instance of WINDET*(*Approval*) in the following way. Let $\Gamma = (\gamma_1, \dots, \gamma_{m+1}, \gamma_{m+2}, \dots, \gamma_{2m+1})$ be a profile such that $\gamma_1 = \dots = \gamma_{m+1} = \varphi$ and $\gamma_{m+2} = p_1, \dots, \gamma_{2m+1} = p_m$.

We have that $\text{Approval}(\Gamma) \subseteq \text{Mod}(\varphi)$, since a strict majority of $\frac{m+1}{2m+1}$ agents already supports all the models of φ . Moreover, note that in this instance of *Approval* precisely the models maximising the number of variables set to true in φ win. In fact, consider a model $v \in \text{Mod}(\varphi)$: as explained, v gets the support of all the first $m+1$ agents whose goal is φ , and then for all the agents in $\{m+2, \dots, 2m+1\}$ it gets the support of those agents whose goal-variable is true in v . Specifically, the support of v is $(m+1) + |\{p_i \mid v(p_i) = 1\}|$.

Hence, only those $v \in \text{Mod}(\varphi)$ with a maximal number of 1s are in the outcome of $\text{Approval}(\Gamma)$. It now suffices to set $S = \{p_i\}$ for p_i the propositional variable in the instance of MAX-MODEL and $\rho(p_i) = 1$. Therefore, a formula φ has a model with a maximal number of variables set to true where p_i is true if and only if WINDET*(*Approval*) returns yes on the constructed input. \square

4.3 Threshold Rules

We study the complexity of finding the outcome of $TrSh^u$ rules for the special case where each model, as well as each agent, has the same weight of 1. We start by studying the following auxiliary problem.

k -MODELSUM

Input: propositional formulas ψ_1, \dots, ψ_ℓ , number $k \in \mathbb{N}$

Question: Is it the case that $\sum_{1 \leq i \leq \ell} |\text{Mod}(\psi_i)| > k$?

We now find the complexity for k -MODELSUM.

Lemma 1. k -MODELSUM is NP-complete.

Proof. To show membership in NP guess k_1, \dots, k_ℓ numbers with $k_i \leq k+1$ for all $1 \leq i \leq \ell$, and guess X_1, \dots, X_ℓ sets, where $X_i \subseteq 2^{\text{Var}}$ for $1 \leq i \leq \ell$ and Var is the set of variables of ψ_1, \dots, ψ_ℓ . The size of each X_i for $1 \leq i \leq \ell$ is bounded by $k+1$, and each X_i corresponds to a set of models. It is then easy to check that $k_1 + \dots + k_\ell > k$, that for all $1 \leq i \leq \ell$ we have $|X_i| = k_i$ and for all $v_i \in X_i$ we have $v_i \in \text{Mod}(\psi_i)$.

For completeness, we reduce from SAT. Let φ^* be the formula whose satisfiability we want to check. Construct now an instance of k -MODELSUM where $\psi_1 = \varphi^*$ and $k = 1$. Formula φ^* is satisfiable if and only if it has at least one model, and SAT can be reduced to k -MODELSUM. \square

Now we can assess the complexity of the rule $TrSh^u$.

Theorem 4. For $\mu_{\gamma_i}(v) = 1$ constant and $w_i = 1$ for all $i \in \mathcal{N}$, $\text{WINDET}(\text{TrSh}^\mu)$ is NP-complete.

Proof. For membership in NP consider a profile $\Gamma = (\gamma_1, \dots, \gamma_n)$ and an issue j . Guess k_1, \dots, k_n numbers with $k_i \leq k+1$ for $1 \leq i \leq n$, and guess X_1, \dots, X_n sets of models where $X_i \subseteq 2^m$ and for each $v \in X_i$ we have $v(j) = 1$ for $i \in \mathcal{N}$. It is then easy to check whether $k_1 + \dots + k_n > q_j$, that for all $i \in \mathcal{N}$ we have $|X_i| = k_i$ and for all $v_i \in X_i$ we have $v_i \in \text{Mod}(\gamma_i)$.

For completeness, we reduce from k -MODELSUM. Let ψ_1, \dots, ψ_ℓ and $k \in \mathbb{N}$ be an instance of this problem. Construct now an instance of $\text{WINDET}(\text{TrSh}^\mu)$ such that $\Gamma = (\gamma_1, \dots, \gamma_\ell)$ where for all $i \in \mathcal{N}$ we have $\gamma_i = \psi_i \oplus \neg p$ for p a fresh variable and \oplus the exclusive or. This is done since the formulas of the k -MODELSUM might be inconsistent, while individual goals are always consistent. Now, we choose $j = p$ and we set $q_j = k$. In this way, every model v such that $v(p) = 1$ is a model of ψ_i for all $i \in \mathcal{N}$, and we can thus count if there are at least k models of each ψ_i , which give $\text{TrSh}^\mu(\Gamma)_p = 1$. \square

While it would be easy to adapt this proof to deal with different values for the individual weights w_i (to be multiplied with the k_i 's), for model weights as the ones in *EQuota* rules it would be necessary to compute the number of models of each goal, thus making it a more difficult problem.

4.4 Majority Rules

We now study the complexity of majority rules. We introduce the complexity class PP, for Probabilistic Polynomial Time, a class of problems that has rarely been encountered in the literature on computational social choice, and we show that the three versions of the majority rule are PP-hard. Membership is an open problem for future work.

Let PP be the class of decision problems solvable by a non-deterministic Turing machine that accepts in strictly more than half of all non-deterministic choices if and only if the answer to the problem is yes [Papadimitriou, 2003]. Consider the following problem:

MAJ-SAT- p

Input: propositional formula φ , variable p of φ

Question: Is it the case that $|\text{Mod}(\varphi \wedge p)| > |\text{Mod}(\varphi \wedge \neg p)|$?

We first show that MAJ-SAT- p is PP-complete by reducing from the PP-complete MAJ-SAT, the problem of deciding whether a formula φ has more models than its negation.

Lemma 2. MAJ-SAT- p is PP-complete.

Proof. We start by showing membership in PP. Consider the non-deterministic Turing machine that guesses a model v for φ . Then, if $v \models \varphi$ the machine accepts with probability $\frac{1}{2}$. If $v \models \varphi \wedge p$ the machine accepts with probability 1 and if $v \models \varphi \wedge \neg p$ the machine accepts with probability 0.

For completeness, we reduce from MAJ-SAT. Consider the formula φ as our instance of MAJ-SAT, and now let $\psi = (\varphi \wedge p) \vee (\neg \varphi \wedge \neg p)$ for p a fresh variable. We can now observe that φ has more models than $\neg \varphi$ if and only if $\psi \wedge p$ has more models than $\psi \wedge \neg p$, concluding the reduction. \square

The next theorem gives a lower bound to computing the outcome of the majority rules. Note that for *TrueMaj* we study a (strict) resolute version *TrueMaj^r*.

Theorem 5. $\text{WINDET}(2sMaj)$, $\text{WINDET}(EMaj)$ and $\text{WINDET}(\text{TrueMaj}^r)$ are PP-hard.

Proof. The proof is analogous for the three rules, so we only prove it for *2sMaj*. We reduce from MAJ-SAT- p . Consider the formula φ and the variable p of φ as our instance of MAJ-SAT- p . Consider now a profile Γ for a single agent such that $\gamma_1 = \varphi$. Since we are dealing with resolute rules, we simply have to fix an issue and ask whether the goal-based voting rule will accept or reject the issue. Given that there is a single agent 1, we have that $2sMaj(\Gamma)_p = 1$ if and only if the set of models of γ_1 accepts p more often than reject it. Therefore, $\varphi \wedge p$ has more models than $\varphi \wedge \neg p$ if and only if $2sMaj(\Gamma)_p = 1$, completing the reduction. \square

While PP is the hardest class studied in this paper, the axiomatic analysis of Section 3 as well as their intuitive definitions make us champion our majority-based rules, and *TrueMaj* in particular. We argue that the class PP is pervasive in propositional goal-based reasoning, calling for the development of good algorithms for problems in this class.

5 Conclusions and Future Work

Starting from the observation that classical judgment aggregation falls short in many examples of collective decision-making in multi-issue domains, such as creating a shared travel plan or collective product configuration, we introduced new rules to aggregate a set of propositional goals into a collectively satisfying alternative. In a quest for resolute rules, we introduced three adaptations of the classical majority rule, as well as other goal-based voting rules, providing an axiomatic characterisation in line with the literature on Social Choice Theory for one of them (*TrueMaj*). We concluded by investigating the computational complexity of determining the outcome of our rules, showing that the use of propositional goals entails harder complexity classes.

Our results open several paths for future research, most notably in studying restrictions on the language of goals that might determine islands of tractability for the winner determination problem, or develop tractable approximations for their computation. Moreover, the quest for more resolute and decisive rules may suggest novel voting procedures in related areas such as non-binary combinatorial domains and more expressive compact languages for preference representation. Finally, we focused on the basic case of no integrity constraints, but it would be interesting to study classes of constraints for which our rules always return consistent results.

Acknowledgments

We are grateful to the AAMAS and IJCAI reviewers, as well as to Jérôme Lang and Ronald de Haan for their helpful comments. This work was partially supported by the project ‘‘Social Choice on Networks’’ of Labex CIMI ANR-11-LABX-0040-CIMI within the program ANR-11-IDEX-0002-02.

References

- [Awad *et al.*, 2017] Edmond Awad, Richard Booth, Fernando Tohmé, and Iyad Rahwan. Judgement aggregation in multi-agent argumentation. *Journal of Logic and Computation*, 27(1):227–259, 2017.
- [Baumeister *et al.*, 2015] Dorothea Baumeister, Gábor Erdélyi, Olivia Johanna Erdélyi, and Jörg Rothe. Complexity of manipulation and bribery in judgment aggregation for uniform premise-based quota rules. *Mathematical Social Sciences*, 76:19–30, 2015.
- [Brams and Fishburn, 2007] Steven Brams and Peter C. Fishburn. *Approval voting*. Springer Science & Business Media, 2007.
- [Brandt *et al.*, 2016] Felix Brandt, Vincent Conitzer, Ulle Endriss, Ariel D Procaccia, and Jérôme Lang. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- [Campbell, 1954] Whitney Campbell. The origin and growth of cumulative voting for directors. *The Business Lawyer*, 10(3):3–16, 1954.
- [Chen and Toda, 1995] Z.Z. Chen and S. Toda. The complexity of selecting maximal solutions. *Information and Computation*, 119(2):231–239, 1995.
- [Dastani and van der Torre, 2002] Mehdi Dastani and Leendert van der Torre. Specifying the merging of desires into goals in the context of beliefs. In *Proceedings of the 1st EurAsian Conference on Information and Communication Technology (ICT-EurAsia)*, 2002.
- [de Haan and Slavkovik, 2017] Ronald de Haan and Marija Slavkovik. Complexity results for aggregating judgments using scoring or distance-based procedures. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 2017.
- [Dietrich and List, 2007] Franz Dietrich and Christian List. Judgment aggregation by quota rules: Majority voting generalized. *Journal of Theoretical Politics*, 19(4):391–424, 2007.
- [Dietrich and List, 2008] Franz Dietrich and Christian List. Judgment aggregation without full rationality. *Social Choice and Welfare*, 31(1):15–39, 2008.
- [Dokow and Holzman, 2010a] Elad Dokow and Ron Holzman. Aggregation of binary evaluations. *Journal of Economic Theory*, 145(2):495–511, 2010.
- [Dokow and Holzman, 2010b] Elad Dokow and Ron Holzman. Aggregation of binary evaluations with abstentions. *Journal of Economic Theory*, 145(2):544–561, 2010.
- [Elkind *et al.*, 2017] Edith Elkind, Martin Lackner, and Dominik Peters. Structured preferences. In Ulle Endriss, editor, *Trends in Computational Social Choice*, chapter 10, pages 187–207. AI Access, 2017.
- [Endriss *et al.*, 2012] Ulle Endriss, Umberto Grandi, and Daniele Porello. Complexity of judgment aggregation. *Journal of Artificial Intelligence Research*, 45:481–514, 2012.
- [Endriss, 2016] Ulle Endriss. Judgment aggregation. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 17, pages 399–426. Cambridge University Press, 2016.
- [Grandi and Endriss, 2011] Umberto Grandi and Ulle Endriss. Binary aggregation with integrity constraints. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, 2011.
- [Harrenstein *et al.*, 2001] Paul Harrenstein, Wiebe van der Hoek, John-Jules Meyer, and Cees Witteveen. Boolean games. In *Proceedings of the 8th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, 2001.
- [Klamler and Pferschy, 2007] Christian Klamler and Ulrich Pferschy. The traveling group problem. *Social Choice and Welfare*, 29(3):429–452, 2007.
- [Konieczny and Pérez, 2002] Sébastien Konieczny and Ramón Pino Pérez. Merging information under constraints: A logical framework. *Journal of Logic and Computation*, 12(5):773–808, 2002.
- [Lang and Slavkovik, 2014] Jérôme Lang and Marija Slavkovik. How hard is it to compute majority-preserving judgment aggregation rules? In *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI)*, 2014.
- [Lang and Xia, 2016] Jérôme Lang and Lirong Xia. Voting in combinatorial domains. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel Procaccia, editors, *Handbook of Computational Social Choice*, chapter 9, pages 197–222. Cambridge University Press, 2016.
- [Lang, 2004] Jérôme Lang. Logical preference representation and combinatorial vote. *Annals of Mathematics and Artificial Intelligence*, 42(1-3):37–71, 2004.
- [Laslier and Sanver, 2010] Jean-François Laslier and Remzi M. Sanver. *Handbook on Approval Voting*. Springer Science & Business Media, 2010.
- [List, 2012] Christian List. The theory of judgment aggregation: an introductory review. *Synthese*, 187(1):179–207, Jul 2012.
- [May, 1952] Kenneth O. May. A set of independent necessary and sufficient conditions for simple majority decision. *Econometrica: Journal of the Econometric Society*, pages 680–684, 1952.
- [Papadimitriou, 2003] Christos H. Papadimitriou. *Computational Complexity*. John Wiley and Sons Ltd., 2003.
- [Qing *et al.*, 2014] Ciyang Qing, Ulle Endriss, Raquel Fernández, and Justin Kruger. Empirical analysis of aggregation methods for collective annotation. In *Proceedings of the 25th International Conference on Computational Linguistics (COLING)*, 2014.
- [Shoham and Leyton-Brown, 2009] Yoav Shoham and Kevin Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, 2009.