



Manipulation in Majoritarian Goal-based Voting

Arianna Novaro, Umberto Grandi, Dominique Longin, Emiliano Lorini

► To cite this version:

Arianna Novaro, Umberto Grandi, Dominique Longin, Emiliano Lorini. Manipulation in Majoritarian Goal-based Voting. 13èmes Journées d'Intelligence Artificielle Fondamentale (JIAF 2019), Plate-Forme Intelligence Artificielle (PFIA 2019), Jul 2019, TOULOUSE, France. hal-02926121

HAL Id: hal-02926121

<https://hal.science/hal-02926121>

Submitted on 31 Aug 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Manipulation in Majoritarian Goal-based Voting

Arianna Novaro¹ Umberto Grandi¹
Dominique Longin² Emiliano Lorini²

¹ IRIT, University of Toulouse

² IRIT, CNRS, Toulouse

[arianna.novaro, umberto.grandi, longin, lorini]@irit.fr

Résumé

Dans le vote par buts les agents s'expriment sur des questions binaires grâce à des formules de logique propositionnelle. Les buts individuels sont agrégés par une fonction qui calcule la décision collective comme un ensemble d'évaluations. Avoir des agents motivés par des buts individuels amène naturellement à des situations de vote stratégique, où un agent peut obtenir un meilleur résultat en déclarant un but insincère. La majorité étant une des règles les plus connues utilisées pour prendre des décisions collectives, nous étudions trois de ses variantes dans le cadre du vote par buts. Nous étudions la manipulation pour ces règles en général, ainsi que pour un ensemble limité d'actions stratégiques ou des restrictions sur le langage des buts. Nous établissons aussi la complexité computationnelle pour qu'un agent puisse trouver une manipulation.

1 Introduction

A key aspect of agent-based architectures is endowing agents with goals [25], and propositional goals in particular are common in models of strategic reasoning. When taking collective decisions in a multi-issue domain, agents share the control over the variables at stake while still holding individual goals, as in the following example :

Example 1.1. *Three automated personal assistants need to arrange a business meal for their owners. They have to decide whether the restaurant should be fancy (F), if it should be in the center (C), and if they should meet for lunch (L) instead of dinner. Each owner gives to their assistant a propositional goal with respect to these issues. The goal of the first agent is that if they go to a restaurant in the suburbs, then they should have*

a casual lunch : $\gamma_1 = \neg C \rightarrow (\neg F \wedge L)$. The second agent wants that the meeting is either in the suburbs or casual, but not both : $\gamma_2 = \neg C \oplus \neg F$. The third agent wants a fancy dinner in the center : $\gamma_3 = F \wedge \neg L \wedge C$. The assistants want to ensure that the final decision satisfies their owners.

In Example 1.1 we need a procedure to make the autonomous agents reach a collective decision on every issue as precisely as possible : if the assistants returned a large set of possible options, the owners would ultimately have to make the choice they wanted to avoid. Secondly, strategic behavior needs to be taken into account, as agents are goal-oriented and nothing bounds them to be truthful. Two frameworks have been proposed in the literature on artificial intelligence to solve analogous problems : belief merging (see, e.g., Konieczny and Pino Pérez [17]) and goal-based voting [22]. Given our primary concern of resoluteness of the voting outcome we choose the latter.

A variety of functions could be studied to handle individual goals : our focus is on majoritarian rules. The appeal of majority lies not only in its intuitive definition and extensive application in real-world scenarios, but also on having been widely studied in the related fields of voting theory and judgment aggregation [20, 3]. However, when moving to goal-based voting many definitions of majority are possible. The three adaptations studied here strike a balance between different needs : that of providing a resolute result, and that of treating each issue independently while still considering the complex structure of propositional goals.

Each of these majoritarian goal-based voting rules will be analyzed with respect to their resistance to

several manipulation strategies. Negative results, i.e., finding that a rule *can* be manipulated, lead us to study the computational complexity of manipulation, as well as restricting the language of individual goals in the hope of discovering niches of strategy-proofness.

Related work. Our starting point is the work on voting in multi-issue domains with compactly represented preferences by Lang [18]. Propositional goals are such an example, linked to the literature on social choice with dichotomous preferences [7, 8]. A related preference language is that of CP-nets, in which preferences that are not necessarily dichotomous can be expressed [1]. The literature on combinatorial voting (see, e.g., the chapter by Lang and Xia [19]) provides solutions to tackle the combinatorial explosion entailed by the structure of the alternatives, such as voting sequentially over issues using tractable voting rules.

Closely related work is the study of strategy-proofness in judgment aggregation [4, 10], where the input is a complete binary choice over all issues rather than a propositional goal, as well as in belief merging [11], which focuses on a specific set of rules defined by axioms inspired from belief revision. The latter is the closest setting from a technical point of view, and in Section 2.5 we clarify the differences between the two models. Manipulation of voting rules has been amply studied in voting theory, starting from the seminal result of Gibbard and Satterthwaite [13, 24] to more recent studies aimed at finding barriers to manipulation (see, e.g., the survey by Conitzer and Walsh [2]).

Propositional goals in a strategic setting have been extensively studied in the literature on boolean games [15, 26]. Here, however, issues are not exclusively controlled by agents, since they express their goals using a common set of variables, a closer model being that of aggregation games [14].

Paper structure. In Section 2 we present the framework of majoritarian goal-based voting for different notions of resoluteness. Section 3 introduces manipulation, providing both theoretical and computational complexity results. Section 4 studies syntactic restrictions on the goal language and analyses strategy-proofness. Section 5 concludes.

2 Formal Framework

We start by presenting goal-based voting, as first introduced by Novaro *et al.* [22], focussing on three variants of issue-wise majority rule with varying degrees of resoluteness. We also include a detailed comparison with the related framework of belief merging [12].

2.1 Goal-Based Voting

A group of *agents*, represented by the finite set $\mathcal{N} = \{1, \dots, n\}$, has to take a collective decision over a number of *issues*, represented by the finite set $\mathcal{I} = \{1, \dots, m\}$ of propositional variables. We call *literal* L any atom $j \in \mathcal{I}$ or its negation $\neg j$. We let $\mathcal{L}_{\mathcal{I}}$ be the propositional language over the atoms in \mathcal{I} , with the usual boolean connectives. Each agent i expresses her *individual goal* by a consistent formula γ_i of $\mathcal{L}_{\mathcal{I}}$, as in Example 1.1. A *goal-profile* $\mathbf{\Gamma} = (\gamma_1, \dots, \gamma_n)$ collects the n agents' goals.

In Section 4 we will study restrictions on the language of goals, i.e., languages \mathcal{L}^{\star} for $\star \in \{\wedge, \vee, \oplus\}$, defined by the following BNF grammars :

$$\varphi := p \mid \neg p \mid \varphi \star \varphi$$

An *interpretation* (or *alternative*) is a function $v : \mathcal{I} \rightarrow \{0, 1\}$ associating a binary value to each variable in \mathcal{I} , where 0 means the issue is rejected and 1 that it is accepted. We assume that there is no integrity constraint, allowing all possible interpretations over the issues. We write $v \models \varphi$ if interpretation v makes φ true (i.e., v is a *model* of φ). The set $\text{Mod}(\varphi) = \{v \mid v \models \varphi\}$ contains all the models of φ . We denote agent i 's choices for issue j in the models of her goal as $v_i(j) = (m_{ij}^1, m_{ij}^0)$, with $m_{ij}^x = |\{v \in \text{Mod}(\gamma_i) \mid v(j) = x\}|$ for $x \in \{0, 1\}$. Abusing notation, we write $v_i(j) = x$ if $|\text{Mod}(\gamma_i)| = 1$ and $m_{ij}^x = 1$.

Numerous procedures can be used to turn individual goals into a group decision. For instance, in Example 1.1 a goal-based voting rule should provide a decision over type and location of the restaurant and the timing of the meal. A *goal-based voting rule* is formally defined as a function $F : (\mathcal{L}_{\mathcal{I}})^n \rightarrow \mathcal{P}(\{0, 1\}^m) \setminus \emptyset$ for any n and m . The input is a profile of n formulas, and the output is a non-empty set of alternatives. The number of acceptances and rejections of issue j in the outcome $F(\mathbf{\Gamma})$ is defined as $F(\mathbf{\Gamma})_j = (F(\mathbf{\Gamma})_j^0, F(\mathbf{\Gamma})_j^1)$, where $F(\mathbf{\Gamma})_j^x = |\{v \in F(\mathbf{\Gamma}) \mid v_j = x\}|$ for $x \in \{0, 1\}$. If $F(\mathbf{\Gamma})_j^x = 0$, we just write $F(\mathbf{\Gamma})_j = 1 - x$.

In this paper we study three generalizations of *majority*. Its definition in judgment aggregation [4, 9], where each agent i express a complete binary ballot B_i over all issues, is $\text{Maj}(\mathbf{B})_j = 1$ iff $\sum_{i \in \mathcal{N}} b_{ij} \geq \lceil \frac{n+1}{2} \rceil$ where $\mathbf{B} = (B_1, \dots, B_n)$.

2.2 Resolute Rules

We begin by presenting two definitions of *resolute* rules, always returning a unique model as their output. The first resolute variant of majority is *EMaj*, a quota rule accepting an issue if and only if more than half of

the total votes are in its favor :

$$EMaj(\mathbf{\Gamma})_j = 1 \text{ iff } \sum_{i \in \mathcal{N}} \left(\sum_{v \in \text{Mod}(\gamma_i)} \frac{v(j)}{|\text{Mod}(\gamma_i)|} \right) \geq \left\lceil \frac{n+1}{2} \right\rceil$$

As the goal formulas of the agents may have a varying number of models satisfying them, to guarantee their *equal* treatment *EMaj* gives a weight to each model of an agent's goal inversely proportional to the total number of models of her goal.

The second resolute variant of majority *2sMaj* proceeds in *two steps* by first applying *Maj* to the models of the agents' goals, and then to the result of the first aggregation step. We write *Maj*(*Mod*(γ_i)) for *Maj*(v_1, \dots, v_k) where *Mod*(γ_i) = $\{v_1, \dots, v_k\}$:

$$2sMaj(\mathbf{\Gamma}) = Maj(Maj(\text{Mod}(\gamma_1)), \dots, Maj(\text{Mod}(\gamma_n)))$$

Both *EMaj* and *2sMaj* are generalizations of *Maj* : they coincide with it when agents have goals in the form of complete conjunctions of literals.

While resolute rules help agents come to a unique decision, other desiderata of unbiasedness cannot be guaranteed at the same time. Consider the following two axioms, as defined by Novaro *et al.* [22].

Let $\varphi[j \mapsto k]$ for $j, k \in \mathcal{I}$ be the replacement of each occurrence of k by j in φ . A goal-based voting rule is *dual* if for all profiles $\mathbf{\Gamma}$, $F(\bar{\gamma}_1, \dots, \bar{\gamma}_n) = \{(1 - v(1), \dots, 1 - v(m)) \mid v \in F(\mathbf{\Gamma})\}$ where $\bar{\gamma} = \gamma[-1 \mapsto 1, \dots, -m \mapsto m]$. A rule F is *anonymous* if for any profile $\mathbf{\Gamma}$ and permutation $\sigma : \mathcal{N} \rightarrow \mathcal{N}$, we have that $F(\gamma_1, \dots, \gamma_n) = F(\gamma_{\sigma(1)}, \dots, \gamma_{\sigma(n)})$.

Unfortunately, these three desirable properties cannot be simultaneously satisfied, as shown by the following theorem :¹

Theorem 2.1. *There is no resolute rule F satisfying both anonymity and duality.*

Démonstration. Consider a rule F and suppose towards a contradiction that F is resolute, anonymous and dual. Take profile $\mathbf{\Gamma}$ for $\mathcal{N} = \{1, 2\}$ and $\mathcal{I} = \{1, 2\}$ where $\gamma_1 = 1 \wedge \neg 2$ and $\gamma_2 = \neg 1 \wedge 2$. By anonymity of F , for profile $\mathbf{\Gamma}' = (\gamma_2, \gamma_1)$ we have $F(\mathbf{\Gamma}) = F(\mathbf{\Gamma}')$. Since F is resolute, $F(\mathbf{\Gamma}) = \{(x, y)\}$, for $x, y \in \{0, 1\}$, and thus $F(\mathbf{\Gamma}') = \{(x, y)\}$. However, note that $\gamma_1 = \bar{\gamma}_2$ and $\gamma_2 = \bar{\gamma}_1$. Hence, $\mathbf{\Gamma}' = \bar{\mathbf{\Gamma}}$ and by duality we must have $F(\mathbf{\Gamma}') = \{(1 - x, 1 - y)\}$. Contradiction. \square

In the next section we will thus define a weaker but more attainable notion of resoluteness, which can be satisfied by anonymous and dual majoritarian rules.

1. A related result in social choice theory states that there exists no resolute, anonymous, and neutral voting procedure for 2 alternatives and an even number of voters (see, e.g., Moulin [21]).

2.3 Weakly Resolute Rules

We call a rule *weakly resolute* if the alternatives in its outcome either accept, reject or abstain on any of the issues. Formally, a rule F is weakly resolute if on every profile $\mathbf{\Gamma}$, $F(\mathbf{\Gamma}) = \text{Mod}(\varphi)$ for some conjunction $\varphi \in \mathcal{L}^\wedge$. We begin by showing that each goal-based voting rule that satisfies an axiom called independence by Novaro *et al.* [22] is weakly resolute. A rule F is *independent* if there are functions $f_j : \mathcal{D}_m^n \rightarrow \mathcal{C}$ for $j \in \mathcal{I}$ such that for all $\mathbf{\Gamma}$ we have $F(\mathbf{\Gamma}) = \Pi_{j \in \mathcal{I}} f_j(m_1(j), \dots, m_n(j))$, for $\mathcal{D}_m = \{(a, b) \mid a, b \in \mathbb{N} \text{ and } a + b \leq 2^m\}$ and $\mathcal{C} = \{\{0\}, \{1\}, \{0, 1\}\}$.

Theorem 2.2. *Each independent goal-based voting rule is weakly resolute.*

Démonstration. Consider an arbitrary $\mathbf{\Gamma}$ and the outcome of an independent rule $F(\mathbf{\Gamma})$. As F is independent, we have $F(\mathbf{\Gamma}) = \Pi_{j \in \mathcal{I}} f_j(m_1(j), \dots, m_n(j))$, where each $m_x(j) \in \{\{0\}, \{1\}, \{0, 1\}\}$. We want to show that F is weakly resolute. We construct a conjunction ψ as follows : for all $j \in \mathcal{I}$, if $f_j(m_1(j), \dots, m_n(j)) = \{0\}$ add conjunct $\neg j$ to ψ ; if $f_j(m_1(j), \dots, m_n(j)) = \{1\}$ add conjunct j to ψ ; if $f_j(m_1(j), \dots, m_n(j)) = \{0, 1\}$ skip. For all $v \in \text{Mod}(\psi)$ and for all $j \in \mathcal{I}$ appearing as conjuncts in ψ , we have $v(j) = 1$ for a positive literal j , and $v(j) = 0$ for a negative literal $\neg j$. Moreover, for all $k \in \mathcal{I}$ which did not appear in ψ we have any possible combination of truth values. Therefore, $\text{Mod}(\psi) = F(\mathbf{\Gamma})$. \square

The other direction of Theorem 2.2 does not hold : consider a rule F that returns $\{(1, 1, \dots, 1)\}$ if in at least one $v \in \text{Mod}(\gamma_1)$ issue 1 is true, and else it returns $\{(0, 0, \dots, 0)\}$.

Novaro *et al.* [22] showed that the only rule satisfying independence, as well as a number of other desirable properties, is the following :

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

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

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

TrueMaj compares issue-by-issue the total acceptances with the total rejections, setting the result to 1 (respectively, 0) if higher (respectively, lower), and to both 0 and 1 if tied.

2.4 Implementation of Majoritarian Rules

To compare the three definitions of majoritarian goal-based voting, we wrote a program to compute their outcome over all 16581375 profiles for 3 agents

and 3 issues : i.e., for any combination of consistent goals and by considering γ_1 and γ_2 equivalent if $\text{Mod}(\gamma_1) = \text{Mod}(\gamma_2)$. Results are shown in Table 1. *TrueMajR* is a resolute version of *TrueMaj* picking a single random alternative in the outcome if multiple are present (we show an average over 10 executions). First, we evaluated the maximization of *social welfare* : i.e., the percentage of profiles on which the outcome satisfies all agents' goals. For *TrueMaj* is on almost 50% of the profiles (40% for the random tie-breaking version), going down to 30% for *2sMaj* and 21% for *EMaj*. Then, we checked the percentage of profiles for which the rules return an outcome that does not satisfy *any* goal. This happens for less than 5% of profiles for all rules, and for less than 0,5% for *TrueMaj*. Finally, we analyzed the percentage of profiles on which agent 1 is satisfied with the result — if they are not satisfied they may have a strategy to manipulate. Agent 1 is satisfied on more than 61% of profiles for *EMaj*, almost 70% for *2sMaj* and almost 80% for *TrueMaj* (75 % for its resolute version). While preliminary, these results provide an overall picture on the performance of majoritarian rules, with *TrueMaj* standing out even when coupled with random tie-breaking.

2.5 Goal-based Voting and Belief Merging

Belief merging was proposed and widely studied as a framework to combine the beliefs of multiple agents [17]. While belief merging rules and axioms differ from those proposed in goal-based voting, both settings are concerned with the problem of combining formulas into sets of interpretations. In what follows we refer to the formulation of postulates by Everaere *et al.* [12].

We begin by observing that the (IC2) postulate, which states that the outcome of a rule should coincide with the conjunction of the goals if they are consistent, is incompatible with both resoluteness, in case $|\text{Mod}(\bigwedge_{i \in \mathcal{N}} \gamma_i)| > 1$, and weak resoluteness, e.g., when $\bigwedge_{i \in \mathcal{N}} \gamma_i = 1 \vee 2$. Our primary concern in resoluteness lead us to choose goal-based voting instead.

As for the other postulates, since in goal-based voting the integrity constraint is absent and $F(\mathbf{\Gamma}) \neq \emptyset$ on all $\mathbf{\Gamma}$, (IC0) and (IC1) are satisfied by default. The principle of irrelevance of syntax (IC3) is implicitly satisfied by goal-based voting rules, as propositional logic is simply used for the compact representation of goals. Postulate (IC4), defined for two agents, is not satisfied by neither *EMaj*, *TrueMaj* nor *2sMaj* : consider a profile $\mathbf{\Gamma}$ for two agents and three issues such that $\gamma_1 = \neg 1 \wedge \neg 2 \wedge \neg 3$ and $\gamma_2 = (1 \wedge \neg 2 \wedge \neg 3) \vee (\neg 1 \wedge 2 \wedge \neg 3) \vee (\neg 1 \wedge \neg 2 \wedge 3)$. We have that $EMaj(\mathbf{\Gamma}) = TrueMaj(\mathbf{\Gamma}) = 2sMaj(\mathbf{\Gamma}) = \{(000)\}$: the outcome is thus only consistent with the goal of agent 1 and not with that of agent 2. Postulates

(IC5) and (IC6) are known in the literature on social choice theory as *reinforcement* [27], which is satisfied by all three majoritarian rules proposed (where $\mathbf{\Gamma} \sqcup \mathbf{\Gamma}' = (\gamma_1, \dots, \gamma_n, \gamma'_1, \dots, \gamma'_n)$ indicates the union of profiles $\mathbf{\Gamma}$ and $\mathbf{\Gamma}'$) :

Theorem 2.3. *For any $\mathbf{\Gamma}$ and $\mathbf{\Gamma}'$, $EMaj$, $2sMaj$ and $TrueMaj$ satisfy $F(\mathbf{\Gamma}) \cap F(\mathbf{\Gamma}') = S \neq \emptyset$ if and only if $F(\mathbf{\Gamma} \sqcup \mathbf{\Gamma}') = S$.*

Démonstration. Consider two arbitrary profiles $\mathbf{\Gamma}$ and $\mathbf{\Gamma}'$. Let $EMaj(\mathbf{\Gamma}) = EMaj(\mathbf{\Gamma}') = \{w\}$. For all $j \in \mathcal{I}$: if $w(j) = 1$, then there were more than $\lceil \frac{n+1}{2} \rceil$ votes for j in both $\mathbf{\Gamma}$ and $\mathbf{\Gamma}'$ (and consequently in $\mathbf{\Gamma} \sqcup \mathbf{\Gamma}'$) ; if $w(j) = 0$, then in $\mathbf{\Gamma}$ and $\mathbf{\Gamma}'$ either there was a tie for j or there were less than $\lceil \frac{n+1}{2} \rceil$ votes for j . Any combination of ties or $< \lceil \frac{n+1}{2} \rceil$ votes for j in $\mathbf{\Gamma}$ and $\mathbf{\Gamma}'$ still leads to $EMaj(\mathbf{\Gamma} \sqcup \mathbf{\Gamma}')_j = 0$. For $2sMaj$ is as for *EMaj*, focusing on the second step only. Let $TrueMaj(\mathbf{\Gamma}) \cap TrueMaj(\mathbf{\Gamma}') = S$. For all $j \in \mathcal{I}$: if there are $w, w' \in S$ such that $w(j) = 1$ and $w'(j) = 0$, then $\mathbf{\Gamma}$ and $\mathbf{\Gamma}'$ had a tie in the votes for j and thus a tie will be in $\mathbf{\Gamma} \sqcup \mathbf{\Gamma}'$ (hence in the outcome). If $w(j) = 1$ for all $w \in S$ (analogously for 0), there may have been a tie in either $\mathbf{\Gamma}'$ or $\mathbf{\Gamma}$ for j , but not both, and so $\mathbf{\Gamma} \sqcup \mathbf{\Gamma}'$ will have no tie for j . \square

Both (IC7) and (IC8) are not applicable to goal-based voting as there is no integrity constraint. Finally, the belief merging postulate (Maj), is not satisfied by *TrueMaj*. Consider goals $\gamma_1 = 1 \wedge 2$ and $\gamma_2 = 1 \oplus 2$: no matter how many times γ_2 is repeated in a profile, in the presence of γ_1 the result will always be $\{(11)\}$.

3 Manipulation : Theory and Complexity

Propositional goals lead to a dichotomous preference relation on alternatives : agents equally prefer any model of their goal to any counter-model. For resolute rules, the *unique* result satisfies an agent if and only if it is a model of their goal. Otherwise, different notions of satisfaction arise depending on how an agent compares two *sets* of interpretations.

Let $sat : \mathcal{L}_{\mathcal{I}} \times (\mathcal{P}(\{0, 1\}^m) \setminus \emptyset) \rightarrow [0, 1]$ be a function expressing the *satisfaction* of agent i towards the outcome of a rule F on profile $\mathbf{\Gamma}$. We simply write $sat(i, F(\mathbf{\Gamma}))$ instead of $sat(\gamma_i, F(\mathbf{\Gamma}))$. The *optimistic*, *pessimistic* and *expected utility maximizer* are three

	<i>EMaj</i>	<i>TrueMaj</i>	<i>TrueMajR</i>	<i>2sMaj</i>
% all agents sat	21,47	48,32	39,77	30,38
% no agent sat	4,91	0,47	0,76	2,18
% agent 1 sat	61,13	79,30	75,19	69,03

Table 1 : For each rule we show the percentage of profiles where all agents/no agent/agent 1 are satisfied with the result. Satisfaction is optimistic.

notions of satisfaction an agent may hold :

$$\begin{aligned}
opt(i, F(\mathbf{\Gamma})) &= \begin{cases} 1 & \text{if } F(\mathbf{\Gamma}) \cap \text{Mod}(\gamma_i) \neq \emptyset \\ 0 & \text{otherwise} \end{cases} \\
pess(i, F(\mathbf{\Gamma})) &= \begin{cases} 1 & \text{if } F(\mathbf{\Gamma}) \subseteq \text{Mod}(\gamma_i) \\ 0 & \text{otherwise} \end{cases} \\
eum(i, F(\mathbf{\Gamma})) &= \frac{|\text{Mod}(\gamma_i) \cap F(\mathbf{\Gamma})|}{|F(\mathbf{\Gamma})|}
\end{aligned}$$

Optimists are satisfied if in the outcome there is at least one model of their goal. Pessimists want all the interpretations in the outcome to be models of their goal (this notion was introduced by Jimeno *et al.* [16]). Expected utility maximizers assume that a unique interpretation will be chosen at random among those tied in the outcome, and the higher the proportion of models of their goal in $F(\mathbf{\Gamma})$ over the total number of interpretations in $F(\mathbf{\Gamma})$, the better.²

Agent i 's preference on outcomes is a complete and transitive relation \succsim_i , whose strict part is \succ_i :

$$F(\mathbf{\Gamma}) \succsim_i F(\mathbf{\Gamma}') \text{ iff } sat(i, F(\mathbf{\Gamma})) \geq sat(i, F(\mathbf{\Gamma}')).$$

For $\mathbf{\Gamma} = (\gamma_i)_{i \in \mathcal{N}}$, let $(\mathbf{\Gamma}_{-i}, \gamma'_i) = (\gamma_1, \dots, \gamma'_i, \dots, \gamma_n)$ be the profile where only agent i changed her goal from γ_i to γ'_i . Agent i has an *incentive* to manipulate by submitting goal γ'_i instead of γ_i if and only if $F(\mathbf{\Gamma}_{-i}, \gamma'_i) \succ_i F(\mathbf{\Gamma})$. A rule F is *strategy-proof* if and only if for all profiles $\mathbf{\Gamma}$ there is no agent i who has an incentive to manipulate.

We focus on three kind of manipulation strategies, following previous work by Everaere *et al.* [11] :

- *Unrestricted* : i can send any γ'_i instead of γ_i
- *Erosion* : i can send γ'_i s.t. $\text{Mod}(\gamma'_i) \subseteq \text{Mod}(\gamma_i)$
- *Dilatation* : i can send γ'_i s.t. $\text{Mod}(\gamma_i) \subseteq \text{Mod}(\gamma'_i)$

3.1 Manipulability of Majority Rules

In judgment aggregation, the issue-by-issue majority rule has been proven to be single-agent strategy-proof by Dietrich and List [4]. Surprisingly, when moving to propositional goals strategy-proofness is not

2. Expected utility maximizers, optimists and pessimists, correspond to the *probabilistic*, *weak drastic* and *strong drastic* satisfaction indexes in the work of Everaere *et al.* [11].

guaranteed anymore for the three adaptations of the majority rule, as shown by the following result :

Theorem 3.1. *EMaj, TrueMaj and 2sMaj can be manipulated by both erosion and dilatation.*

Démonstration. We provide goal-profiles where an agent can get a better result by submitting an untruthful goal. For ease of presentation we display the models of the agents' goals, but the input of a rule F consists of propositional formulas. Consider the profiles $\mathbf{\Gamma}$, $\mathbf{\Gamma}'$ and $\mathbf{\Gamma}''$ for three agents and three issues, together with the results of *EMaj*, *TrueMaj* and *2sMaj* :

	$\mathbf{\Gamma}$	$\mathbf{\Gamma}'$	$\mathbf{\Gamma}''$	$\mathbf{\Gamma}^*$	$\mathbf{\Gamma}^{**}$
$\text{Mod}(\gamma_1)$	(111)	(111)	(111)	(111)	(111)
$\text{Mod}(\gamma_2)$	(001)	(001)	(001)	(111) (011) (100)	(111) (011) (100)
$\text{Mod}(\gamma_3)$	(101) (010) (000)	(101)	(101) (010) (000) (100) (110)	(011) (010) (101)	(011) (010) (101) (001)
<i>EMaj</i>	(001)	(101)	(001)	(111)	(011)
<i>TrueMaj</i>	(001)	(101)	(101)	-	-
<i>2sMaj</i>	(001)	(101)	(101)	-	-

Let $\mathbf{\Gamma}$ be the profile where agents submit their truthful goal : $\gamma_1 = 1 \wedge 2 \wedge 3$, $\gamma_2 = \neg 1 \wedge \neg 2 \wedge 3$, $\gamma_3 = (\neg 1 \wedge \neg 3) \vee (1 \wedge \neg 2 \wedge 3)$. For erosion manipulation, agent 3 prefers the result of *EMaj*, *2sMaj* and *TrueMaj* (which they happen to coincide) when applied to $\mathbf{\Gamma}'$ rather than when applied to $\mathbf{\Gamma}$. For dilatation manipulation, agent 3 prefers the result of *TrueMaj* and *2sMaj* when applied to $\mathbf{\Gamma}''$ rather than to $\mathbf{\Gamma}$. For *EMaj* and dilatation manipulation, agent 3 can get a better result by manipulating $\mathbf{\Gamma}^*$ and moving to $\mathbf{\Gamma}^{**}$. \square

Theorem 3.1 is thus in sharp contrast with the result of judgment aggregation. Since the profiles used in the proof give singleton outcomes, the theorem holds for expected utility maximizers, optimists and pessimists.

3.2 Computational Complexity

Majoritarian goal-based voting rules are manipulable, as shown by Theorem 3.1, but how difficult it is for an agent to find another goal allowing her to get a better outcome? If we restrict to resolute rules, the problem definition is analogous to existing work in judgment aggregation [10] :

MANIP(F)

Input : Profile $\mathbf{\Gamma} = (\gamma_1, \dots, \gamma_n)$, agent i

Question : If $\text{Mod}(\gamma_i) \cap F(\mathbf{\Gamma}) = \emptyset$, is there γ'_i such that $\text{Mod}(\gamma'_i) \cap F(\gamma_1, \dots, \gamma'_i, \dots, \gamma_n) \neq \emptyset$?

Let Probabilistic Polynomial Time (PP) be the class of problems that can be solved in nondeterministic polynomial time with acceptance condition that more than half of the computations accept [23]. In what follows we refer to the PP-complete problem MAJ-SAT- p , asking whether $|\text{Mod}(\varphi \wedge p)| > |\text{Mod}(\varphi \wedge \neg p)|$ for propositional formula φ and one of its variables p [22].

Theorem 3.2. MANIP($2sMaj$) is PP-hard.

Démonstration. We reduce from MAJ-SAT- p . Take an instance of MAJ-SAT- p with a formula $\varphi[p_1, \dots, p_k]$ and p_1 one of its variables. Construct an instance of MANIP($2sMaj$)³ where $\mathcal{I} = \{p_1, \dots, p_k, q, r\}$, and a profile $\mathbf{\Gamma}$ with $\gamma_1 = (p_2 \wedge \dots \wedge p_k) \wedge p_1 \wedge q \wedge r$, and $\gamma_2 = \varphi \wedge \neg q \wedge \neg r$, and $\gamma_3 = (p_2 \wedge \dots \wedge p_k) \wedge (p_1 \oplus q) \wedge (r \rightarrow q)$.

We show that $|\text{Mod}(\varphi \wedge p_1)| > |\text{Mod}(\varphi \wedge \neg p_1)|$ if and only if agent 3 can manipulate $2sMaj$ on $\mathbf{\Gamma}$. The following table represents some features of $\mathbf{\Gamma}$, where question marks represent the possibly many models of φ over p_1, \dots, p_k :

	$p_2 \dots p_k$	p_1	q	r
Mod(γ_1)	$1 \dots 1$	1	1	1
	?	?	0	0
Mod(γ_2)	\vdots	\vdots	0	0
	?	?	0	0
Mod(γ_3)	$1 \dots 1$	0	1	1
	$1 \dots 1$	0	1	0
	$1 \dots 1$	1	0	0

The result on p_2, \dots, p_k is decided by agents 1 and 3 : all issues will be accepted regardless of the vote of agent 2.

Let us now focus on p_1, q and r . Applying (strict) majority to the models of γ_3 leads to the first-step result (010). Agent 3 is pivotal on issues q and r (since agents 1 and 2 will give one vote for and one vote against them after the first step). There are now two cases to consider :

3. For ease of presentation and to avoid confusion we write the issues as $\mathcal{I} = \{p, q, r, \dots\}$ instead of $\mathcal{I} = \{1, 2, 3, \dots\}$.

- If $|\text{Mod}(\varphi \wedge p_1)| > |\text{Mod}(\varphi \wedge \neg p_1)|$, the result of $2sMaj(\mathbf{\Gamma})$ is $(1 \dots 1, 110)$, that is not a model of γ_3 . However, by submitting $\gamma'_3 = (p_2 \wedge p_k) \wedge p_1 \wedge \neg q \wedge \neg r$, we have $2sMaj(\mathbf{\Gamma}') = (1 \dots 1, 100)$ which is a model of γ_3 . Hence, agent 3 has an incentive to manipulate.
- If $|\text{Mod}(\varphi \wedge p_1)| \leq |\text{Mod}(\varphi \wedge \neg p_1)|$, we have that $2sMaj(\mathbf{\Gamma}) = (1 \dots 1, 010)$, which is a model of γ_3 . Agent 3 has thus no incentive to manipulate.

This completes the reduction from MAJ-SAT- p , showing that MANIP($2sMaj$) is PP-hard. \square

A similar reduction allows us to show the following :

Theorem 3.3. MANIP($EMaj$) is PP-hard.

Démonstration. We construct an instance of MANIP($EMaj$) from a given instance (φ, p_1) of MAJ-SAT- p . Let $\mathcal{I} = \{p_1, \dots, p_k, q, r\}$, and profile $\mathbf{\Gamma} = (\gamma_1, \gamma_2, \gamma_3)$ with $\gamma_1 = (p_2 \wedge \dots \wedge p_k) \wedge p_1 \wedge q \wedge r$, $\gamma_2 = \varphi(\overline{p_1}) \wedge \neg q \wedge r$ and $\gamma_3 = (p_2 \wedge \dots \wedge p_k) \wedge (p_1 \oplus q)$, where $\varphi(\overline{p_1})$ is a formula where each occurrence of p_1 in φ has been replaced by $\neg p_1$ and vice-versa. The final part of the proof, showing that $|\text{Mod}(\varphi \wedge p_1)| > |\text{Mod}(\varphi \wedge \neg p_1)|$ if and only if agent 3 can manipulate $EMaj$ on $\mathbf{\Gamma}$, is analogous to that of Theorem 3.2 and thus omitted. \square

We leave the extension of the above results to *TrueMaj* for each attitude of the agent (optimist, utility maximizer, or pessimist) for future work, conjecturing it to be PP-hard as well.

4 Language Restrictions

We study three restrictions on the goal-language : conjunctions, disjunctions and exclusive disjunctions, as defined by the corresponding languages \mathcal{L}^* defined in Section 2.1. Results are presented in Table 2.

4.1 Conjunctions

The formulas of the *language of conjunctions* \mathcal{L}^\wedge are conjunctions of literals over issues in \mathcal{I} . \mathcal{L}^\wedge captures the framework of judgment aggregation with abstentions [5, 6], as agents have definite opinions over the issues appearing as literals in their goal and they do not care about the other issues. We find positive results of strategy-proofness :

Theorem 4.1. An agent i with $\gamma_i \in \mathcal{L}^\wedge$ has no incentive to manipulate unrestrictedly $2sMaj$ and $EMaj$.

Démonstration. Take $\mathbf{\Gamma}$ with $\gamma_i \in \mathcal{L}^\wedge$. Since $2sMaj$ and $EMaj$ are resolute, we have a unique outcome $\{w\}$ on $\mathbf{\Gamma}$. Suppose that $w \notin \text{Mod}(\gamma_i)$. As $\gamma_i = L_1 \wedge \dots \wedge L_k$ for $k \leq m$, we have for all $j \in \mathcal{I}$:

- $v_i(j) = (\frac{|\text{Mod}(\gamma_i)|}{2}, \frac{|\text{Mod}(\gamma_i)|}{2})$ if γ_i has no L_j .
- If $L_j = j$ appears in γ_i , then $v_i(j) = (|\text{Mod}(\gamma_i)|, 0)$, and if $L_j = \neg j$, then $v_i(j) = (0, |\text{Mod}(\gamma_i)|)$.

Therefore, if $w \notin \text{Mod}(\gamma_i)$ there must be ℓ literals with $\ell \leq k$ such that $w \models \neg L_1 \wedge \dots \wedge \neg L_\ell$. Take an arbitrary such L_x .

2sMaj. Let $\text{Maj}(\gamma_i) = \{w_i\}$ be the result of the first step of majority applied to γ_i . We have that $w_i(x) = 1 - w(x)$, and therefore agent i cannot influence the outcome towards $w_i(x)$, and more generally towards her goal.

EMaj. If $w(x) = 1$ (similarly for 0), then $L_x = \neg x$. Since $v_i(x) = (0, |\text{Mod}(\gamma_i)|)$, we have $\sum_{v \in \text{Mod}(\gamma_i)} \frac{v(x)}{|\text{Mod}(\gamma_i)|} = 0$ and thus $\sum_{k \in \mathcal{N} \setminus \{i\}} \sum_{v \in \text{Mod}(\gamma_k)} \frac{v_k(x)}{|\text{Mod}(\gamma_k)|} \geq \lceil \frac{n+1}{2} \rceil$. Agent i is already giving no support to x and yet x is accepted in the outcome. Therefore, *EMaj* cannot be manipulated. \square

The result for *TrueMaj* has a similar proof, which is omitted for space constraints, also considering optimists, pessimists and expected utility maximizers.

Theorem 4.2. *If agent i has $\gamma_i \in \mathcal{L}^\wedge$ she has no incentive to manipulate unrestrictedly the rule *TrueMaj*.*

A consequence of Theorem 4.1 and 4.2 is that goals in \mathcal{L}^\wedge make the three majorities strategy-proof :

Corollary 4.1. *For any \mathcal{I} and \mathcal{N} , if $\gamma_i \in \mathcal{L}^\wedge$ for all $i \in \mathcal{N}$ then *EMaj*, *TrueMaj* and *2sMaj* are strategy-proof for unrestricted manipulation.*

4.2 Disjunctions

In the language of clauses \mathcal{L}^\vee , formulas are disjunctions of literals. Unfortunately, this restriction does not guarantee strategy-proofness for two of our rules :

Theorem 4.3. *There exists Γ and $i \in \mathcal{N}$ with $\gamma_i \in \mathcal{L}^\vee$ such that agent i has an incentive to manipulate *EMaj* and *TrueMaj* by erosion.*

Démonstration. Profiles are for three agents and two issues. For *EMaj*, consider Γ with $\gamma_1 = 1$, $\gamma_2 = \neg 1$ and $\gamma_3 = 1 \vee 2$. By submitting $\gamma'_3 = 1 \wedge 2$ agent 3 can change $\text{EMaj}(\Gamma) = \{(00)\}$ into $\{(11)\}$: hence, they have an incentive to lie. For *TrueMaj*, consider Γ with $\gamma_1 = \neg 1 \wedge \neg 2$ and $\gamma_2 = \gamma_3 = 1 \vee 2$. The result is $\text{TrueMaj}(\Gamma) = \{(00)\}$, but if agent 3 submits $\gamma'_3 = 1 \wedge 2$ we get $\text{TrueMaj}(\Gamma') = \{(11)\}$ and thus agent 3 has an incentive to manipulate. \square

We can obtain positive results if we restrict the set of available manipulation strategies to dilatation :

Theorem 4.4. *For any Γ with $\gamma_i \in \mathcal{L}^\vee$ for $i \in \mathcal{N}$, agent i has no incentive to manipulate *EMaj* and *TrueMaj* by dilatation.*

Proof sketch. Take Γ with $\gamma_i \in \mathcal{L}^\vee$ for some $i \in \mathcal{N}$, i.e., $\gamma_i = L_1 \vee \dots \vee L_k$. Consider $\text{EMaj}(\Gamma) = \{w\}$, such that $w \notin \text{Mod}(\gamma_i)$. As the agent is restricted to dilatation strategies, the goals γ'_i such that $\text{Mod}(\gamma_i) \subseteq \text{Mod}(\gamma'_i)$ agent i can use are those whose models would lower i 's support for each literal L_ℓ in γ_i . Thus she cannot manipulate.

Let now $\text{TrueMaj}(\Gamma) = \{w_1, \dots, w_k\}$. Reasoning as above, an optimist agent i does not have any dilatation strategy to include one of the models of γ_i into the outcome. An expected utility maximizer i however, may want to remove some $w_k \in \text{TrueMaj}(\Gamma)$ such that $w_k \models \neg L_j$ for all L_j in γ_i . But this is only possible when $\{w_1, \dots, w_\ell\} \cap \text{Mod}(\gamma_i) = \emptyset$, since the agent is restricted to dilatation strategies, and other reported goal would have more models increasing the votes against the literals present in her sincere goal γ_i . Therefore, this would not constitute a profitable deviation. A similar reasoning applies to pessimists. \square

Theorem 4.5. *For any profile where $\gamma_i \in \mathcal{L}^\vee$ for $i \in \mathcal{N}$, agent i has no incentive to manipulate unrestrictedly *2sMaj*.*

Démonstration. Since $\gamma_i \in \mathcal{L}^\vee$, the result w_i of $\text{Maj}(\text{Mod}(\gamma_i))$ is such that $w_i(x) = 1$ if $L_x = x$ appears in γ_i , and $w_i(x) = 0$ if $L_x = \neg x$. Hence, it coincides with the result of $\text{Maj}(\text{Mod}(\gamma'_i))$ for $\gamma'_i \in \mathcal{L}^\wedge$ where every occurrence of \vee in γ_i has been replaced by \wedge in γ'_i . The proof of Theorem 4.1 can thus be applied, since agent i is already maximizing her chances of getting γ_i satisfied by submitting γ_i . \square

By combining the results of Theorem 4.4 and Theorem 4.5 we get the following corollary :

Corollary 4.2. *If $\gamma_i \in \mathcal{L}^\vee$ for all $i \in \mathcal{N}$ then *EMaj* and *TrueMaj* are strategy-proof for dilatation manipulation and *2sMaj* is strategy-proof for unrestricted manipulation.*

4.3 Exclusive Disjunctions

In the language of exclusive disjunctions \mathcal{L}^\oplus , each formula is an exclusive disjunction of literals (cf. agent 2's goal in Example 1.1). We prove the following :

Theorem 4.6. *There exists profile $\Gamma^0, \Gamma^1, \Gamma^2, \Gamma^3$, and $i \in \mathcal{N}$ with $\gamma_i \in \mathcal{L}^\oplus$ such that agent i has an incentive to manipulate rules *2sMaj*, *EMaj* and *TrueMaj*, by erosion and dilatation.*

	\mathcal{L}^\wedge		\mathcal{L}^\vee		\mathcal{L}^\oplus	
	E	D	E	D	E	D
<i>EMaj</i>	SP	SP	M	SP	M	M
<i>TrueMaj</i>	SP	SP	M	SP	M	M
<i>2sMaj</i>	SP	SP	SP	SP	M	M

Table 2 : *E* stands for erosion, *D* for dilatation, *SP* for strategy-proof and *M* for manipulable.

Démonstration. All profiles are for three agents and two issues.

2sMaj. For erosion, consider profile $\mathbf{\Gamma}^0$ where $\gamma_1 = 1 \wedge 2$, $\gamma_2 = \neg 1 \wedge \neg 2$ and $\gamma_3 = 1 \oplus 2$. We have that $2sMaj(\mathbf{\Gamma}) = \{(00)\}$. Consider now $\gamma'_3 = \neg 1 \wedge 2$. The result is $\{(01)\}$, and thus agent 3 has an incentive to manipulate. For dilatation, consider the profile $\mathbf{\Gamma}^1$ where $\gamma_1 = \neg 1 \wedge 2$, $\gamma_2 = \neg 1 \wedge \neg 2$ and $\gamma_3 = 1 \oplus 2$. We have that $2sMaj(\mathbf{\Gamma}^1) = \{(00)\}$. If we consider $\gamma_3^* = 1 \vee 2$, the result is $2sMaj(\mathbf{\Gamma}^{1*}) = \{(01)\}$, and thus agent 3 has again an incentive to manipulate.

EMaj. For erosion, the results of *2sMaj* and *EMaj* coincide on $\mathbf{\Gamma}^0$ and $\mathbf{\Gamma}'$. For dilatation, take $\mathbf{\Gamma}^2$ with $\gamma_1 = \neg 1 \wedge 2$, $\gamma_2 = \neg 1 \vee \neg 2$ and $\gamma_3 = 1 \oplus 2$. We have $EMaj(\mathbf{\Gamma}^2) = \{(00)\}$. Agent 3 can submit $\gamma_3^* = 1 \vee 2$ to obtain $\{(01)\}$.

TrueMaj. Take $\mathbf{\Gamma}^3$ with $\gamma_1 = 1 \wedge 2$, $\gamma_2 = 1 \wedge \neg 2$ and $\gamma_3 = 1 \oplus 2$. We have $TrueMaj(\mathbf{\Gamma}^3) = \{(10), (11)\}$. Agent 3 can manipulate by erosion with $\gamma_3^* = 1 \wedge \neg 2$, and by dilatation with $\gamma_3^{**} = \neg 1 \vee \neg 2$. In both cases the result is $\{(10)\}$. \square

5 Conclusions

In this paper we studied the strategic component of the framework of goal-based voting [22], related (yet different) to both judgment aggregation and belief merging. Our focus was on three rules that have been proposed as adaptations of the issue-wise majority rule in this setting, with varying degrees of resoluteness. We find that all the majoritarian rules are not immune from manipulation, even when the manipulator can only apply limited strategies on their truthful goal (i.e., erosion and dilatation). We also find that, although not strategy-proof in general, *EMaj* and *2sMaj* are PP-hard for an agent to manipulate, as hard as their winner determination problem. Moreover, restricting the language of an agent's goal to conjunctions makes manipulation impossible, as well as dilatation manipulation for the language of disjunctions, suggesting promising directions for further research on minimal restrictions to the goal language to guarantee strategy-proofness of majoritarian rules.

Références

- [1] Boutilier, Craig, Ronen I. Brafman, Carmel Domshlak, Holger H. Hoos et David Poole: *CP-nets : A Tool for Representing and Reasoning with Conditional Ceteris Paribus Preference Statements*. Journal of Artificial Intelligence Research, 21 :135–191, 2004.
- [2] Conitzer, Vincent et Toby Walsh: *Barriers to Manipulation in Voting*. Dans Brandt, F., V. Conitzer, U. Endriss, J. Lang et A. D. Procaccia (éditeurs) : *Handbook of Computational Social Choice*, chapitre 6. Cambridge University Press, 2016.
- [3] Dietrich, Franz et Christian List: *Judgment aggregation by quota rules : Majority voting generalized*. Journal of Theoretical Politics, 19(4) :391–424, 2007.
- [4] Dietrich, Franz et Christian List: *Strategy-proof judgment aggregation*. Economics & Philosophy, 23(3) :269–300, 2007.
- [5] Dietrich, Franz et Christian List: *Judgment Aggregation Without Full Rationality*. Social Choice and Welfare, 31(1) :15–39, 2008.
- [6] Dokow, Elad et Ron Holzman: *Aggregation of Binary Evaluations with Abstentions*. Journal of Economic Theory, 145(2) :544–561, 2010.
- [7] Elkind, Edith et Martin Lackner: *Structure in Dichotomous Preferences*. Dans *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, 2015.
- [8] Elkind, Edith, Martin Lackner et Dominik Peters: *Structured Preferences*. Dans Endriss, Ulle (éditeur) : *Trends in Computational Social Choice*, chapitre 10, pages 187–207. AI Access, 2017.
- [9] Endriss, Ulle: *Judgment Aggregation*. Dans Brandt, F., V. Conitzer, U. Endriss, J. Lang et A. D. Procaccia (éditeurs) : *Handbook of Computational Social Choice*, chapitre 17. Cambridge University Press, 2016.
- [10] Endriss, Ulle, Umberto Grandi et Daniele Porello: *Complexity of Judgment Aggregation*. J. of Artificial Intelligence Research, 45 :481–514, 2012.
- [11] Everaere, Patricia, Sébastien Konieczny et Pierre Marquis: *The Strategy-Proofness Landscape of Merging*. Journal of Artificial Intelligence Research, 28 :49–105, 2007.
- [12] Everaere, Patricia, Sébastien Konieczny et Pierre Marquis: *An Introduction to Belief Merging and its Links with Judgment Aggregation*. Dans Endriss, Ulle (éditeur) : *Trends in Computational Social Choice*, chapitre 7, pages 123–143. AI Access, 2017.

- [13] Gibbard, Allan: *Manipulation of voting schemes : a general result*. *Econometrica*, pages 587–601, 1973.
- [14] Grandi, Umberto, Davide Grossi et Paolo Turini: *Equilibrium Refinement through Negotiation in Binary Voting*. Dans *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 540–546, 2015.
- [15] Harrenstein, Paul, Wiebe van der Hoek, John-Jules Ch. Meyer et Cees Witteveen: *Boolean games*. Dans *Proceeding of the 8th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, 2001.
- [16] Jimeno, José L., Joaquín Pérez et Estefanía García: *An extension of the Moulin No Show Paradox for voting correspondences*. *Social Choice and Welfare*, 33(3) :343–359, 2009.
- [17] Konieczny, Sébastien et Ramón Pino Pérez: *Merging Information Under Constraints : A Logical Framework*. *Journal of Logic and Computation*, 12(5) :773–808, 2002.
- [18] Lang, Jérôme: *Logical Preference Representation and Combinatorial Vote*. *Annals of Mathematics and Artificial Intelligence*, 42(1-3) :37—71, 2004.
- [19] Lang, Jérôme et Lirong Xia: *Voting in Combinatorial Domains*. Dans Brandt, F., V. Conitzer, U. Endriss, J. Lang et A. D. Procaccia (éditeurs) : *Handbook of Computational Social Choice*, chapitre 9. Cambridge University Press, 2016.
- [20] May, Kenneth O: *A set of independent necessary and sufficient conditions for simple majority decision*. *Econometrica : Journal of the Econometric Society*, pages 680–684, 1952.
- [21] Moulin, H.: *The strategy of social choice*. North-Holland, 1983.
- [22] Novaro, Arianna, Umberto Grandi, Dominique Longin et Emiliano Lorini: *Goal-based Voting with Propositional Goals*. Dans *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, 2018.
- [23] Papadimitriou, Christos H.: *Computational Complexity*. John Wiley and Sons, 2003.
- [24] Satterthwaite, Mark Allen: *Strategy-proofness and Arrow's conditions : Existence and correspondence theorems for voting procedures and social welfare functions*. *Journal of Economic Theory*, 10(2) :187–217, 1975.
- [25] Wooldridge, Michael: *An introduction to multi-agent systems*. John Wiley and Sons, 2009.
- [26] Wooldridge, Michael, Ullé Endriss, Sarit Kraus et Jérôme Lang: *Incentive Engineering in Boolean Games*. *Artificial Intelligence*, 195 :418–439, 2013.
- [27] Young, H.P: *An axiomatization of Borda's rule*. *Journal of Economic Theory*, 9(1) :43 – 52, 1974.