Egalitarian Judgment Aggregation

Sirin Botan
University of Amsterdam, The Netherlands
sirin.botan@uva.nl

Ronald de Haan
University of Amsterdam, The Netherlands
me@ronalddehaan.eu

Marija Slavkovik
University of Bergen, Norway
marija.slavkovik@uib.no

Zoi Terzopoulou
University of Amsterdam, The Netherlands
z.terzopoulou@uva.nl

ABSTRACT

Egalitarian considerations play a central role in many areas of social choice theory. Applications of egalitarian principles range from ensuring everyone gets an equal share of a cake when deciding how to divide it, to guaranteeing balance with respect to gender or ethnicity in committee elections. Yet, the egalitarian approach has received little attention in judgment aggregation—a powerful framework for aggregating logically interconnected issues. We make the first steps towards filling that gap. We introduce axioms capturing two classical interpretations of egalitarianism in judgment aggregation and situate these within the context of existing axioms in the pertinent framework of belief merging. We then explore the relationship between these axioms and several notions of strategyproofness from social choice theory at large. Finally, a novel egalitarian judgment aggregation rule stems from our analysis; we present complexity results concerning both outcome determination and strategic manipulation for that rule.

KEYWORDS

Social Choice Theory, Judgment Aggregation, Egalitarianism, Strategic Manipulation, Computational Complexity

ACM Reference Format:
Sirin Botan, Ronald de Haan, Marija Slavkovik, and Zoi Terzopoulou. 2021. Egalitarian Judgment Aggregation. In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), Online, May 3–7, 2021*, IFAAMAS, 14 pages.

1 INTRODUCTION

Judgment aggregation is an area of social choice theory concerned with turning the individual binary judgments of a group of agents over logically related issues into a collective judgment [23]. Being a flexible and widely applicable framework, judgment aggregation provides the foundations for collective decision making settings in various disciplines, like philosophy, economics, legal theory, and artificial intelligence [39]. The purpose of judgment aggregation methods (*rules*) is to find those collective judgments that better represent the group as a whole. Following the utilitarian approach in social choice, an "ideal" such collective judgment has traditionally been considered the will of the majority. In this paper we challenge this perspective, introducing a more egalitarian point of view.

In economic theory, utilitarian approaches are often contrasted with egalitarian ones [55]. In the context of judgment aggregation, an egalitarian rule must take into account whether the collective outcome achieves equally distributed satisfaction among agents and ensure that agents enjoy equal consideration. A rapidly growing application domain of egalitarian judgment aggregation (that also concerns multiagent systems with practical implications like in the construction of self-driving cars) is the aggregation of moral choices [15], where utilitarian approaches do not always offer appropriate solutions [4, 57]. One of the drawbacks of majoritarianism is that a strong enough majority can cancel out the views of a minority, which is questionable in several occasions.

For example, suppose that the president of a student union has secured some budget for the decoration of the union's office and she asks her colleagues for their opinions on which paintings to buy (perhaps imposing some constraints on the combinations of paintings that can be simultaneously selected, due to clashes on style). If the members of the union largely consist of pop-art enthusiasts that the president tries to satisfy, then a few members with diverting taste will find themselves in an office that they detest; an arguably more viable strategy would be to ensure that—as much as possible—no-one is strongly dissatisfied. But then, consider a similar situation in which a kindergarten teacher needs to decide what toys to complement the existing playground with. In that case, the teacher's goal is to select toys that equally (dis)satisfy all kids involved, so that no extra tension is created due to envy, which the teacher will have to resolve—if the kids disagree a lot, then the teacher may end up choosing toys that none of them really likes.

In order to formally capture scenarios like the above, this paper introduces two fundamental properties (also known as *axioms*) of egalitarianism to judgment aggregation, inspired by the theory of justice. The first captures the idea behind the so-called *veil of ignorance* of Rawls [60], while the second speaks about how happy agents are with the collective outcome relative to each other.

Our axioms closely mirror properties in other areas of social choice theory. In *belief merging*, egalitarian axioms and merging operators have been studied by Everaere et al. [29]. The nature of their axioms is in line with the interpretation of egalitarianism in this paper, although the two main properties they study are logically weaker than ours, as we further discuss in Section 3.1. In *resource allocation*, fairness has been interpreted both as maximising the share of the worst off agent [11] as well as eliminating envy between agents [32]. In *multiwinner elections*, egalitarianism is present in diversity [22] and in proportional representation [2, 20] notions.

Unfortunately, egalitarian considerations often come at a cost. A central concern in many areas of social choice theory, of which judgement aggregation does not constitute an exception, is that agents may have incentives to *manipulate*, i.e., to misrepresent their

Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), U. Endriss, A. Nowé, F. Dignum, A. Lomuscio (eds.), May 3–7, 2021, Online.

judgments aiming for a more preferred outcome [18]. Frequently, it is impossible to simultaneously be fair and avoid strategic manipulation. For both variants of fairness in resource allocation, rules satisfying them usually are susceptible to strategic manipulation [1, 9, 13, 54]. The same type of results have recently been obtained for multiwinner elections [49, 58]. It is not easy to be egalitarian while disincentivising agents from taking advantage of it.

Inspired by notions of manipulation stemming from voting theory, we explore how our egalitarian axioms affect the agents' strategic behaviour within judgment aggregation. Our most important result in this vein is showing that the two properties of egalitarianism defined in this paper clearly differ in terms of strategyproofness.

Our axioms give rise to two concrete egalitarian rules—one that has been previously studied, and one that is new to the literature. For the latter, we are interested in exploring how computationally complex its use is in the worst-case scenario. This kind of question, first addressed by Endriss et al. [28], is regularly asked in the literature of judgment aggregation [5, 25, 51]. As Endriss et al. [26] wrote recently, the problem of determining the collective outcome of a given judgment aggregation rule is "the most fundamental algorithmic challenge in this context".

The remainder of this paper is organised as follows. Section 2 reviews the basic model of judgment aggregation, while Section 3 introduces our two original axioms of egalitarianism and the rules they induce. Section 4 analyses the relationship between egalitarianism and strategic manipulation in judgment aggregation, and Section 5 focuses on relevant computational aspects: although the general problems of outcome determination and of strategic manipulation are proven to be very difficult, we propose a way to confront them with the tools of *Answer Set Programming* [36].

2 BASIC MODEL

Our framework relies on the standard formula-based model of judgment aggregation [52], but for simplicity we also use notation commonly employed in binary aggregation [38].

Let $\mathbb{N}$ denote the (countably infinite) set of all agents that can potentially participate in a judgment aggregation setting. In every specific such setting, a finite set of agents $N \subset \mathbb{N}$ of size $n \geq 2$ express judgments on a finite and nonempty set of *issues* (formulas in propositional logic) $\Phi = \{\varphi_1, \ldots, \varphi_m\}$, called the *agenda*. $\mathcal{J}(\Phi) \subseteq \{0,1\}^m$ denotes the set of all admissible opinions on Φ. Then, a *judgment* J is a vector in $\mathcal{J}(\Phi)$, with 1 (0) in position k meaning that the issue φ_k is accepted (rejected). $\overline{J}$ is the *antipodal* judgment of J: for all $\varphi \in \Phi$, φ is accepted in $\overline{J}$ if and only if it is rejected in J.

A *profile* $\boldsymbol{J} = (J_1, \ldots J_n) \in \mathcal{J}(\Phi)^n$ is a vector of individual judgments, one for each agent in a group N. We write $\boldsymbol{J}' =_{-i} \boldsymbol{J}$ when the profiles $\boldsymbol{J}$ and $\boldsymbol{J}'$ are the same, besides the judgment of agent i. We write $\boldsymbol{J}_{-i}$ to denote the profile $\boldsymbol{J}$ with agent i's judgment removed, and $(\boldsymbol{J}, J) \in \mathcal{J}(\Phi)^{n+1}$ to denote the profile $\boldsymbol{J}$ with judgment J added. A *judgment aggregation rule* F is a function that maps every possible profile $\boldsymbol{J} \in \mathcal{J}(\Phi)^n$, for every group N and agenda Φ, to a nonempty set $F(\boldsymbol{J})$ of collective judgments in $\mathcal{J}(\Phi)$. Note that a judgment aggregation rule is defined over groups and agendas of variable size, and may return several, tied, collective judgments.

The agents that participate in a judgment aggregation scenario will naturally have preferences over the outcome produced by the aggregation rule. First, given an agent i's truthful judgment J_i, we need to determine when agent i would prefer a judgment J over a different judgment J'. The most prevalent type of such preferences considered in the judgment aggregation literature is that of *Hamming distance* preferences [6–8, 64].

The Hamming distance between two judgments J and J' equals the number of issues on which these judgments disagree—concretely, it is defined as $H(J, J') = \sum_{\varphi \in \Phi} |J(\varphi) - J'(\varphi)|$, where $J(\varphi)$ denotes the binary value in the position of φ in J. For example, $H(100, 111) = 2$. Then, the (weak, and analogously strict) preference of agent i over judgments is defined by the relation $\succeq_i$ (where $J \succeq_i J'$ means that i's utility from J is higher than that from J'):

$$J \succeq_i J' \text{ if and only if } H(J_i, J) \leq H(J_i, J').$$

But an aggregation rule often outputs more than one judgment, and thus we also need to determine agents' preferences over sets of judgments.[1] We define two requirements guaranteeing that the preferences of the agents over sets of judgments are consistent with their preferences over single judgments. To that end, let $\mathring{\succeq}_i$ (with strict part $\mathring{\succ}_i$) denote agent i's preferences over sets $X, Y \subseteq \mathcal{J}(\Phi)$. We require that $\mathring{\succeq}_i$ is related to $\succeq_i$ as follows:

- $J \succeq_i J'$ if and only if $\{J\} \mathring{\succeq}_i \{J'\}$, for any $J, J' \in \mathcal{J}(\Phi)$;
- $X \mathring{\succ}_i Y$ implies that there exist some $J \in X$ and $J' \in Y$ such that $J \succ_i J'$ and $\{J, J'\} \not\subseteq X \cap Y$.

The above conditions hold for almost all well-known preference extensions. For example, they hold for the *pessimistic* preference ($X \succ^{pess} Y$ if and only if there exists $J' \in Y$ such that $J \succ J'$ for all $J \in X$) and the *optimistic* preference ($X \succ^{opt} Y$ if and only if there exists $J \in X$ such that $J \succ J'$ for all $J' \in Y$) of Duggan and Schwartz [19], as well as the preference extensions of Gärdenfors [33] and Kelly [43]. The results provided in this paper abstract away from specific preference extensions.

3 EGALITARIAN AXIOMS AND RULES

This section focuses on two axioms of egalitarianism in judgment aggregation. We examine them in relation to each other and to existing properties from belief merging, as well as to the standard majority property defined below. Most of the well-known judgment aggregation rules return the majority opinion, when that opinion is logically consistent [24].[2]

Let $m(\boldsymbol{J})$ be the judgment that accepts exactly those issues accepted by a strict majority of agents in $\boldsymbol{J}$. A rule F is *majoritarian* when for all profiles $\boldsymbol{J}$, $m(\boldsymbol{J}) \in \mathcal{J}(\Phi)$ implies that $F(\boldsymbol{J}) = \{J\}$.

Our first axiom with an egalitarian flavour is the *maximin property*, suggesting that we should aim at maximising the utility of those agents that will be worst off in the outcome. Assuming that everyone submits their truthful judgment during the aggregation process, this means that we should try to minimise the distance of the agents that are furthest away from the outcome. Formally:

- ► A rule F satisfies the **maximin** property if for all profiles $\boldsymbol{J} \in \mathcal{J}(\Phi)^n$ and judgments $J \in F(\boldsymbol{J})$ there do not exist judgment

[1] Various approaches have been taken within the area of social choice theory in order to extend preferences over objects to preferences over sets of objects —see Barberà et al. [3] for a review.

[2] A central problem in judgment aggregation concerns the fact that the issue-wise majority is not always logically consistent [52].

$J' \in \mathcal{J}(\Phi)$ and agent $j \in N$ such that

$$H(J_i, J') < H(J_j, J) \text{ for all } i \in N.$$

Although the maximin property is quite convincing, there are settings like those motivated in the Introduction where it does not offer sufficient egalitarian guarantees. We thus consider a different property next, which we call the *equity property*. This axiom requires that the gaps in the agents' satisfaction be minimised. In other words, no two agents should find themselves in very different distances with respect to the collective outcome. Formally:

- A rule F satisfies the **equity** property if for all profiles $\boldsymbol{J} \in \mathcal{J}^n$ and judgments $J \in F(\boldsymbol{J})$, there do not exist judgment $J' \in \mathcal{J}(\Phi)$ and agents $i', j' \in N$ such that

$$|H(J_i, J') - H(J_j, J')| < |H(J_{i'}, J) - H(J_{j'}, J)| \text{ for all } i, j \in N.$$

No rule that satisfies either the maximin- or equity property can be majoritarian.[3] As an illustration, in a profile of only two agents who disagree on some issues, any egalitarian rule will try to reach a compromise, and this compromise will not be affected if any agents holding one of the two initial judgments are added to the profile—in contrast, a majoritarian rule will simply conform to the crowd.

Proposition 1 shows that it is also impossible for the maximin property and the equity property to simultaneously hold. Therefore, we have established the logical independence of all three axioms discussed so far: maximin, equity, and majoritarianism.

Proposition 1. *No judgment aggregation rule can satisfy both the maximin property and the equity property.*

Proof. Take an agenda Φ where $\mathcal{J}(\Phi)$ consists of the nodes in the graph below and consider the profile $\boldsymbol{J} = (J_1, J_2)$. Each edge is labelled with the Hamming distance between the judgments.

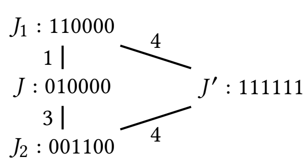

Every aggregation rule satisfying the maximin property will return $\{J\}$, as this judgment maximises the utility of the worst off agent—in this case, agent 2. However every rule satisfying the equity property will return $\{J'\}$, as this judgment minimises the difference in utility between the best off and worst off agents. Thus, there is no rule that can satisfy the two properties at the same time. ◻

From Proposition 1, we also know now that the two properties of egalitarianism generate two disjoint classes of aggregation rules. In particular, in this paper we focus on the *maximal* rule that meets each property: a rule F is the maximal one of a given class if, for every profile $\boldsymbol{J}$, the outcomes obtained by any other rule in that class are always outcomes of F too.[4]

The maximal rule satisfying the maximin property is the rule *MaxHam* (see, e.g., Lang et al., 2011). For all profiles $\boldsymbol{J} \in \mathcal{J}(\Phi)^n$,

$$\text{MaxHam}(\boldsymbol{J}) = \underset{J \in \mathcal{J}(\Phi)}{\operatorname{argmin}} \max_{i \in N} H(J_i, J).$$

Analogously, we define a rule new to the judgment aggregation literature, which is the maximal one satisfying the equity property. For all profiles $\boldsymbol{J} \in \mathcal{J}(\Phi)^n$,

$$\text{MaxEq}(\boldsymbol{J}) = \underset{J \in \mathcal{J}(\Phi)}{\operatorname{argmin}} \max_{i,j \in N} |H(J_i, J) - H(J_j, J)|.$$

To better understand these rules, consider an agenda with six issues: $p, q, r \equiv p \wedge q$, and their negations. Suppose that there are only two agents in a profile $\boldsymbol{J}$, holding judgments $J_1 = (111)$ and $J_2 = (010)$. Then, we have that $\text{MaxHam}(\boldsymbol{J}) = \{(111), (010)\}$, while $\text{MaxEq} = \{(000), (100)\}$. In this example, the difference in spirit between the two rules of our interest is evident. Although the MaxHam rule is able to fully satisfy exactly one of the agents without causing much harm to the other, it still creates greater unbalance than the MaxEq rule, which ensures that the two agents are equally happy with the outcome (under Hamming-distance preferences). In that sense, MaxEq is better suited for a group of agents that do not want any of them to feel particularly put upon, while MaxHam seems more desirable when a minimum level of happiness is asked for.

MaxHam generalises minimax approval voting [10], which is the special case without logical constraint on the judgments, meaning agents may approve any subset of issues. Brams et al. [10] show that MaxHam remains manipulable in this special case. As finding the outcome of minimax is computationally hard, Caragiannis et al. [12] provide approximation algorithms that circumvent this problem. They also demonstrate the interplay between manipulability and lower bounds for the approximation algorithm—establishing strategyproofness results for approximations of minimax.

3.1 Relations with Egalitarian Belief Merging

A framework closely related to ours is that of belief merging [45], which is concerned with how to aggregate several (possibly inconsistent) sets of beliefs into one consistent belief set.[5] Egalitarian belief merging is studied by Everaere et al. [29], who examine interpretations of the *Sen-Hammond equity condition* [63] and the *Pigou-Dalton transfer principle* [17]—two properties that are logically incomparable.[6] We situate our egalitarian axioms within the context of these egalitarian axioms from belief merging; we reformulate these axioms into our framework.

- Fix an arbitrary profile $\boldsymbol{J}$, agents i, j, and any three judgment sets $J, J' \in \mathcal{J}(\Phi)$. An aggregation rule F satisfies the **Sen-Hammond equity property** if whenever

$$H(J_i, J) < H(J_i, J') < H(J_j, J') < H(J_j, J)$$

and $H(J_{i'}, J) = H(J_{i'}, J')$ for all other agents $i' \in N \setminus \{i, j\}$, then $J \in F(\boldsymbol{J})$ implies $J' \in F(\boldsymbol{J})$.

Proposition 2. *If a rule satisfies either the maximin property or the equity property, then it will satisfy the Sen-Hammond equity property.*

[3] This includes popular rules like the median rule [56]—known under a number of other names, notably *distance-based rule* [59], *Kemeny rule* [24], and *prototype rule* [53].

[4] Of course, several natural refinements of these rules can be defrined, with respect to various other axiomatic properties that we may find desirable. Identifying and studying such rules is an interesting direction for future research.

[5] We refer to Everaere et al. [30] for a detailed comparison of the two frameworks.

[6] Another egalitarian property in belief merging is the *arbitration postulate*. We do not go into detail on this postulate, but refer the reader to Konieczny and Pérez [45].

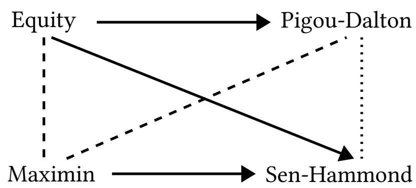

Figure 1: Dashed lines denote incompatibility, dotted lines incomparability, and arrows implication relations.

.

Proof (sketch). Let $\boldsymbol{J} = (J_i, J_j)$ be a profile such that $H(J, J_i) < H(J', J_i) < H(J', J_j) < H(J, J_j)$, and $H(J_{i'}, J) = H(J_{i'}, J')$ for all other agents $i' \in N \setminus \{i, j\}$. Suppose F satisfies the equity property—if there is some agent i' such that $|H(J_i, J) - H(J_{i'}, J)| > |H(J_i, J) - H(J_j, J)|$, then $J \in F(\boldsymbol{J})$ if and only if $J' \in F(\boldsymbol{J})$, as the maximal difference in distance will be the same for the two judgments. If this is not the case, then agents i and j determine the outcome regarding J and J' so clearly $J \in F(\boldsymbol{J})$ implies $J' \in F(\boldsymbol{J})$. The argument for other cases proceeds similarly.

If F satisfies the maximin property, then a similar argument tells us that if membership of J and J' in the outcome is determined by an agent other than i or j, we will either have both or neither. If i, and j are the determining factor then $J \in F(\boldsymbol{J})$ implies $J' \in F(\boldsymbol{J})$. □

- Given a profile $\boldsymbol{J} = (J_1, \ldots, J_n)$ and agents i and j such that:
  - $H(J_i, J) < H(J_i, J') \leq H(J_j, J') < H(J_j, J)$,
  - $H(J_i, J') - H(J_i, J) = H(J_j, J') - H(J_j, J)$, and
  - $H(J_{i^*}, J) = H(J_{i^*}, J')$ for all other agents $i^* \in N \setminus \{i, j\}$,

  F satisfies the **Pigou-Dalton transfer principle** if $J' \in F(\boldsymbol{J})$ implies $J \notin F(\boldsymbol{J})$.

We refer to these axioms simply as *Sen-Hammond*, and *Pigou-Dalton*. Note that Pigou-Dalton is also a weaker version of our equity property, as it stipulates that the difference between utility in agents should be lessened under certain conditions, while the equity property always aims to minimise this distance.

While we can find a rule that satisfies both the equity property and a weakening of the maximin property, Sen-Hammond, we cannot do the same by weakening the equity property.

Proposition 3. *No judgment aggregation rule can satisfy both the maximin property and Pigou-Dalton.*

Proof. Consider the domain $\mathcal{J}(\Phi) = \{J_1, J_2, J_3, J, J'\}$ with the following Hamming distances between judgment sets.[7]

	J	J'	J_1	J_2	J_3
J_1	2	4	0	4	8
J_2	6	4	4	0	10
J_3	6	6	8	10	0

Let $\boldsymbol{J} = (J_1, J_2, J_3)$. If F satisfies the maximin property, $\{J, J'\} \subseteq F(\boldsymbol{J})$, as we can see from the grey cells. This means Pigou-Dalton is violated in this profile, as $J' \in F(\boldsymbol{J})$ should imply $J \notin F(\boldsymbol{J})$. □

We summarise the observations of this section in Figure 1.

[7] One such domain would be the following, where J = 00000000111, J' = 00000001110, J_1 = 00000010011, J_2 = 00000111000, and J_3 = 11111001111.

4 STRATEGIC MANIPULATION

This section provides an account of strategic manipulation with respect to the egalitarian axioms defined in Section 3. We start off with presenting the most general notion of strategic manipulation in judgment aggregation, introduced by Dietrich and List [18].[8] We assume Hamming preferences throughout this section.

Definition 1. *A rule F is susceptible to **manipulation** by agent i in profile $\boldsymbol{J}$, if there exists a profile $\boldsymbol{J}' =_{-i} \boldsymbol{J}$ such that $F(\boldsymbol{J}') \mathrel{\mathring{>}_i} F(\boldsymbol{J})$.*

We say that F is *strategyproof* in case F is not manipulable by any agent $i \in N$ in any profile $\boldsymbol{J} \in \mathcal{J}(\Phi)^n$.

Proposition 4 shows an important fact: In judgment aggregation, egalitarianism is incompatible with strategyproofness.[9]

Proposition 4. *If an aggregation rule is strategyproof, it cannot satisfy the maximin property or the equity property.*

Proof. We show the contrapositive. Let Φ be an agenda such that $\mathcal{J}(\Phi) = \{000000, 110000, 111000, 111111\}$. Consider the following two profiles $\boldsymbol{J}$ (left) and $\boldsymbol{J}'$ (right).

J_i	111000
J_j	000000
$F(\boldsymbol{J})$	110000

J'_i	111111
J'_j	000000
$F(\boldsymbol{J}')$	111000

In profile $\boldsymbol{J}$, both the maximin and the equity properties prescribe that 110000 should be returned as the single outcome, while in profile $\boldsymbol{J}'$ they agree on 111000. Because $\boldsymbol{J}' = (\boldsymbol{J}_{-i}, J'_i)$, and 111000 $>_i$ 110000, this is a successful manipulation. Thus, if F satisfies the maximin or the equity property, it fails strategyproofness. □

Strategyproofness according to Definition 1 is a strong requirement, which many known rules fail [8]. We investigate two more nuanced notions of strategyproofness that are novel to judgment aggregation, yet have familiar counterparts in voting theory.

First, *no-show manipulation* happens when an agent can achieve a preferable outcome simply by not submitting any judgment, instead of reporting a truthful or an untruthful one.

Definition 2. *A rule F is susceptible to **no-show manipulation** by agent i in profile $\boldsymbol{J}$ if $F(\boldsymbol{J}_{-i}) \mathrel{\mathring{>}_i} F(\boldsymbol{J})$.*

We say that F satisfies *participation* if it is not susceptible to no-show manipulation by any agent $i \in N$ in any profile.[10]

Second, *antipodal strategyproofness* poses another barrier against manipulation, by stipulating that an agent cannot change the outcome towards a better one for herself by reporting a totally untruthful judgment. This is a strictly weaker requirement than full strategyproofness, serving as a protection against excessive lying.

Definition 3. *A rule F is susceptible to **antipodal manipulation** by agent i in profile $\boldsymbol{J}$ if $F(\boldsymbol{J}_{-i}, \overline{J_i}) \mathrel{\mathring{>}_i} F(\boldsymbol{J})$.*

We say that F satisfies *antipodal strategyproofness* if it not susceptible to antipodal manipulation by any agent $i \in N$ in any profile. As is the case for participation, antipodal strategyproofness is a weaker notion of strategyproofness as far as the MaxHam and the MaxEq rules are concerned.

[8] The original definition of Dietrich and List [18] concerned single-judgment collective outcomes, and a type of preferences that covers Hamming-distance ones.

[9] This in in line with Brams et al.'s work on the minimax rule in approval voting.

[10] cf. the no-show paradox in voting [31].

In voting theory, Sanver and Zwicker [61] show that participation implies antipodal strategyproofness (or *half-way monotonicity*, as called in that framework) for rules that output a single winning alternative. Notably, this is not always the case in our model (see Example 1). This is not surprising, as obtaining such a result independently of the preference extension would be significantly stronger than the result by Sanver and Zwicker [61]. We are, however, able to reproduce this relationship between participation and strategyproofness in Theorem 1, for a specific type of preferences.

Example 1. *We present a rule that satisfies participation but violates antipodal strategyproofness. The other direction admits a similar example, and is thus omitted. Note that the rule demonstrated is quite unnatural for simplicity of the presentation.*

Consider an agenda Φ with $\mathcal{J}(\Phi) = \{00, 01, 11\}$.[11] *We construct an anonymous rule F that is only sensitive to* which *judgments are submitted and not to their quantity:*

$F(00) = F(11) = F(01, 00) = F(00, 11) = \{01, 11\}$;
$F(01) = \{00, 11\}$; $F(01, 11) = F(01, 00, 11) = \{01\}$.

For the pessimistic preference, no agent can be strictly better off by abstaining. However, compare the profiles $(01, 00)$ and $(01, 11)$: agent 2 with truthful judgment 00 can move from outcome $\{01, 11\}$ to outcome $\{01\}$, which is strictly better for her.

While the two axioms are independent in the general case, participation implies antipodal strategyproofness (Theorem 1) if we stipulate that

- $X \mathrel{\mathring{\succ}_i} Y$ **if and only if** there exist some $J \in X$ and $J' \in Y$ such that $J \succ_i J'$ and $\{J, J'\} \not\subseteq X \cap Y$.

If a preference satisfies the above condition, we say that it is *decisive*. This condition gives rise to a preference extension equivalent to the *large preference extension* of Kruger and Terzopoulou [48]. Note that a decisive preference is not necessarily acyclic—in fact, it may even be symmetric. The interpretation of such a preference extension is slightly different than the usual one; when we say that a rule is strategyproof for a decisive preference where both $J \mathrel{\mathring{\succ}} J'$ and $J' \mathrel{\mathring{\succ}} J$ hold, we mean that no agent i with $J \mathrel{\mathring{\succ}_i} J'$ and no agent $j \neq i$ with $J' \mathrel{\mathring{\succ}_j} J$ will ever have an incentive to manipulate.

Using Lemma 1, we can now prove a result analogous to the one in voting theory, to give a complete picture of how these axioms relate to each other in judgment aggregation.

Lemma 1. *For judgment sets J, J' and J'': $H(J, J') > H(J, J'')$, if and only if $H(\overline{J}, J') < H(\overline{J}, J'')$.*

Proof. For judgment sets $J, J' \in \mathcal{J}(\Phi)$, $H(\overline{J}, J') = m - H(J, J')$. Suppose $H(J, J') > H(J, J'')$. Then $H(\overline{J}, J') = m - H(J, J') < m - H(J, J'') = H(\overline{J}, J'')$. The other direction is analogous. □

Theorem 1. *For decisive preferences over sets of judgments, participation implies antipodal strategyproofness.*

Proof. Working on the contrapositive, suppose that F is susceptible to antipodal manipulation. We will prove that F is susceptible to no-show manipulation too. We know that there exists $i \in N$ such that $F(\boldsymbol{J}_{-i}, \overline{J_i}) \mathrel{\mathring{\succ}_i} F(\boldsymbol{J}_{-i}, J_i)$, for some profile $\boldsymbol{J}$. This means that there exist $J' \in F(\boldsymbol{J}_{-i}, \overline{J_i})$ and $J \in F(\boldsymbol{J}_{-i}, J_i)$ with $J' \succ_i J$. Equivalently,

$$H(J_i, J') < H(J_i, J) \tag{1}$$

Next, consider a judgment $J'' \in F(\boldsymbol{J}_{-i})$.

If $H(\overline{J_i}, J'') < H(\overline{J_i}, J')$, then F is susceptible to no-show manipulation by agent i in the profile $(\boldsymbol{J}_{-i}, \overline{J_i})$.

Otherwise, $H(\overline{J_i}, J') \leq H(\overline{J_i}, J'')$. Then Lemma 1 implies that $H(J_i, J'') \leq H(J_i, J')$. So, together with Inequality (1), we have that $H(J_i, J'') < H(J_i, J)$. This means that F is susceptible to no-show manipulation by agent i in the profile $(\boldsymbol{J}_{-i}, J_i)$. □

We next prove that any rule satisfying the maximin property is immune to both no-show manipulation and antipodal manipulation (Theorem 2), while this is not true for the equity property (Proposition 5).[12] We emphasise that the theorem holds for *all* preference extensions. These results—holding for two independent notions of strategyproofness—are significant for two reasons. First, they bring to light the conditions under which we can have our cake and eat it too, simultaneously satisfying an egalitarian property and a degree of strategyproofness. In addition, they provide a further way to distinguish between the properties of maximin and equity: the former is better suited in contexts where we may worry about the agents' strategic behaviour.

Theorem 2. *The maximin property implies participation and antipodal strategyproofness.*

Proof. We prove the participation case; the proof for antipodal strategyproofness is analogous, and utilises Lemma 1.

Suppose for contradiction that F is a rule that satisfies the maximin property but violates participation. Then there must exist agent $i \in N$ and profile $\boldsymbol{J}$ where J_i is agent i's truthful judgment, such that $F(\boldsymbol{J}_{-i}) \mathrel{\mathring{\succ}_i} F(\boldsymbol{J})$. This means there must exist judgments $J \in F(\boldsymbol{J})$ and $J' \in F(\boldsymbol{J}_{-i})$ such that $J' \succ_i J$ and $\{J, J'\} \not\subseteq F(\boldsymbol{J}) \cap F(\boldsymbol{J}_{-i})$. Because agent i strictly prefers J' to J, this means that $H(J_i, J) > H(J_i, J')$. We consider two cases.

Case 1: Suppose that $J' \notin F(\boldsymbol{J})$. Let k be the distance between the worst off agent's judgment in $\boldsymbol{J}$ and any judgment in $F(\boldsymbol{J})$. Then,

$$H(J_{j'}, J) \leq k \text{ for all } j' \in N. \tag{2}$$

We know that $H(J_i, J') < k$ because $H(J_i, J) \leq k$, and agent i strictly prefers J' to J. From Inequality (2), this means that if J' is not among the outcomes in $F(\boldsymbol{J})$, there has to be some $j \in N \setminus \{i\}$ such that $H(J_j, J') > k$. But all judgments submitted to profile $(\boldsymbol{J}_{-i})$ by agents in $N \setminus \{i\}$ are at most at distance k from J by Inequality (2), so J would be selected by any rule satisfying the maximin property will select J as an outcome of $F(\boldsymbol{J}_{-i})$—instead of J', a contradiction.

Case 2: Suppose that $J' \in F(\boldsymbol{J})$, meaning that $J \notin F(\boldsymbol{J}_{-i})$. Analogously to the first case, let k' be the distance between the worst off agent's judgment in $\boldsymbol{J}_{-i}$ and any judgment in $F(\boldsymbol{J}_{-i})$. Then,

$$H(J_{j'}, J') \leq k' \text{ for all } j' \in N \setminus \{i\}. \tag{3}$$

Moreover, since $J \notin F(\boldsymbol{J}_{-i})$, it is the case that

$$H(J_j, J) > k' \text{ for some } j \neq i. \tag{4}$$

[11] For other agendas we can simply take the rule to be constant.

[12] Note that antipodal strategyproofness is not so weak a requirement that is immediately satisfied by all "utilitarian" aggregation rules. For example, the Copeland voting rule fails the analogous axiom of half-way monotonicity [66].

In profile $\boldsymbol{J}$, Inequalities (3) and (4) still hold. In addition, we have that $H(J_i, J) > H(J_i, J')$ because agent i strictly prefers J' to J. So, for any rule satisfying the maximin property, judgment J' will be better as an outcome of $F(\boldsymbol{J})$ than J, a contradiction. □

Corollary 1. *The rule MaxHam satisfies antipodal strategyproofness and participation.*

Proposition 5. *No rule that satisfies the equity property can satisfy participation or antipodal strategyproofness .*

Proof. The following is a counterexample for antipodal strategyproofness. A similar one exists for participation.

Consider the following profiles $\boldsymbol{J} = \{Ji, J_j\}$ and $\boldsymbol{J}' = (\boldsymbol{J}_{-i}, \overline{J_i})$. We give a visual representation of the profiles as well as the outcomes under an arbitrary rule F that satisfies the equity principle. We specify that $\mathcal{J}(\Phi) = \{00110, 00000, 01110, 10000, 11111\}$.

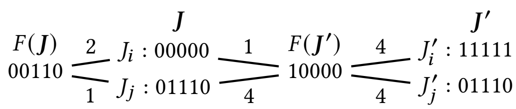

Each edge from an individual judgment to a collective one is labelled with the Hamming distance between the two. It is clear that agent i will benefit from her antipodal manipulation, as her true judgment is much closer to the singleton outcome in $\boldsymbol{J}'$ than the singleton outcome in $\boldsymbol{J}$. □

Corollary 2. *The rule MaxEq does not satisfy participation or antipodal strategyproofness.*

5 COMPUTATIONAL ASPECTS

We have discussed two aggregation rules that reflect desirable egalitarian principles—i.e., the MaxHam and MaxEq rules—and examined whether they give agents incentives to misrepresent their truthful judgments. In this section we consider how complex it is, computationally, to employ these rules, and the complexity of determining whether an agent can manipulate the collective outcome.

The MaxHam rule has been considered from a computational perspective before [40–42]. Here, we extend this analysis to the MaxEq rule, and we compare the two rules with each other on their computational properties. Concretely, we primarily establish some computational complexity results; motivated by these results, we then illustrate how some computational problems related to these rules can be solved using the paradigm of Answer Set Programming.

5.1 Computational Complexity

We investigate some computational complexity aspects of the judgment aggregation rules that we have considered. Due to space constraints, we will only describe the main lines of these results—for full details, we refer to the accompanying Appendix.

Consider the problem of outcome determination (for a rule F). This is most naturally modelled as a search problem, where the input consists of an agenda Φ and a profile $\boldsymbol{J} = (J_1, \ldots, J_n) \in \mathcal{J}(\Phi)^n$. The problem is to produce some judgment set $J^* \in F(\boldsymbol{J})$. We will show that for the MaxEq rule, this problem can be solved in polynomial time with a logarithmic number of calls to an oracle for NP search problems (where the oracle also produces a witness for yes answers—also called an FNP witness oracle). Said differently, the outcome determination problem for the the MaxEq rule lies in the complexity class $\mathsf{FP}^{\mathsf{NP}}[\text{log,wit}]$. We also show that the problem is complete for this class (using the standard type of reductions used for search problems: polynomial-time Levin reductions).

Theorem 3. *The outcome determination problem for the MaxEq rule is* $\mathsf{FP}^{\mathsf{NP}}[\text{log,wit}]$*-complete under polynomial-time Levin reductions.*

Proof (sketch). Membership in $\mathsf{FP}^{\mathsf{NP}}[\text{log,wit}]$ can be shown by giving a polynomial-time algorithm that solves the problem by querying an FNP witness oracle a logarithmic number of times. The algorithm first finds the minimum value k of $\max_{J',J'' \in \boldsymbol{J}} |H(J, J') - H(J, J'')|$ by means of binary search—requiring a logarithmic number of oracle queries. Then, with one additional oracle query, the algorithm can produce some $J^* \in \mathcal{J}(\Phi)$ with $\max_{J',J'' \in \boldsymbol{J}} |H(J^*, J') - H(J^*, J'')| = k$.

To show $\mathsf{FP}^{\mathsf{NP}}[\text{log,wit}]$-hardness, we reduce from the problem of finding a satisfying assignment of a (satisfiable) propositional formula ψ that sets a maximum number of variables to true [14, 44]. This reduction works roughly as follows. Firstly, we produce 3CNF formulas $\psi_1, \ldots, \psi_v$ where each ψ_i is 1-in-3-satisfiable if and only if there exists a satisfying assignment of ψ that sets at least i variables to true. Then, for each i, we transform ψ_i to an agenda Φ_i and a profile $\boldsymbol{J}_i$ such that there is a judgment set with equal Hamming distance to each $J \in \boldsymbol{J}_i$ if and only if ψ_i is 1-in-3-satisfiable. Finally, we put the agendas Φ_i and profiles $\boldsymbol{J}_i$ together into a single agenda Φ and a single profile $\boldsymbol{J}$ such that we can—from the outcomes selected by the MaxEq rule—read off the largest i for which ψ_i is 1-in-3-satisfiable, and thus, the maximum number of variables set to true in any truth assignment satisfying ψ. This last step involves duplicating issues in $\Phi_1, \ldots, \Phi_v$ different numbers of times, and creating logical dependencies between them. Moreover, we do this in such a way that from any outcome selected by the MaxEq rule, we can reconstruct a truth assignment satisfying ψ that sets a maximum number of variables to true. □

The result of Theorem 3 means that the computational complexity of computing outcomes for the MaxEq rule lies at the Θ_2^{p}-level of the Polynomial Hierarchy. This is in line with previous results on the computational complexity of the outcome determination problem for the MaxHam rule—De Haan and Slavkovik [41] showed that a decision variant of the outcome determination problem for the MaxHam rule is Θ_2^{p}-complete. Notably, our proof (presented in detail in the Appendix) brings out an intriguing fact about a problem that is at first glance simpler than outcome determination for MaxEq: Given an agenda Φ and a profile $\boldsymbol{J}$, deciding whether the minimum value of $\max_{i,j \in N} |H(J_i, J) - H(J_j, J)|$ for $J \in \mathcal{J}(\Phi)$—the value that the MaxEq rule minimizes—is divisible by 4, is Θ_2^{p}-complete (Proposition 6). Intuitively, merely computing the minimum value that is relevant for MaxEq is Θ_2^{p}-hard.

Proposition 6. *Given an agenda Φ and a profile $\boldsymbol{J}$, deciding whether the minimal value of* $\max_{J',J'' \in \boldsymbol{J}} |H(J^*, J') - H(J^*, J'')|$ *for $J^* \in \mathcal{J}(\Phi)$, is divisible by 4, is a Θ_2^{p}-complete problem.*

Interestingly, we found that the problem of deciding if there exists a judgment set $J^* \in \mathcal{J}(\Phi)$ that has the exact same Hamming distance to each judgment set in the profile is NP-hard, even when the agenda consists of logically independent issues.

Proposition 7. *Given an agenda Φ and a profile $\boldsymbol{J}$, the problem of deciding whether there is some $J^* \in \mathcal{J}(\Phi)$ with $\max_{J',J'' \in \boldsymbol{J}} |H(J^*, J') - H(J^*, J'')| = 0$ is* NP-complete. Moreover, NP-hardness holds even for the case where Φ consists of logically independent issues—i.e., the case where $\mathcal{J}(\Phi) = \{0, 1\}^m$ for some m.

This is also in line with previous results for the MaxHam rule—De Haan [40] showed that computing outcomes for the MaxHam rule is computationally intractable even when the agenda consists of logically independent issues.

Next, we turn our attention to the problem of strategic manipulation. Specifically, we show that—for the case of decisive preferences over sets of judgment sets—the problem of deciding if an agent i can strategically manipulate is in the complexity class Σ_2^{p}.

Proposition 8. *Let $\succeq$ be a preference relation over judgment sets that is polynomial-time computable, and let $\mathring{\succeq}$ be a decisive extension over sets of judgment sets. Then the problem of deciding if a given agent i can strategically manipulate under the MaxEq rule—i.e., given Φ and $\boldsymbol{J}$, deciding if there exists some $\boldsymbol{J}' =_{-i} \boldsymbol{J}$ with MaxEq$(\boldsymbol{J}')$ $\mathring{\succ}_i$ MaxEq$(\boldsymbol{J})$—is in the complexity class Σ_2^{p}.*

Proof (sketch). To show membership in $\Sigma_2^{\mathrm{p}} = \mathrm{NP}^{\mathrm{NP}}$, we describe a nondeterministic polynomial-time algorithm with access to an NP oracle that solves the problem. The algorithm firstly guesses a new judgment set J_i' for agent i in the new profile $\boldsymbol{J}'$, and guesses a truth assignment witnessing that J_i' is consistent. Then, using the NP oracle, it computes the values $k = \max_{J',J'' \in \boldsymbol{J}} |H(J, J') - H(J, J'')|$ and $k' = \max_{J',J'' \in \boldsymbol{J}'} |H(J, J') - H(J, J'')|$, for $J \in \mathcal{J}(\Phi)$. Finally, it guesses some $J, J' \in \mathcal{J}(\Phi)$, together with truth assignments witnessing consistency, and it verifies that $J' >_i J$, that $J' \in \mathrm{MaxEq}(\boldsymbol{J}')$, that $J \in \mathrm{MaxEq}(\boldsymbol{J})$, and that $\{J, J'\} \not\subseteq \mathrm{MaxEq}(\boldsymbol{J}) \cap \mathrm{MaxEq}(\boldsymbol{J}')$. Since these final checks can all be done in polynomial time—using the previously guessed and computed information—one can verify that this can be implemented by an $\mathrm{NP}^{\mathrm{NP}}$ algorithm. □

This Σ_2^{p}-membership result can straightforwardly be extended to other variants of the manipulation problem (e.g., no-show manipulation and antipodal manipulation) and to other preferences, as well as to the MaxHam rule. Due to space constraints, we omit further details on this. Still, we shall mention that results demonstrating that strategic manipulation is very complex are generally more welcome than analogous ones regarding outcome determination. If manipulation is considered a negative side-effect of the agents' strategic behaviour, knowing that it is hard for the agents to materialise it is good news.[13] In Section 5.2 we will revisit these concerns from a different angle.

5.2 ASP Encoding for the MaxEq Rule

The complexity results in Section 5.1 leave no doubt that applying our egalitarian rules is computationally difficult. Nevertheless, they also indicate that a useful approach for computing outcomes of the MaxEq rule in practice would be to encode this problem into the paradigm of Answer Set Programming (ASP) [36], and to use ASP solving algorithms. ASP offers an expressive automated reasoning framework that typically works well for problems at the Θ_2^{p} level of the Polynomial Hierarchy. In this section, we will show how this encoding can be done—similarly to an ASP encoding for the MaxHam rule [42]. Due to space restrictions, we refer to the literature for details on the syntax and semantics of ASP—e.g., [34, 36].

We use the same basic setup that De Haan and Slavkovik [42] use to represent judgment aggregation scenarios—with some simplifications and modifications for the sake of readability. In particular, we use the predicate `voter/1` to represent individuals, we use `issue/1` to represent issues in the agenda, and we use `js/2` to represent judgment sets—both for the individual voters and for a dedicated agent `col` that represents the outcome of the rule.

With this encoding of judgment aggregation scenarios, one can add further constraints on the predicate `js/2` that express which judgment sets are consistent, based on the logical relations between the issues in the agenda Φ—as done by De Haan and Slavkovik [42]. We refer to their work for further details on how this can be done.

Now, we show how to encode the MaxEq rule into ASP, similarly to the encoding of the MaxHam rule by De Haan and Slavkovik [42]. We begin by defining a predicate `dist/2` to capture the Hamming distance `D` between the outcome and the judgment set of an agent `A`.

```
1 dist(A,D) :- voter(A),
      D = #count { X : issue(X), js(col,X), js(A,-X) }.
```

Then, we define predicates `maxdist/1`, `mindist/1` and `inequity/1` that capture the maximum Hamming distance from the outcome to any judgment set in the profile, the minimum such Hamming distance, and the difference between the maximum and minimum (or *inequity*), respectively.

```
2 maxdist(Max) :- Max = #max { D : dist(A,D) }.
3 mindist(Min) :- Min = #min { D : dist(A,D) }.
4 inequity(Max-Min) :- maxdist(Max), mindist(Min).
```

Finally, we add an optimization constraint that states that only outcomes should be selected that minimize the inequity.[14]

```
5 #minimize { I@30 : inequity(I) }.
```

For any answer set program that encodes a judgment aggregation setting, combined with Lines 1–5, it then holds that the optimal answer sets are in one-to-one correspondence with the outcomes selected by the MaxEq rule.

Interestingly, we can readily modify this encoding to capture refinements of the MaxEq rule. An example of this is the refinement that selects (among the outcomes of the MaxEq rule) the outcomes that minimize the maximum Hamming distance to any judgment set in the profile. We can encode this example refinement by adding the following optimization statement that works at a lower priority level than the optimization in Line 5.

```
6 #minimize { Max@20 : maxdist(Max) }.
```

5.3 Encoding Strategic Manipulation

We now show how to encode the problem of strategic manipulation into ASP. The value of this section's contribution should be viewed from the perspective of the modeller rather than from that of the agents. That is, even if we do not wish for the agents to be able to easily check whether they can be better off by lying, it may be

[13] Note though that hardness results regarding manipulation of our egalitarian rules remain an open question.

[14] The expression "`@30`" in Line 5 indicates the priority level of this optimization statement (we used the arbitrary value of 30, and priority levels lexicographically).

reasonable, given a profile of judgments, to externally determine whether a certain agent can benefit from being untruthful.

The simplest way to achieve this is with the meta-programming techniques developed by Gebser et al. [35]. Their meta-programming approach allows one to additionally express optimization statements that are based on subset-minimality, and to transform programs with this extended expressivity to standard (disjunctive) answer set programs. We use this to encode the problem of strategic manipulation.

Due to space reasons, we will not spell out the full ASP encoding needed to do so. Instead, we will highlight the main steps, and describe how these fit together. We will use the example of MaxEq, but the exact same approach would work for any other judgment aggregation rule that can be expressed in ASP efficiently using regular (cardinality) optimization constraints—in other words, for all rules for which the outcome determination problem lies at the Θ_2^{p} level of the Polynomial Hierarchy. Moreover, we will use the example of a decisive preference $\mathring{\succ}$ over sets of judgment sets that is based on a polynomial-time computable preference $\succ$ over judgment sets. The approach can be modified to work with other preferences as well.

We begin by guessing a new judgment set J_i' for the individual i that is trying to manipulate—and we assume, w.l.o.g., that $i = 1$.

```
voter(prime(1)).
1 { js(prime(1),X), js(prime(1),-X) } 1 :- issue(X).
```

Then, we express the outcomes of the MaxEq rule, both for the non-manipulated profile $\boldsymbol{J}$ and for the manipulated profile $\boldsymbol{J}'$, using the dedicated agents `col` (for $\boldsymbol{J}$) and `prime(col)` (for $\boldsymbol{J}'$). This is done exactly as in the encoding of the problem of outcome determination (so for the case of MaxEq, as described in Section 5.2)—with the difference that optimization is expressed in the right format for the meta-programming method of Gebser et al. [35].

We express the following subset-minimality minimization statement (at a higher priority level than all other optimization constraints used so far). This will ensure that every possible judgment set J_i' will be considered as a subset-minimal solution.

```
_criteria(40,1,js(prime(1),X)) :- js(prime(1),S).
_optimize(40,1,incl).
```

To encode whether or not the guessed manipulation was successful, we have to define a predicate `successful/0` that is true if and only if (i) $J' \succ_i J$ and (ii) J and J' are not both selected as outcome by the MaxEq rule for both $\boldsymbol{J}$ and $\boldsymbol{J}'$, where J' is the outcome encoded by the statements `js(prime(col),X)` and J is the outcome encoded by the statements `js(col,X)`. Since we assume that $\succ_i$ is computable in polynomial time, and since we can efficiently check using statements in the answer set whether J and J' are selected by the MaxEq rule for $\boldsymbol{J}$ and $\boldsymbol{J}'$, we know that we can define the predicate `successful/0` correctly and succinctly in our encoding. For space reasons, we omit further details on how to do this.

Then, we express another minimization statement (at a lower priority level than all other optimization statements used so far), that states that we should make `successful` true whenever possible. Intuitively, we will use this to filter our guessed manipulations that are unsuccessful.

```
unsuccessful :- not successful.
successful :- not unsuccessful.
_criteria(10,1,unsuccessful) :- unsuccessful.
_optimize(10,1,card).
```

Finally, we feed the answer set program P that we constructed so far into the meta-programming method, resulting in a new (disjunctive) answer set program P' that uses no optimization statements at all, and whose answer sets correspond exactly to the (lexicographically) optimized answer sets of our program P. Since the new program P' does not use optimization, we can add additional constraint to P' to remove some of the answer sets. In particular, we will filter out those answer sets that correspond to an unsuccessful manipulation—i.e., those containing the statement `unsuccessful`. Effectively, we add the following constraint to P':

```
:- unsuccessful.
```

As a result the only answer sets of P' that remain correspond exactly to successful manipulations J_i' for agent i.

The meta-programming technique that we use uses the full disjunctive answer set programming language. For this full language, finding answer sets is a Σ_2^{p}-complete problem [21]. This is in line with our result of Proposition 8 where we show that the problem of strategic manipulation is in Σ_2^{p}.

The encoding that we described can straightforwardly be modified for various variants of strategic manipulation (e.g., antipodal manipulation). To make this work, one needs to express additional constraints on the choice of the judgment set J_i'. To adapt the encoding for other preference relations $\mathring{\succ}$, one needs to adapt the definition of `successful/0`, expressing under what conditions an act of manipulation is successful.

Our encoding using meta-programming is relatively easily understandable, since we do not need to tinker with the encoding of complex optimization constraints in full disjunctive answer set programming ourselves—this we outsource to the meta-programming method. If one were to do this manually, there is more space for tailor-made optimizations, which might lead to a better performance of ASP solving algorithms for the problem of strategic manipulation. It is an interesting topic for future research to investigate this, and possibly to experimentally test the performance of different encodings, when combined with ASP solving algorithms.

6 CONCLUSION

We have introduced the concept of egalitarianism into the framework of judgment aggregation and have presented how egalitarian and strategyproofness axioms interact in this setting. Importantly, we have shown that the two main interpretations of egalitarianism give rise to rules with differing levels of protection against manipulation. In addition, we have looked into various computational aspects of the egalitarian rules that arise from our axioms, in a twofold manner: First, we have provided worst-case complexity results; second, we have shown how to solve the relevant hard problems using Answer Set Programming.

While we have axiomatised two prominent egalitarian principles, it remains to be seen whether other egalitarian axioms can provide stronger barriers against manipulation. For example, in parallel to majoritarian rules, one could define rules that minimise the distance to some egalitarian ideal. Moreover, as is the case in judgment aggregation, there is an obvious lack of voting rules designed with egalitarian principles in mind. We hope this paper opens the door for similar explorations in voting theory.

A APPENDIX

In this appendix, we will show that the outcome determination problem for the MaxEq rule boils down to a Θ_2^{p}-complete problem. In particular, we will show that the outcome determination problem for the MaxEq rule, when seen as a search problem, is complete (under polynomial-time Levin reductions) for $\mathsf{FP}^{\mathsf{NP}}[\log,\mathrm{wit}]$—which is a search variant of Θ_2^{p}. Then, we show that the problem of finding the minimum difference between two agents' satisfaction, and deciding if this value is divisible by 4, is a Θ_2^{p}-complete problem. Along the way, we show that deciding if there is a judgment set $J^* \in \mathcal{J}(\Phi)$ that has the same Hamming distance to each judgment set in a given profile $\boldsymbol{J}$ is NP-hard, even in the case where the agenda consists of logically independent issues.

We begin, in Section A.1, by recalling some notions from computational complexity theory—in particular, notions related to search problems. Then, in Section A.2, we establish the computational complexity results mentioned above.

A.1 Additional complexity-theoretic preliminaries

We will consider search problems. Let Σ be an alphabet. A *search problem* is a binary relation R over strings in Σ^*. For any input string $x \in \Sigma^*$, we let $R(x) = \{ y \in \Sigma^* \mid (x,y) \in R \}$ denote the set of *solutions* for x. We say that a Turing machine T *solves* R if on input $x \in \Sigma^*$ the following holds: if there exists at least one y such that $(x,y) \in R$, then T accepts x and outputs some y such that $(x,y) \in R$; otherwise, T rejects x. With any search problem R we associate a decision problem S_R, defined by $S_R = \{ x \in \Sigma^* \mid \text{there exists some } y \in \Sigma^* \text{ such that } (x,y) \in R \}$. We will use the following notion of reductions for search problems. A *polynomial-time Levin reduction* from one search problem R_1 to another search problem R_2 is a pair of polynomial-time computable functions (g_1, g_2) such that:

- the function g_1 is a many-one reduction from S_{R_1} to S_{R_2}, i.e., for every $x \in \Sigma^*$ it holds that $x \in S_{R_1}$ if and only if $g_1(x) \in S_{R_2}$.
- for every string $x \in S_{R_1}$ and every solution $y \in R_2(g_1(x))$ it holds that $(x, g_2(x,y)) \in R_1$.

One could also consider other types of reductions for search problems, such as Cook reductions (an algorithm that solves R_1 by making one or more queries to an oracle that solves the search problem R_2). For more details, we refer to textbooks on the topic—e.g., [37].

we will use complexity classes that are based on Turing machines that have access to an oracle. Let C be a complexity class with decision problems. A Turing machine T with access to a *yes-no C oracle* is a Turing machine with a dedicated *oracle tape* and dedicated states q_{oracle}, q_{yes} and q_{no}. Whenever T is in the state q_{oracle}, it does not proceed according to the transition relation, but instead it transitions into the state q_{yes} if the oracle tape contains a string x that is a yes-instance for the problem C, i.e., if $x \in C$, and it transitions into the state q_{no} if $x \notin C$. Let C be a complexity class with search problems. Similarly, a Turing machine with access to a *witness C oracle* has a dedicated oracle tape and dedicated states q_{oracle}, q_{yes} and q_{no}. Also, whenever T is in the state q_{oracle} it transitions into the state q_{yes} if the oracle tape contains a string x such that there exists some y such that $C(x,y)$, and in addition the contents of the oracle tape are replaced by (the encoding of) such an y; it transitions into the state q_{no} if there exists no y such that $C(x,y)$. Such transitions are called *oracle queries*.

We consider the following complexity classes that are based on oracle machines.

- The class $\mathsf{P}^{\mathsf{NP}}[\log]$ consists of all decision problems that can be decided by a deterministic polynomial-time Turing machine that has access to a yes-no NP oracle, and on any input of length n queries the oracle at most $O(\log n)$ many times. This class coincides with the class $\mathsf{P}_{\parallel}^{\mathsf{NP}}$ (spoken: "parallel access to NP"), and is also known as Θ_2^{p}.
  Incidentally, allowing the algorithms access to a witness FNP oracle instead of access to a yes-no NP oracle leads to the same class of problems, i.e., the class $\mathsf{P}^{\mathsf{NP}}[\log,\mathrm{wit}]$ that coincides with $\mathsf{P}^{\mathsf{NP}}[\log]$ (cf. [46, Corollary 6.3.5]).
- The class $\mathsf{FP}^{\mathsf{NP}}[\log,\mathrm{wit}]$ consists of all search problems that can be solved by a deterministic polynomial-time Turing machine that has access to a witness FNP oracle, and on any input of length n queries the oracle at most $O(\log n)$ many times. In a sense, it is the search variant of $\mathsf{P}^{\mathsf{NP}}[\log]$.
  This complexity class happens to coincide with the class $\mathsf{FNP}/\!/\mathsf{OptP}[\log]$, which is defined as the set of all search problems that are solvable by a nondeterministic polynomial-time Turing machine that receives as advice the answer to one "NP optimization" computation [14, 44].

A.2 Complexity proofs for the MaxEq rule

We define the search problem of outcome determination for a judgment aggregation rule F as follows. The input for this problem consists of an agenda Φ, a profile $\boldsymbol{J} = (J_1, \dots, J_n) \in \mathcal{J}(\Phi)^n$. The problem is to output some judgment set $J^* \in F(\boldsymbol{J})$. In other words, the problem is the relation R that consists of all pairs $((\Phi, \boldsymbol{J}), J^*)$ such that $\boldsymbol{J} \in \mathcal{J}(\Phi)^n$ is a profile and $J^* \in F(\boldsymbol{J})$. We will show that the outcome determination problem for the MaxEq rule is complete for the complexity class $\mathsf{FP}^{\mathsf{NP}}[\log,\mathrm{wit}]$ under polynomial-time Levin reductions.

In order to do so, we begin with establishing a lemma that will be useful for the $\mathsf{FP}^{\mathsf{NP}}[\log,\mathrm{wit}]$-hardness proof. This lemma uses the notion of 1-in-3-satisfiability. Let ψ be a propositional logic formula in 3CNF, i.e., $\psi = c_1 \wedge \dots \wedge c_m$, where each c_i is a clause containing exactly three literals. Then ψ is *1-in-3-satisfiable* if there exists a truth assignment α that satisfies exactly one of the three literals in each clause c_i.

Lemma A.1. *Let ψ be a 3CNF formula with clauses $c_1, \dots, c_b$ that are all of size exactly 3 and with n variables $x_1, \dots, x_n$, such that (1) no clause of ψ contains complementary literals, and (2) there exists some $x^* \in \{x_1, \dots, x_n\}$ and a partial truth assignment $\beta : \{x_1, \dots, x_n\} \setminus \{x^*\} \to \{0,1\}$ that satisfies exactly one literal in each clause where x^* or $\neg x^*$ occurs, and satisfies no literal in each clause where x^* or $\neg x^*$ occurs. We can, in polynomial time given ψ, construct an agenda Φ on m issues such that $\mathcal{J}(\Phi) = \{0,1\}^m$ and a profile $\boldsymbol{J}$ over Φ, such that:*

- $\Phi = \{ y_i, y_i' \mid 1 \le i \le n \} \cup \{z_1, \dots, z_5\}$;

- *there exists a judgment set $J \in \mathcal{J}(\Phi)$ that has the same Hamming distance to each $J' \in \boldsymbol{J}$ if and only if ψ is 1-in-3-satisfiable;*
- *if ψ is not 1-in-3-satisfiable, then for each judgment set $J \in \mathcal{J}(\Phi)$ it holds that $\max_{J',J'' \in \boldsymbol{J}} |H(J,J') - H(J,J'')| \geq 2$, and there exists some $J \in \mathcal{J}(\Phi)$ such that $\max_{J',J'' \in \boldsymbol{J}} |H(J,J') - H(J,J'')| = 2$;*
- *the above two properties hold also when restricted to judgment sets J that contain exactly one of y_i and y'_i for each $1 \leq i \leq n$, and that contain $\neg z_1, \dots, \neg z_4, z_5$; and*
- *the number of judgment sets in the profile $\boldsymbol{J}$ only depends on n and b.*

Proof. We let the agenda Φ consist of the $2n+5$ issues $y_1, \dots, y_n, y'_1, \dots, y'_n, z_1, \dots, z_5$. It follows directly that $\mathcal{J}(\Phi) = \{0,1\}^{2n+5}$. Then, we start by constructing $2n$ judgment sets $J_1, \dots, J_{2n}$ over these issues, defined as depicted in Figure 2.

	y_1	y_2	$\cdots$	y_n	y'_1	y'_2	$\cdots$	y'_n	z_1	$\cdots$	z_5
J_1	1	0	$\cdots$	0	1	0	$\cdots$	0	0	$\cdots$	0
J_2	0	1	$\cdots$	0	0	1	$\cdots$	0	0	$\cdots$	0
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
J_n	0	0	$\cdots$	1	0	0	$\cdots$	1	0	$\cdots$	0
J_{n+1}	0	1	$\cdots$	1	0	1	$\cdots$	1	0	$\cdots$	0
J_{n+2}	1	0	$\cdots$	1	1	0	$\cdots$	1	0	$\cdots$	0
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
J_{2n}	1	1	$\cdots$	0	1	1	$\cdots$	0	0	$\cdots$	0

Figure 2: Construction of the judgment sets $J_1, \dots, J_{2n}$ in the proof of Lemma A.1.

Then, for each clause c_k of ψ, we introduce three judgment sets $J_{k,1}$, $J_{k,2}$, and $J_{k,3}$ that are defined as follows. The judgment set $J_{k,1}$ contains $y_i, \neg y'_i$ for each positive literal x_i occurring in c_k, and contains $y'_i, \neg y_i$ for each negative literal $\neg x_i$ occurring in c_k. Conversely, the judgment sets $J_{k,2}, J_{k,3}$ contain $y'_i, \neg y_i$ for each positive literal x_i occurring in c_k, and contain $y_i, \neg y'_i$ for each negative literal $\neg x_i$ occurring in c_k. For each variable x_j that does not occur in c_k, all three of $J_{k,1}, J_{k,2}, J_{k,3}$ contain $\neg y_j, \neg y'_j$. Finally, the judgment set $J_{k,1}$ contains $\neg z_1, \dots, \neg z_4, z_5$, the judgment set $J_{k,2}$ contains $\neg z_1, \neg z_2, z_3, z_4, z_5$, and the judgment set $J_{k,3}$ contains $z_1, z_2, \neg z_3, \neg z_4, z_5$. This is illustrated in Figure 3 for the example clause $(x_1 \vee \neg x_2 \vee x_3)$.

	y_1	y_2	y_3	$\cdots$	y'_1	y'_2	y'_3	$\cdots$	z_1	z_2	z_3	z_4	z_5
$J_{k,1}$	1	0	1	$\cdots$	0	1	0	$\cdots$	0	0	0	0	1
$J_{k,2}$	0	1	0	$\cdots$	1	0	1	$\cdots$	0	0	1	1	1
$J_{k,3}$	0	1	0	$\cdots$	1	0	1	$\cdots$	1	1	0	0	1

Figure 3: Illustration of the construction of the judgment sets $J_{k,1}, J_{k,2}, J_{k,3}$ for the example clause $c_k = (x_1 \vee \neg x_2 \vee x_3)$ in the proof of Lemma A.1.

The profile $\boldsymbol{J}$ then consists of the judgment sets $J_1, \dots, J_{2n}$, as well as the judgment sets $J_{k,1}, J_{k,2}, J_{k,3}$ for each $1 \leq k \leq m$. In order to prove that $\boldsymbol{J}$ has the required properties, we consider the following observations and claims (and prove the claims).

Observation 1. *If a judgment set J contains exactly one of y_i and y'_i for each i, then the Hamming distance from J to each of $J_1, \dots, J_{2n}$—restricted to the issues $y_1, \dots, y_n, y'_1, \dots, y'_n$—is exactly n.*

Claim 2. *If a judgment set J contains both y_i and y'_i for some i, or neither y_i nor y'_i for some i, then there are at least two judgment sets among $J_1, \dots, J_{2n}$ such that the Hamming distance from J to these two judgment sets differs by at least 2.*

Proof of Claim 2. We argue that this is the case for y_1 and y'_1, i.e., for the case of $i = 1$. For other values of i, an entirely similar argument works.

Picking both y_1 and y'_1 to be part of a judgment set J adds to the Hamming distances between J, on the one hand, and $J_1, \dots, J_{2n}$, on the other hand, according to the vector v_1^+:

$$v_1^+ = (0, \underbrace{2, \dots, 2}_{n-1}, 2, \underbrace{0, \dots, 0}_{n-1}).$$

Picking both $\neg y_1$ and $\neg y'_1$ to be part of the set J adds to the Hamming distances between J and $J_1, \dots, J_{2n}$ according to the vector v_1^-:

$$v_1^- = (2, \underbrace{0, \dots, 0}_{n-1}, 0, \underbrace{2, \dots, 2}_{n-1}).$$

Picking exactly one of y_1 and y'_1 and exactly one of $\neg y_1$ and $\neg y'_1$ to be part of J corresponds to the all-ones vector $\mathbf{1} = (1, \dots, 1)$. More generally, both y_i and y'_i to be part of J adds to the Hamming distances according to the vector v_i^+:

$$v_i^+ = (\underbrace{2, \dots, 2}_{i-1}, 0, \underbrace{2, \dots, 2}_{n-i+2}, \underbrace{0, \dots, 0}_{i-1}, 2, \underbrace{0, \dots, 0}_{n-i+2}),$$

picking both $\neg y_i$ and $\neg y'_i$ corresponds to the vector v_i^-:

$$v_i^- = (\underbrace{0, \dots, 0}_{i-1}, 2, \underbrace{0, \dots, 0}_{n-i+2}, \underbrace{2, \dots, 2}_{i-1}, 0, \underbrace{2, \dots, 2}_{n-i+2}),$$

and picking exactly one of y_i and y'_i and exactly one of $\neg y_i$ and $\neg y'_i$ to be part of J corresponds to the all-ones vector $\mathbf{1} = (1, \dots, 1)$. For each $1 \leq i \leq n$, the vectors $v_1^-, \dots, v_n^-$ and v_i^+ are linearly independent, and the vectors $v_1^+, \dots, v_n^+$ and v_i^- are linearly independent.

Suppose now that we pick J to contain both y_1 and y'_1. Suppose, moreover, to derive a contradiction, that the Hamming distance from J to each of the judgment sets $J_1, \dots, J_{2n}$ is the same. This means that there is some way of choosing $s_2, \dots, s_n$ such that $v_1^- = \sum_{1<i\leq n} v_i^{s_i}$, which contradicts the fact that $v_1^+, \dots, v_n^+$ and v_i^- are linearly independent—since each v_j^- can be expressed as $v_j^- = v_i^+ + v_i^- - v_j^+$. Thus, we can conclude that there exist at least two judgment sets among $J_1, \dots, J_{2n}$ such that the Hamming distance from J to these two judgment sets differs. Moreover, since all vectors contain only even numbers, and the coefficients in the sum are integers (in fact, either 0 or 1), we know that the difference must be even and thus at least 2.

An entirely similar argument works for the case where we pick J to contain both $\neg y_1$ and $\neg y'_1$. ⊣

Observation 3. *Let $\alpha : \{x_1, \dots, x_n\} \to \{0,1\}$ be a truth assignment that satisfies exactly one literal in each clause of ψ. Consider the judgment set $J_\alpha = \{\, y_i, \neg y_i' \mid 1 \leq i \leq n, \alpha(x_i) = 1 \,\} \cup \{\, y_i', \neg y_i \mid 1 \leq i \leq n, \alpha(x_i) = 0 \,\} \cup \{\neg z_1, \dots, \neg z_4, z_5\}$. Then the Hamming distance from J_α to each judgment set in the profile $\boldsymbol{J}$ is exactly $n+1$.*

Claim 4. *Suppose that ψ is not 1-in-3-satisfiable. Then for each judgment set J that contains exactly one of y_i and y_i' for each i, there is some clause c_k of ψ such that the difference in Hamming distance from J to (two of) $J_{k,1}, J_{k,2}, J_{k,3}$ is at least 2.*

Proof of Claim 4. Take an arbitrary judgment set J that contains exactly one of y_i and y_i' for each i. This judgment set J corresponds to the truth assignment $\alpha_J : \{x_1, \dots, x_n\} \to \{0,1\}$ defined such that for each $1 \leq i \leq n$ it is the case that $\alpha(x_i) = 1$ if $y_i \in J$ and $\alpha(x_i) = 0$ if $y_i \notin J$. Since ψ is not 1-in-3-satisfiable, we know that there exists some clause c_ℓ such that α_J does not satisfy exactly one literal in c_ℓ. We distinguish several cases: either (i) α_J satisfies no literals in c_ℓ, or (ii) α_J satisfies two literals in c_ℓ, or (iii) α_J satisfies three literals in c_ℓ. In each case, the Hamming distances from J, on the one hand, and $J_{k,1}, J_{k,2}, J_{k,3}$, on the other hand, must differ by at least 2. This can be verified case by case—and we omit a further detailed case-by-case verification of this. ⊣

Claim 5. *If ψ is not 1-in-3-satisfiable, then there exists a judgment set J such that $\max_{J',J'' \in \boldsymbol{J}} |H(J,J') - H(J,J'')| = 2$.*

Proof of Claim 5. Suppose that ψ is not 1-in-3-satisfiable. We know that there exists a variable $x^* \in \{x_1, \dots, x_n\}$ and a partial truth assignment $\beta : \{x_1, \dots, x_n\} \setminus \{x^*\} \to \{0,1\}$ that satisfies exactly one literal in each clause where x^* or $\neg x^*$ does not occur, and satisfies no literal in each clause where x^* or $\neg x^*$ occurs. Without loss of generality, suppose that $x^* = x_1$. Now consider the judgment set $J_\beta = \{\neg y_1, \neg y_1'\} \cup \{\, y_i, \neg y_i' \mid 1 < i \leq n, \beta(x_i) = 1 \,\} \cup \{\, y_i', \neg y_i \mid 1 < i \leq n, \beta(x_i) = 0 \,\} \cup \{\neg z_1, \dots, \neg z_4, z_5\}$. One can verify that the Hamming distances from J_β, on the one hand, and $J' \in \boldsymbol{J}$, on the other hand, differ by at most 2—and for some $J', J'' \in \boldsymbol{J}$ it holds that $|H(J,J') - H(J,J'')| = 2$. ⊣

We now use the above observations and claims to show that Φ and $\boldsymbol{J}$ have the required properties. If ψ is 1-in-3-satisfiable, by Observation 3, there is some $J \in \mathcal{J}(\Phi)$ that has the same Hamming distance to each $J' \in \boldsymbol{J}$. Suppose, conversely, that ψ is not 1-in-3-satisfiable. Then by Claims 2 and 4, there exist two judgment sets $J', J'' \in \boldsymbol{J}$ such that $H(J,J')$ and $H(J,J'')$ differ (by at least 2). Thus, ψ is 1-in-3-satisfiable if and only if there exists a judgment set $J \in \mathcal{J}(\Phi)$ that has the same Hamming distance to each judgment set in the profile $\boldsymbol{J}$.

Suppose that ψ is not 1-in-3-satisfiable. Then by Claims 2 and 4, we know that for each $J \in \mathcal{J}(\Phi)$ it holds that $\max_{J',J'' \in \boldsymbol{J}} |H(J,J') - H(J,J'')| \geq 2$. Moreover, by Claim 5, there exists a judgment set $J \in \mathcal{J}(\Phi)$ such that $\max_{J',J'' \in \boldsymbol{J}} |H(J,J') - H(J,J'')| = 2$.

We already observed that that $\mathcal{J}(\Phi) = \{0,1\}^m$ for some m. Moreover, one can straightforwardly verify that the statements after the first two bullet points in the statement of the lemma also hold when restricted to judgment sets J that contain exactly one of y_i and y_i' for each $1 \leq i \leq n$ and that contain $\neg z_1, \dots, \neg z_4, z_5$. This concludes the proof of the lemma. □

Now that we have established the lemma, we continue with the $\mathsf{FP}^{\mathsf{NP}}[\log,\mathrm{wit}]$-completeness proof.

Theorem 3. *The problem of outcome determination for the MaxEq rule is $\mathsf{FP}^{\mathsf{NP}}[\log,\mathrm{wit}]$-complete under polynomial-time Levin reductions.*

Proof. To show membership in $\mathsf{FP}^{\mathsf{NP}}[\log,\mathrm{wit}]$, we describe an algorithm with access to a witness FNP oracle that solves the problem in polynomial time by making at most a logarithmic number of oracle queries. This algorithm will use an oracle for the following FNP problem: given some $k \in \mathbb{N}$ and given the agenda Φ and the profile $\boldsymbol{J}$, compute a judgment set $J \in \mathcal{J}(\Phi)$ such that $\max_{J',J'' \in \boldsymbol{J}} |H(J,J') - H(J,J'')| \leq k$, if such a J exists, and return "none" otherwise. By using $O(\log|\Phi|)$ queries to this oracle, one can compute the minimum value $k_{\min}$ of $\max_{J',J'' \in \boldsymbol{J}} |H(J,J') - H(J,J'')|$ where the minimum is taken over all judgment sets $J \in \mathcal{J}(\Phi)$. Then, with a final query to the oracle, using $k = k_{\min}$, one can use the oracle to produce a judgment set $J^* \in \mathrm{MaxEq}(\boldsymbol{J})$.

We will show $\mathsf{FP}^{\mathsf{NP}}[\log,\mathrm{wit}]$-hardness by giving a polynomial-time Levin reduction from the $\mathsf{FP}^{\mathsf{NP}}[\log,\mathrm{wit}]$-complete problem of finding a satisfying assignment for a (satisfiable) propositional formula ψ that sets a maximum number of variables to true (among any satisfying assignment of ψ) [14, 44, 47, 65]. Let ψ be an arbitrary satisfiable propositional logic formula with v variables. Without loss of generality, assume that v is even and that there is a satisfying truth assignment for ψ that sets at least one variable to true. We will construct an agenda Φ and a profile $\boldsymbol{J}$, such that the minimum value of $\max_{J',J'' \in \boldsymbol{J}} |H(J^*,J') - H(J^*,J'')|$, for any $J^* \in \mathcal{J}(\Phi)$, is divisible by 4 if and only if the maximum number of variables set to true in any satisfying assignment of ψ is odd. Moreover, we will construct Φ and $\boldsymbol{J}$ in such a way that from any $J^* \in \mathrm{MaxEq}(\Phi)$ we can construct, in polynomial time, a satisfying assignment of ψ that sets a maximum number of variables to true. This—together with the fact that Φ and $\boldsymbol{J}$ can be constructed in polynomial time—suffices to exhibit a polynomial-time Levin reduction, and thus to show $\mathsf{FP}^{\mathsf{NP}}[\log,\mathrm{wit}]$-hardness. We proceed in several stages (i–iv).

(i) We begin, in the first stage, by constructing a 3CNF formula ψ_i with certain properties, for each $1 \leq i \leq v$.

Claim 6. *We can construct in polynomial time, for each $1 \leq i \leq v$, a 3CNF formula ψ_i, that is 1-in-3-satisfiable if and only if there is a truth assignment that satisfies ψ and that sets at least i variables in ψ to true. Moreover, we construct these formulas ψ_i in such a way that they all contain exactly the same variables $x_1, \dots, x_n$ and exactly the same number b of clauses, and such that each formula ψ_i has the properties:*

- *that it contains no clause with complementary literals, and*
- *that it contains a variable x^* and a partial truth assignment $\beta : \{x_1, \dots, x_n\} \setminus \{x^*\} \to \{0,1\}$ that satisfies exactly one literal in each clause where x^* or $\neg x^*$ does not occur, and satisfies no literal in each clause where x^* or $\neg x^*$ occurs.*

Proof of Claim 6. Consider the problems of deciding if a given truth assignment α to the variables in ψ satisfies ψ, and deciding if α satisfies at least i variables, for some $1 \leq i \leq v$. These problems are both polynomial-time solvable. Therefore, by using standard techniques from the proof of the Cook-Levin Theorem [16], we can construct in polynomial time a propositional formula χ in

3CNF containing (among others) the variables $t_1, \dots, t_v$, the variable $x^\dagger$ and the variables in ψ such that any truth assignment to the variables $x^\dagger, t_1, \dots, t_v$ and the variables in ψ can be extended to a satisfying truth assignment for χ if and only if either (i) it sets $x^\dagger$ to true, or (ii) it satisfies ψ and for each $1 \le i \le v$ it sets t_i to false if and only if it sets at least i variables among the variables in ψ to true. Then, we can transform this 3CNF formula χ to another 3CNF formula χ' with a similar property—namely that any truth assignment to the variables $x^\dagger, t_1, \dots, t_v$ and the variables in ψ can be extended to a truth assignment that satisfies exactly one literal in each clause of χ if and only if either (i) it sets $x^\dagger$ to true, or (ii) it satisfies ψ and for each $1 \le i \le v$ it sets t_i to false if and only if it sets at least i variables among the variables in ψ to true. We do so by using the the polynomial-time reduction from 3SAT to 1-IN-3-SAT given by Schaefer [62].

Then, for each particular value of i, we add two clauses that intuitively serve to ensure that variable t_i must be set to false in any 1-in-3-satisfying truth assignment. We add $(s_0 \lor s_1 \lor t_i)$ and $(\neg s_0 \lor \neg s_1 \lor t_i)$, where s_0 and s_1 are fresh variables—the only way to satisfy exactly one literal in both of these clauses is to set exactly one of s_0 and s_1 to true, and to set t_i to false.

Moreover, we add the clauses $(r_0 \lor r_1 \lor x^\dagger)$, $(\neg r_0 \lor r_2 \lor \neg x^\dagger)$, $(x^* \lor r_3 \lor r_0)$, and $(\neg x^* \lor r_4 \lor r_0)$, where $r_0, \dots, r_4$ and x^* are fresh variables. These clauses serve to ensure the property that there always exists a partial truth assignment $\beta : \{x_1, \dots, x_n\} \setminus \{x^*\} \to \{0, 1\}$ that satisfies exactly one literal in each clause where x^* or $\neg x^*$ does not occur, and satisfies no literal in each clause where x^* or $\neg x^*$ occurs. Moreover, these added clauses preserve 1-in-3-satisfiability if and only if there exists a 1-in-3-satisfying truth assignment for the formula without these added clauses that sets $x^\dagger$ to true.

Putting all this together, we have constructed a 3CNF formula χ''—consisting of χ' with the addition of the six clauses mentioned in the above two paragraphs—that has the right properties mentioned in the statement of the claim. In particular, χ'' is 1-in-3-satisfiable if and only if there is a truth assignment that satisfies ψ and that sets at least i variables in ψ to true. Moreover, the constructed formula χ'' has the same variables and the same number of clauses, regardless of the value of i chosen in the construction. ⊣

Since the formulas ψ_i, as described in Claim 6, satisfy the requirements for Lemma A.1, we can construct agendas $\Phi_1, \dots, \Phi_v$ and profiles $\boldsymbol{J}_1, \dots, \boldsymbol{J}_v$ such that for each $1 \le i \le v$, the agenda Φ_i and the profile $\boldsymbol{J}_i$ satisfy the conditions mentioned in the statement of Lemma A.1. Moreover, we can construct the agendas $\Phi_1, \dots, \Phi_v$ in such a way that they are pairwise disjoint. For each $1 \le i \le v$, let $y_{i,1}, \dots, y_{i,n}, y'_{i,1}, \dots, y'_{i,n}, z_{i,1}, \dots, z_{i,5}$ denote the issues in Φ_i and let $\boldsymbol{J}_i = (J_{i,1}, \dots, J_{i,u})$.

(ii) Then, in the second stage, we will use the profiles $\Phi_1, \dots, \Phi_v$ and the profiles $\boldsymbol{J}_1, \dots, \boldsymbol{J}_v$ to construct a single agenda Φ and a single profile $\boldsymbol{J}$.

We let Φ contain the issues $y_{i,j}, y'_{i,j}$ and $z_{i,1}, \dots, z_{i,5}$, for each $1 \le i \le v$ and each $1 \le j \le n$, as well as issues $w_{i,\ell,k}$ for each $1 \le i \le v$, each $1 \le \ell \le u$ and each $1 \le k \le n$. We let $\boldsymbol{J}$ contain judgment sets $J'_{i,\ell}$ for each $1 \le i \le v$ and each $1 \le \ell \le u$, that we will define below. Intuitively, for each i, the sets $J'_{i,\ell}$ will contain the judgment sets $J_{i,1}, \dots, J_{i,u}$ from the profile $\boldsymbol{J}_i$.

Take an arbitrary $1 \le i \le v$, and an arbitrary $1 \le \ell \le u$. We let $J'_{i,\ell}$ agree with $J_{i,\ell}$ (from $\boldsymbol{J}_i$) on all issues from Φ_i—i.e., the issues $y_{i,j}, y'_{i,j}$ for each $1 \le j \le n$ and $z_{i,1}, \dots, z_{i,5}$. On all issues φ from each $\Phi_{i'}$, for $1 \le i' \le v$ with $i' \neq i$, we let $J'_{i,\ell}(\varphi) = 0$. Then, we let $J'_{i,\ell}(w_{i',\ell',k}) = 1$ if and only if $i = i'$ and $\ell = \ell'$. In other words, $J'_{i,\ell}$ agrees with $J_{i,\ell}$ on the issues from Φ_i, it sets every issue from each other $\Phi_{i'}$ to false, it sets all the issues $w_{i,\ell,k}$ to true, and it sets all other issues $w_{i',\ell',k}$ to false.

(iii) In the third stage, we will replace the logically independent issues in Φ by other issues, in order to place restrictions on the different judgment sets that are allowed.

We start by describing a constraint—in the form of a propositional logic formula Γ on the original (logically independent) issues in Φ—and then we describe how this constraint can be used to produce replacement formulas for the issues in Φ. We define $\Gamma = \Gamma_1 \lor \Gamma_2$ by:

$$\Gamma_1 = \bigvee_{\substack{1 \le i \le v \\ 1 \le \ell \le u}} \left(\bigwedge_{1 \le k \le n} w_{i,\ell,k} \land \bigwedge_{\substack{1 \le i' \le v, 1 \le \ell \le u \\ i \neq i' \text{ or } \ell \neq \ell'}} \neg w_{i',\ell',k} \right)$$

$$\Gamma_2 = \left(\bigwedge_{\substack{1 \le i \le v, 1 \le \ell \le u \\ 1 \le k \le n}} \neg w_{i,\ell,k} \right) \land \left(\bigwedge_{\substack{1 \le i \le v \\ 1 \le j \le n}} (y_{i,j} \leftrightarrow \neg y'_{i,j}) \right)$$

In other words, Γ requires that either (1) for some i, ℓ all issues $w_{i,\ell,k}$ are set to true, and all other issues $w_{i',\ell',k}$ are set to false, or (2) all issues $w_{i,\ell,k}$ are set to false and for each i, j exactly one of $y_{i,j}$ and $y'_{i,j}$ is set to true.

Because we can in polynomial time compute a satisfying truth assignment for Γ, we know that we can also in polynomial time compute a replacement φ' for each $\varphi \in \Phi$, resulting in an agenda Φ', so that the logically consistent judgment sets $J \in \mathcal{J}(\Phi')$ correspond exactly to the judgment sets $J \in \mathcal{J}(\Phi)$ that satisfy the constraint Γ [27, Proposition 3]. In the remainder of this proof, we will use Φ (with the restriction that judgment sets should satisfy Γ) interchangeably with Φ'.

(iv) Finally, in the fourth stage, we will duplicate some issues $\varphi \in \Phi$ a certain number of times, by adding semantically equivalent (yet syntactically different) issues to the agenda Φ, and by updating the judgment sets in the profile $\boldsymbol{J}$ accordingly.

For each $1 \le i \le v$, we will define a number c_i and will make sure that there are c_i (syntactically different, logically equivalent) copies of each issue in Φ that originated from the agenda Φ_i. In other words, we duplicate each agenda Φ_i a certain number of times, by adding $c_i - 1$ copies of each issue that originated from Φ_i. For each $1 \le i \le v$, we let $c_i = (v - i)$.

This concludes the description of our reduction—i.e., of the profile Φ and the profile $\boldsymbol{J}$. What remains is to show that the minimum value of $\max_{J',J'' \in \boldsymbol{J}} |H(J^*, J') - H(J^*, J'')|$, for any $J^* \in \mathcal{J}(\Phi)$, is divisible by 4 if and only if the maximum number of variables set to true in any satisfying assignment of ψ is odd. To do so, we begin with stating and proving the following claims.

Claim 7. *For any judgment set $J^* \in \mathcal{J}(\Phi)$ that satisfies Γ_1, the value of $\max_{J',J'' \in \boldsymbol{J}} |H(J^*, J') - H(J^*, J'')|$ is at least $2n$.*

Proof of Claim 7. Take some J^* that satisfies Γ_1. Then there must exist some i, ℓ such that J sets all $w_{i,\ell,k}$ to true and all other $w_{i',\ell',k}$ to false. Therefore, $|H(J^*, J'_{i,\ell}) - H(J^*, J_{i',\ell'})|$ is at least $2n$, for any i', ℓ' such that $(i, \ell) \neq (i', \ell')$. ⊣

Claim 8. *For any judgment set $J^* \in \mathcal{J}(\Phi)$ that satisfies Γ_2, the value of* $\max_{J',J'' \in \boldsymbol{J}} |H(J^*, J') - H(J^*, J'')|$ *is strictly less than $2v$.*

Proof of Claim 8. Take some J^* that satisfies Γ_2. Then J^* sets each $w_{i,\ell,k}$ to false, and for each $1 \leq i \leq v$ the judgment set J contains $\neg z_{i,1}, \dots, \neg z_{i,4}, z_{i,5}$ and contains exactly one of $y_{i,j}$ and $y'_{i,j}$ for each $1 \leq j \leq n$. Moreover, without loss of generality, we may assume that J^*, for each $1 \leq i \leq v$, assigns truth values to the issues originating from Φ_i in a way that corresponds either (a) to a truth assignment to the variables in ψ_i witnessing that ψ_i is 1-in-3-satisfiable, or (b) to a partial truth assignment β : $\text{var}(\psi_i) \setminus \{x_1\} \to \{0, 1\}$ that satisfies exactly one literal in each clause where x_1 or $\neg x_1$ occurs, and satisfies no literal in each clause where x_1 or $\neg x_1$ occurs. If this were not the case, we could consider another J^* instead that does satisfy these properties and that has a value of $\max_{J',J'' \in \boldsymbol{J}} |H(J^*, J') - H(J^*, J'')|$ that is at least as small.

Then, by the construction of $\boldsymbol{J}$—and the profiles $\boldsymbol{J_i}$ used in this construction—for each $J'_{i,j}$ it holds that $H(J^*, J'_{i,j}) = \sum_{1 \leq i \leq v} c_i(n + 1) + n \pm d_i$, where $d_i = 0$ if ψ_i is 1-in-3-satisfiable and $d_i = (v - i)$ otherwise. From this it follows that the value of $\max_{J',J'' \in \boldsymbol{J}} |H(J^*, J') - H(J^*, J'')|$ is strictly less than $2v$, since ψ_1 is 1-in-3-satisfiable. ⊣

By Claims 7 and 8, and because we may assume without loss of generality that $n \geq v$, we know that any judgment set $J^* \in \mathcal{J}(\Phi)$ that minimizes $\max_{J',J'' \in \boldsymbol{J}} |H(J^*, J') - H(J^*, J'')|$ must satisfy Γ_2. Moreover, by a straightforward modification of the arguments in the proof of Claim 8, we know that the minimal value, over judgment sets $J^* \in \mathcal{J}(\Phi)$, of $\max_{J',J'' \in \boldsymbol{J}} |H(J^*, J') - H(J^*, J'')|$ is $2(v - i)$ for the smallest value of i such that ψ_i is not 1-in-3-satisfiable, which coincides with $2(v + 1 - i)$ for the largest value of i such that ψ_i is 1-in-3-satisfiable. Since v is even (and thus $v + 1$ is odd), we know that $2(v + 1 - i)$ is divisible by 4 if and only if i is odd. Therefore, the minimal value of $\max_{J',J'' \in \boldsymbol{J}} |H(J^*, J') - H(J^*, J'')|$ is divisible by 4 if and only if the maximum number of variables set to true in any satisfying assignment of ψ is odd.

Moreover, it is straightforward to show that from any $J^* \in \text{MaxEq}(\Phi)$ we can construct, in polynomial time, a satisfying assignment of ψ that sets a maximum number of variables to true.

This concludes our description and analysis of the polynomial-time Levin reduction, and thus of our hardness proof. □

Proposition 5. *Given an agenda Φ and a profile $\boldsymbol{J}$, deciding whether the minimal value of* $\max_{J',J'' \in \boldsymbol{J}} |H(J^*, J') - H(J^*, J'')|$ *for $J^* \in \mathcal{J}(\Phi)$, is divisible by 4, is a Θ_2^{P}-complete problem.*

Proof. This follows from the proof of Theorem 3. The proof of membership in $\text{FP}^{\text{NP}}[\log,\text{wit}]$ for the outcome determination problem for the MaxEq rule can directly be used to show membership in Θ_2^{P} for this problem. Moreover, the polynomial-time Levin reduction used in the $\text{FP}^{\text{NP}}[\log,\text{wit}]$-hardness proof can be seen as a polynomial-time (many-to-one) reduction from the Θ_2^{P}-complete problem of deciding if the maximum number of variables set to true in any satisfying assignment of a (satisfiable) propositional formula ψ is odd [47, 65]. □

Proposition 6. *Given an agenda Φ and a profile $\boldsymbol{J}$, the problem of deciding whether there is some $J^* \in \mathcal{J}(\Phi)$ with* $\max_{J',J'' \in \boldsymbol{J}} |H(J^*, J') - H(J^*, J'')| = 0$ *is* NP-complete. Moreover, NP-hardness holds even for the case where Φ consists of logically independent issues—i.e., the case where $\mathcal{J}(\Phi) = \{0, 1\}^m$ for some m.

Proof (sketch). The problem of deciding if there exists a truth assignment that satisfies a given 3CNF formula ψ and that sets at least, say, 2 variables among ψ to true is an NP-complete problem. Then, by combining Claim 6 in the proof of Theorem 3 with Lemma A.1, we directly get that the problem of deciding whether there is some $J^* \in \mathcal{J}(\Phi)$ with $\max_{J',J'' \in \boldsymbol{J}} |H(J^*, J') - H(J^*, J'')| = 0$, for a given agenda Φ and a given profile $\boldsymbol{J}$ is NP-hard, even for the case where Φ consists of logically independent issues. Membership in NP follows from the fact that one can guess a set J^* (together with a truth assignment witnessing that it is consistent) in polynomial time, after which checking whether $J^* \in \mathcal{J}(\Phi)$—i.e., whether it is indeed consistent—and checking whether $\max_{J',J'' \in \boldsymbol{J}} |H(J^*, J') - H(J^*, J'')| = 0$ can then be done in polynomial time. □

REFERENCES

[1] Georgios Amanatidis, Georgios Birmpas, and Evangelos Markakis. 2016. On Truthful Mechanisms for Maximin Share Allocations. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*.

[2] Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. 2017. Justified Representation in Approval-based Committee Voting. *Social Choice and Welfare* 48, 2 (2017), 461–485.

[3] Salvador Barberà, Walter Bossert, and Prasanta K Pattanaik. 2004. Ranking Sets of Objects. In *Handbook of Utility Theory*. Springer, 893–977.

[4] Seth D. Baum. 2017. Social Choice Ethics in Artificial Intelligence. *AI & Society* (2017), 1–12.

[5] Dorothea Baumeister, Gábor Erdélyi, Olivia Johanna Erdélyi, and Jörg Rothe. 2013. Computational aspects of manipulation and control in judgment aggregation. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT)*.

[6] Dorothea Baumeister, Gábor Erdélyi, Olivia J Erdélyi, and Jörg Rothe. 2015. Complexity of Manipulation and Bribery in Judgment Aggregation for Uniform Premise-Based Quota Rules. *Mathematical Social Sciences* 76 (2015), 19–30.

[7] Dorothea Baumeister, Jörg Rothe, and Ann-Kathrin Selker. 2017. Strategic Behavior in Judgment Aggregation. In *Trends in Computational Social Choice*. Lulu.com, 145–168.

[8] Sirin Botan and Ulle Endriss. 2020. Majority-Strategyproofness in Judgment Aggregation. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

[9] Steven J Brams, Michael A Jones, and Christian Klamler. 2008. Proportional Pie-Cutting. *International Journal of Game Theory* 36, 3-4 (2008), 353–367.

[10] Steven J Brams, D Marc Kilgour, and M Remzi Sanver. 2007. A Minimax Procedure for Electing Committees. *Public Choice* 132, 3 (2007), 401–420.

[11] Eric Budish. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.

[12] Ioannis Caragiannis, Dimitris Kalaitzis, and Evangelos Markakis. 2010. Approximation algorithms and mechanism design for minimax approval voting. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*, Vol. 24.

[13] Yiling Chen, John K Lai, David C Parkes, and Ariel D Procaccia. 2013. Truth, Justice, and Cake Cutting. *Games and Economic Behavior* 77, 1 (2013), 284–297.

[14] Zhi-Zhong Chen and Seinosuke Toda. 1995. The Complexity of Selecting Maximal Solutions. *Information and Computation* 119 (1995), 231–239. Issue 2.

[15] Vincent Conitzer, Walter Sinnott-Armstrong, Jana Schaich Borg, Yuan Deng, and Max Kramer. 2017. Moral Decision Making Frameworks for Artificial Intelligence. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*.

[16] Stephen A. Cook. 1971. The Complexity of Theorem-Proving Procedures. In *Proceedings of the 3rd Annual ACM Symposium on Theory of Computing*. Shaker Heights, Ohio, 151–158.

[17] H Dalton. 1920. The Measurement of the Inequality of Incomes. *Economic Journal* 30, 119 (1920), 348–461.

[18] Franz Dietrich and Christian List. 2007. Strategy-proof Judgment Aggregation. *Economics & Philosophy* 23, 3 (2007), 269–300.

[19] John Duggan and Thomas Schwartz. 2000. Strategic Manipulability without Resoluteness or Shared Beliefs: Gibbard-Satterthwaite Generalized. *Social Choice and Welfare* 17, 1 (2000), 85–93.

[20] Michael Dummet. 1984. *Voting Procedures*. Oxford University Press.

[21] Thomas Eiter and Georg Gottlob. 1995. On the Computational Cost of Disjunctive Logic Programming: Propositional Case. *Annals of Mathematics and Artifficial Intelligence* 15, 3-4 (1995), 289–323.

[22] Edith Elkind, Piotr Faliszewski, Piotr Skowron, and Arkadii Slinko. 2017. Properties of Multiwinner Voting Rules. *Social Choice and Welfare* 48, 3 (2017), 599–632.

[23] Ulle Endriss. 2016. Judgment Aggregation. In *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia (Eds.). Cambridge University Press.

[24] Ulle Endriss. 2016. Judgment Aggregation. In *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia (Eds.). Cambridge University Press.

[25] Ulle Endriss and Ronald de Haan. 2015. Complexity of the Winner Determination Problem in Judgment Aggregation: Kemeny, Slater, Tideman, Young. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

[26] Ulle Endriss, Ronald de Haan, Jérôme Lang, and Marija Slavkovik. 2020. The Complexity Landscape of Outcome Determination in Judgment Aggregation. *Journal of Artificial Intelligence Research* (2020).

[27] Ulle Endriss, Umberto Grandi, Ronald de Haan, and Jérôme Lang. 2016. Succinctness of Languages for Judgment Aggregation. In *Proceedings of the 15th International Conference on the Principles of Knowledge Representation and Reasoning (KR)*.

[28] Ulle Endriss, Umberto Grandi, and Daniele Porello. 2012. Complexity of Judgment Aggregation. *Journal of Artificial Intelligence Research* 45 (2012), 481–514.

[29] Patricia Everaere, Sébastien Konieczny, and Pierre Marquis. 2014. On Egalitarian Belief Merging. In *Proceedings of the 14th International Conference on the Principles of Knowledge Representation and Reasoning (KR)*.

[30] Patricia Everaere, Sébastien Konieczny, and Pierre Marquis. 2017. Belief Merging and its Links with Judgment Aggregation. In *Trends in Computational Social Choice*, Ulle Endriss (Ed.). AI Access Foundation, 123–143.

[31] Peter C Fishburn and Steven J Brams. 1983. Paradoxes of Preferential Voting. *Mathematics Magazine* 56, 4 (1983), 207–214.

[32] Duncan K Foley. 1967. Resource Allocation and the Public Sector. *Yale Economic Essays* 7, 1 (1967), 45–98.

[33] Peter Gärdenfors. 1976. Manipulation of Social Choice Functions. *Journal of Economic Theory* 13, 2 (1976), 217–228.

[34] Martin Gebser, Roland Kaminski, Benjamin Kaufmann, and Torsten Schaub. 2012. *Answer Set Solving in Practice*. Morgan & Claypool Publishers.

[35] Martin Gebser, Roland Kaminski, and Torsten Schaub. 2011. Complex Optimization in Answer Set Programming. *Theory and Practics of Logic Programming* 11, 4-5 (2011), 821–839.

[36] Michael Gelfond. 2006. Answer Sets. In *Handbook of Knowledge Representation*, Frank van Harmelen, Vladimir Lifschitz, and Bruce Porter (Eds.). Elsevier.

[37] Oded Goldreich. 2010. *P, NP, and NP-Completeness: The Basics of Complexity Theory*. Cambridge University Press.

[38] Umberto Grandi and Ulle Endriss. 2010. Lifting Rationality Assumptions in Binary Aggregation. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*.

[39] Davide Grossi and Gabriella Pigozzi. 2014. *Judgment Aggregation: A Primer*. Synthesis Lectures on Artificial Intelligence and Machine Learning, Vol. 8. Morgan & Claypool Publishers. 1–151 pages.

[40] Ronald de Haan. 2018. Hunting for Tractable Languages for Judgment Aggregation. In *Proceedings of the 16th International Conference on the Principles of Knowledge Representation and Reasoning (KR)*.

[41] Ronald de Haan and Marija Slavkovik. 2017. Complexity Results for Aggregating Judgments using Scoring or Distance-Based Procedures. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS*.

[42] Ronald de Haan and Marija Slavkovik. 2019. Answer Set Programming for Judgment Aggregation. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*.

[43] Jerry S Kelly. 1977. Strategy-proofness and Social Choice Functions without Singlevaluedness. *Econometrica* 45, 2 (1977), 439–446.

[44] Johannes Köbler and Thomas Thierauf. 1990. Complexity Classes with Advice. In *Proceedings of the 5th Annual Structure in Complexity Theory Conference*.

[45] Sébastien Konieczny and Ramón Pino Pérez. 2011. Logic Based Merging. *Journal of Philosophical Logic* 40, 2 (2011), 239–270.

[46] Jan Krajicek. 1995. *Bounded arithmetic, propositional logic and complexity theory*. Cambridge University Press.

[47] Mark W. Krentel. 1988. The Complexity of Optimization Problems. *J. Comput. System Sci.* 36, 3 (1988), 490–509.

[48] Justin Kruger and Zoi Terzopoulou. 2020. Strategic Manipulation with Incomplete Preferences: Possibilities and Impossibilities for Positional Scoring Rules. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

[49] Martin Lackner and Piotr Skowron. 2018. Approval-Based Multi-Winner Rules and Strategic Voting. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*.

[50] Jérôme Lang, Gabriella Pigozzi, Marija Slavkovik, and Leendert van der Torre. 2011. Judgment Aggregation Rules Based on Minimization. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*.

[51] Jérôme Lang and Marija Slavkovik. 2014. How Hard is it to Compute Majority-Preserving Judgment Aggregation Rules?. In *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI)*.

[52] Christian List and Philip Pettit. 2002. Aggregating Sets of Judgments: An Impossibility Result. *Economics and Philosophy* 18, 1 (2002), 89–110.

[53] Michael K. Miller and Daniel Osherson. 2009. Methods for Distance-based Judgment Aggregation. *Social Choice and Welfare* 32, 4 (2009), 575–601.

[54] Elchanan Mossel and Omer Tamuz. 2010. Truthful Fair Division. In *Proceedings of the 3rd International Symposium on Algorithmic Game Theory (SAGT)*.

[55] Hervé Moulin. 1988. *Axioms of Cooperative Decision Making*. Econometric Society Monographs, Vol. 15. Cambridge University Press.

[56] Klaus Nehring, Marcus Pivato, and Clemens Puppe. 2014. The Condorcet Set: Majority Voting over Interconnected Propositions. *Journal of Economic Theory* 151 (2014), 268–303.

[57] Ritesh Noothigattu, Snehalkumar S Gaikwad, Edmond Awad, Sohan Dsouza, Iyad Rahwan, Pradeep Ravikumar, and Ariel D Procaccia. 2018. A Voting-Based System for Ethical Decision Making. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*.

[58] Dominik Peters. 2018. Proportionality and Strategyproofness in Multiwinner Elections. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

[59] Gabriella Pigozzi. 2006. Belief Merging and the Discursive Dilemma: An Argument-based Account to Paradoxes of Judgment Aggregation. *Synthese* 152, 2 (2006), 285–298.

[60] John Rawls. 1971. *A Theory of Justice*. Belknap Press.

[61] M Remzi Sanver and William S Zwicker. 2009. One-way Monotonicity as a Form of Strategy-proofness. *International Journal of Game Theory* 38, 4 (2009), 553–574.

[62] Thomas J. Schaefer. 1978. The complexity of satisfiability problems. In *Conference Record of the 10th Annual ACM Symposium on Theory of Computing*. ACM.

[63] Amartya Sen. 1997. *Choice, Welfare and Measurement*. Harvard University Press.

[64] Zoi Terzopoulou and Ulle Endriss. 2018. Modelling Iterative Judgment Aggregation. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*.

[65] Klaus W. Wagner. 1990. Bounded Query Classes. *SIAM J. Comput.* 19, 5 (1990), 833–846.

[66] William S. Zwicker. 2016. Introduction to the Theory of Voting. In *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia (Eds.). Cambridge University Press.