

Judgment Aggregation with Rationality and Feasibility Constraints

Ulle Endriss

Institute for Logic, Language and Computation
University of Amsterdam

ABSTRACT

I introduce a model of judgment aggregation that allows for an explicit distinction between rationality and feasibility constraints. The former are assumed to be satisfied by the individual agents; the latter must be met by the collective decision returned by the aggregation rule in use. Using this model, I characterise the class of combinations of rationality and feasibility constraints for which the majority rule can guarantee feasible outcomes and I propose several majoritarian aggregation rules that, in some sense, approximate the ideal of the majority when using the majority rule itself is not feasible. Finally, to illustrate the power and flexibility of the model, I show how it can be used to simulate several common voting rules in a simple and elegant manner. This includes the well-known Borda rule, for which finding a natural counterpart in judgment aggregation has long been an elusive quest.

KEYWORDS

Judgment Aggregation; Social Choice Theory; Voting Rules

ACM Reference Format:

Ulle Endriss. 2018. Judgment Aggregation with Rationality and Feasibility Constraints. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018*, IFAAMAS, 9 pages.

1 INTRODUCTION

The field of judgment aggregation is concerned with the design and analysis of procedures for combining the judgments of several agents regarding the truth of a number of logically related statements into a collective judgment [23, 34]. Inspired by questions arising in legal theory [27], it has received significant attention in both philosophy and economics, starting with the seminal contribution of List and Pettit [33]. More recently, due to its potential for applications in areas such as multiagent systems and crowd-sourcing, and due to the interesting algorithmic questions it raises, judgment aggregation has also received increasing attention in computer science and artificial intelligence [1, 13, 21, 47].

I propose to enrich the standard model of judgment aggregation by distinguishing between *rationality constraints* (to be respected by the individual agents when supplying their judgments) and *feasibility constraints* (to be respected by the outcomes returned by an aggregation rule). In contrast to this proposal, in essentially all existing work on judgment aggregation there is only a single type of constraint (which sometimes is explicitly represented and

sometimes left implicit), governing what is permissible for both the input and the output.

This new model will allow us to handle a wider range of application scenarios than what has hitherto been possible. For example, we may consider it irrational for an individual to support both a law to subsidise local health education and a law providing tax breaks to international fast food chains, while we may very well accept this combination as a compromise when returned by an aggregation rule (i.e., in some contexts, rationality constraints may be more demanding than feasibility constraints). At the same time, for the outcomes of an aggregation rule we may be bound by budget considerations, while we may not want to think of individuals as being concerned with such feasibility constraints when pondering their personal judgments (i.e., in other contexts, feasibility constraints may be more demanding than rationality constraints).

EXAMPLE 1. The 5-member local council of a medium-sized town has to decide on the funding for three projects: refurbishing the primary school (φ_1), organising a “Summer of Culture” (φ_2), and building a second parking lot next to the shopping mall (φ_3). The budget is limited and it is not feasible to fund all three projects, i.e., $\neg(\varphi_1 \wedge \varphi_2 \wedge \varphi_3)$ is a feasibility constraint. However, the individual councillors are not expected to keep this constraint in mind when making their personal judgments regarding the merit of each project. Instead, each of them can be assumed to want to please at least one type of clientele, i.e., it would be irrational for a councillor not to recommend any of the projects for funding. Thus, $\varphi_1 \vee \varphi_2 \vee \varphi_3$ is a rationality constraint. Suppose they vote as follows:

	φ_1	φ_2	φ_3
Councillor 1	1	1	0
Councillor 2	0	0	1
Councillor 3	1	0	1
Councillor 4	1	1	0
Councillor 5	1	1	1

Thus, all councillors respect the rationality constraint. If we decide on each issue by majority, then the outcome is (1, 1, 1), in violation of the feasibility constraint. A better rule might be to select from amongst the feasible outcomes one that maximises overall agreement, resulting in a tie between (1, 1, 0) and (1, 0, 1). \triangle

Existing work on judgment aggregation speaks of either “consistency” [33] or “rationality” [22] to define what individual judgments are permissible. Feasibility is not discussed as a separate concept in this literature. Rather, the common assumption is that for the output one would want to impose the same requirements as for the input. This is known as *collective rationality*. On the other hand, the concept of a feasibility constraint on aggregation outcomes is

prominent in work on logic-based belief merging [26], a framework that is closely related to judgment aggregation [17]. But in belief merging one usually makes no assumptions regarding individual rationality (other than respecting the laws of classical logic).

While the explicit distinction between rationality and feasibility constraints is a new idea, similar ideas—arguably—are implicit in at least some prior work on judgment aggregation. For instance, in a proposal for a model of judgment aggregation to handle scenarios that require reasoning about a number of premises and a single conclusion, Miller [36] suggests that individual agents may have subjective views on the logical relationship between the premises and the conclusion. Using my terminology, this would amount to each individual agent being subject to a possibly different rationality constraint. In related work, Benamara et al. [2] allow for the possibility that a given agent may or may not accept the (shared) rationality constraint. Another example is the work of Lang and Slavkovik [31]. While—in line with much of the literature—they only specify a single constraint (which conceptually takes on the role of a feasibility constraint), some of their results relate to scenarios where all individual agents happen to conform to a stronger constraint than the feasibility constraint imposed explicitly. That stronger constraint corresponds to the rationality constraint in my model. Costantini et al. [7] speculate about the possibility of learning feasibility constraints from data: if most individual agents happen to respect a certain constraint, even if that constraint is not explicitly imposed as a rationality constraint, then it arguably makes sense to try to respect this constraint in the output of a judgment aggregation rule. Finally, my proposal to clearly distinguish rationality from feasibility constraints is conceptually (though not technically) similar to an idea due to Porello [43, 44] who proposes to use different logical calculi to assess the consistency of individual judgments (which is what I call rationality) and the consistency of collective judgments (which is what I call feasibility).

Besides the conceptual contribution of proposing an enriched model of judgment aggregation and formulating a number of natural aggregation rules for this model, the technical contribution of this paper is twofold. First, I fully characterise the set of all combinations of rationality and feasibility constraints that avoid the “majority paradox” illustrated by Example 1, where the majority rule applied to individual judgments that are rational returns a collective outcome that fails to be feasible. Second, I show how to embed several voting rules into the model in a simple and elegant manner. This includes, amongst others, the Borda rule (for which finding an appropriate counterpart in judgment aggregation has long been an open problem and for which there still is no broad consensus as to what would be the “right” kind of embedding). It also includes the Uncovered Set (which appears not to have been discussed in the context of judgment aggregation before). Refining an idea originally due to Lang and Slavkovik [31], my approach makes it possible to simulate several common voting rules using the same basic judgment aggregation rules by only varying the feasibility constraint.

The remainder of the paper is organised as follows. Section 2 is a presentation of my model of judgment aggregation with rationality and feasibility constraints. In Section 3, I characterise the range of aggregation problems for which the majority rule is well-behaved and then define several other majoritarian rules that approximate

the ideal represented by the majority whilst guaranteeing the feasibility of outcomes by definition. My approach to simulating voting rules in judgment aggregation is introduced in Section 4. This approach is then applied to the majoritarian aggregation rules defined earlier, yielding several common voting rules. Section 5 concludes.

2 THE MODEL

In this section, I introduce the model of judgment aggregation with rationality and feasibility constraints. It heavily borrows from the standard model of judgment aggregation going back to the work of List and Pettit [33], particularly in its incarnation known as “binary aggregation with integrity constraints” [21, 22].

2.1 Notational Preliminaries

For any given set S , we use $\mathcal{P}(S)$ to denote its *powerset*, i.e., the set of all subsets of S , and $\mathcal{P}_+(S)$ to denote the set of *nonempty* subsets of S , i.e., $\mathcal{P}_+(S) = \mathcal{P}(S) \setminus \{\emptyset\}$.

Recall the familiar *argmax*-operator, which for a given set S and a given function $f : S \rightarrow \mathbb{R}$ can be used to denote the set of elements x in S for which $f(x)$ is maximal:

$$\operatorname{argmax}_{x \in S} f(x) = \{x \in S \mid f(y) > f(x) \text{ for no } y \in S\}$$

In analogy to *argmax*, we define the *argsetmax*-operator, which for two given sets S and S' as well as a function $f : S \rightarrow \mathcal{P}(S')$ can be used to denote the set of elements x in S for which $f(x)$ is maximal with respect to set-inclusion:

$$\operatorname{argsetmax}_{x \in S} f(x) = \{x \in S \mid f(y) \supset f(x) \text{ for no } y \in S\}$$

2.2 Judgments and Aggregation Rules

Let $N = \{1, \dots, n\}$ with $n > 1$ be a finite set of *agents*. For ease of exposition, throughout this paper we assume that n is odd (so we can avoid having to consider tied majorities).

Each agent is asked to answer a number of questions with either “yes” or “no”. We model these questions as a finite set Φ , referred to as the *agenda*. A *judgment* is a function $J : \Phi \rightarrow \{0, 1\}$, mapping each agenda item $\varphi \in \Phi$ to either 0 (indicating rejection) or 1 (indicating acceptance). We often write $\{0, 1\}^\Phi$ as a shorthand for $\Phi \rightarrow \{0, 1\}$, the space of all possible judgments. For any two judgments $J : \Phi \rightarrow \{0, 1\}$ and $J' : \Phi \rightarrow \{0, 1\}$, let their *equaliser* be defined as the set of agenda items on which they agree:

$$\operatorname{Eq}(J, J') = \{\varphi \in \Phi \mid J(\varphi) = J'(\varphi)\}$$

A *profile* $\mathbf{J} = (J_1, \dots, J_n) \in (\{0, 1\}^\Phi)^n$ is a vector of judgments, one for each agent. That is, in profile \mathbf{J} every agent $i \in N$ provides us with her judgment J_i . The *majority judgment* $\operatorname{Maj}(\mathbf{J}) : \Phi \rightarrow \{0, 1\}$ derived from profile \mathbf{J} is defined as follows, for every $\varphi \in \Phi$:

$$\operatorname{Maj}(\mathbf{J})(\varphi) = \begin{cases} 0 & \text{if } |\{i \in N \mid J_i(\varphi) = 0\}| > \frac{n}{2} \\ 1 & \text{if } |\{i \in N \mid J_i(\varphi) = 1\}| > \frac{n}{2} \end{cases}$$

Note that, due to our assumption that n is odd, $\operatorname{Maj}(\mathbf{J})$ is always a *complete* judgment, taking a well-defined position on every single agenda item.

An *aggregation rule* is a function F that takes as input a profile and that returns as output a single judgment that is supposed to represent a suitable compromise between the judgments made by the individual agents. In fact, most rules allow for the possibility of

ties between two or more judgments in the output. Thus, formally an aggregation rule is a function $F : (\{0, 1\}^\Phi)^n \rightarrow \mathcal{P}_+(\{0, 1\}^\Phi)$, mapping any given profile to a nonempty set of judgments.¹

An example for an aggregation rule is the *majority rule*, defined as $F_{\text{Maj}} : J \mapsto \{\text{Maj}(J)\}$. That is, for any given profile, it returns a singleton set containing only the majority judgment.

2.3 Rationality and Feasibility Constraints

By a slight abuse of notation, we identify each agenda item $\varphi \in \Phi$ with a propositional variable. Let $\mathcal{L}(\Phi)$ be the *propositional language* over this set of variables. That is, $\mathcal{L}(\Phi)$ is the set of all well-formed formulas of propositional logic that we can construct using the propositional variables in Φ and the familiar connectives \neg (negation), \wedge (conjunction), \vee (disjunction), and \rightarrow (implication).

We are going to think of such formulas as constraints on judgments. In particular, we say that a judgment $J : \Phi \rightarrow \{0, 1\}$ *satisfies* a constraint $\Gamma \in \mathcal{L}(\Phi)$, denoted $J \models \Gamma$, if Γ evaluates to *true* under the assignment of propositional variables to truth values induced by J in the natural manner, for the usual semantics of classical propositional logic [8]. In other words, this notion of satisfaction is defined recursively as follows:

- $J \models \varphi$ for propositional variables $\varphi \in \Phi$ if and only if $J(\varphi) = 1$
- $J \models \neg\Gamma$ if and only if $J \models \Gamma$ is not the case
- $J \models \Gamma \wedge \Gamma'$ if and only if both $J \models \Gamma$ and $J \models \Gamma'$
- $J \models \Gamma \vee \Gamma'$ if and only if $J \models \Gamma$ or $J \models \Gamma'$ (or both)
- $J \models \Gamma \rightarrow \Gamma'$ if and only if $J \not\models \Gamma$ or $J \models \Gamma'$ (or both)

Thus, for example, if J is given by $(\varphi_1 \mapsto 1, \varphi_2 \mapsto 0, \varphi_3 \mapsto 1)$, or $(1, 0, 1)$ for short, then $J \models \varphi_1 \vee \neg\varphi_3$ but $J \not\models \varphi_2$. We write $\Gamma \models \Gamma'$ in case Γ *entails* Γ' , i.e., in case $J \models \Gamma$ implies $J \models \Gamma'$ for every judgment $J \in \{0, 1\}^\Phi$. For a given formula $\Gamma \in \mathcal{L}(\Phi)$, the set of *models* of Γ is the set of judgments that satisfy Γ :

$$\text{Mod}(\Gamma) = \{J \in \{0, 1\}^\Phi \mid J \models \Gamma\}$$

We make use of constraints in two complementary ways. First, we use *rationality constraints* $\Gamma \in \mathcal{L}(\Phi)$ to constrain the range of profiles we consider relevant. A profile $\mathbf{J} = (J_1, \dots, J_n) \in (\{0, 1\}^\Phi)^n$ is called Γ -*rational* if $J_i \models \Gamma$ for all $i \in N$, i.e., if $\mathbf{J} \in \text{Mod}(\Gamma)^n$. Second, we use *feasibility constraints* $\Gamma' \in \mathcal{L}(\Phi)$ to define what outcomes we consider acceptable. We are now ready to state the central definition of this paper.

DEFINITION 1. *An aggregation rule F is said to **guarantee** Γ' -feasible outcomes on Γ -rational profiles, if for every profile $\mathbf{J} \in \text{Mod}(\Gamma)^n$ it is the case that $F(\mathbf{J}) \subseteq \text{Mod}(\Gamma')$.*

In other words, F guarantees Γ' -feasible outcomes on Γ -rational profiles, if $J \models \Gamma'$ for all $J \in F(\mathbf{J})$ for every profile \mathbf{J} with $J_i \models \Gamma$ for all $i \in N$. In the remainder of this paper, Γ always denotes a rationality constraint and Γ' always denotes a feasibility constraint.

3 MAJORITARIAN AGGREGATION RULES

The majority rule is very appealing on normative grounds, as it treats all agents in a “fair” manner. However, as we are going to see,

¹How to break ties so as to select a single judgment from the set of judgments returned by an aggregation rule is a question of some interest that, however, is beyond the scope of this paper. For instance, in the case of Example 1, we may favour $(1, 0, 1)$ over $(1, 1, 0)$, as the latter would leave Councillor 2 extremely unhappy.

it cannot always guarantee that the outcomes it returns will be feasible. Whether this problem arises depends on both the rationality and the feasibility constraint. In this section, I first prove a characterisation theorem that settles for which combinations of rationality and feasibility constraints the majority rule is well-behaved.

When the conditions of this characterisation theorem are not met, then we should not use the majority rule. Nevertheless, we may want to preserve some of its appealing features. In the second part of this section, I therefore define several aggregation rules that are “as close as possible” to the majority rule, whilst nevertheless being able to guarantee the feasibility of outcomes. These rules differ in how they interpret the qualification “as close as possible”.²

3.1 The Majority Rule

Recall that, for any given profile \mathbf{J} , the majority rule F_{Maj} returns a set containing only the majority judgment $\text{Maj}(\mathbf{J})$.

EXAMPLE 2. Consider the case of $\Phi = \{\varphi_1, \varphi_2, \varphi_3\}$ and $\Gamma = \Gamma' = (\varphi_1 \vee \varphi_2 \vee \varphi_3)$. For this scenario, the majority rule F_{Maj} does *not* guarantee feasible outcomes on all rational profiles, as demonstrated by the following counterexample $\mathbf{J} = (J_1, J_2, J_3)$:

	φ_1	φ_2	φ_3	
Agent 1	1	0	0	$J_1 \models \varphi_1 \vee \varphi_2 \vee \varphi_3$
Agent 2	0	1	0	$J_2 \models \varphi_1 \vee \varphi_2 \vee \varphi_3$
Agent 3	0	0	1	$J_3 \models \varphi_1 \vee \varphi_2 \vee \varphi_3$
Majority	0	0	0	$\text{Maj}(\mathbf{J}) \not\models \varphi_1 \vee \varphi_2 \vee \varphi_3$

This is a variant of the famous *doctrinal paradox*, showing that the majority rule may violate constraints (“doctrines”) that are satisfied by every single individual [27]. △

Of course, this problem does not occur for *every* combination of constraints. So when can we guarantee feasible outcomes, given rational profiles? For the standard model of judgment aggregation (with $\Gamma = \Gamma'$) the answer to this question is well understood. To be able to state it here, we require some further terminology.

Recall that a *clause* is a disjunction of literals. Let us call a formula $\Gamma \in \mathcal{L}(\Phi)$ *simple*, if it is logically equivalent to a conjunction of clauses with at most two literals each. Otherwise Γ is called *nonsimple*. Note that these definitions can also be applied to clauses: a simple clause has at most two literals and a nonsimple clause cannot be simplified to a clause with fewer than three literals.

By a known result, in case $\Gamma = \Gamma'$, the majority rule will guarantee Γ' -feasible outcomes on Γ -rational profiles if and only if Γ is simple [22, Theorem 28].³ Thus, to use the language of that paper, the majority rule can “lift” the constraint Γ from the individual to the collective level if and only if Γ is simple.

We are now going to state and prove a generalisation of this result for the model of judgment aggregation with rationality and feasibility constraints. This requires some additional machinery.

²All of the rules defined here are known rules for the special case of $\Gamma = \Gamma'$. Some were originally inspired by similar rules for preference aggregation. There unfortunately is no consensus regarding terminology and a variety of different—often confusing—names are used to refer to these rules in the literature. I propose here a very simple naming scheme that directly alludes to what is being optimised by each rule.

³This is the “syntactic counterpart” of a seminal result in judgment aggregation, first proved by Nehring and Puppe [41] in an algebraic framework, which characterises the domains on which the majority rule is well-behaved as the so-called *median spaces*.

Recall that a *prime implicate* of a formula $\Gamma \in \mathcal{L}(\Phi)$ is a clause $\pi \in \mathcal{L}(\Phi)$ such that (i) $\Gamma \models \pi$ and (ii) π is logically equivalent to every clause $\pi' \in \mathcal{L}(\Phi)$ with $\Gamma \models \pi'$ and $\pi' \models \pi$ [35]. Thus, the prime implicates of Γ are the logically strongest clauses that are entailed by Γ . For example, if $\Gamma = (p \vee q) \wedge (p \rightarrow r)$, then the set of prime implicates of Γ includes $(p \vee q)$, $(\neg p \vee r)$, and $(q \vee r)$. Observe that a formula Γ is simple if and only if all its prime implicates are simple.⁴ We are going to make use of the following well-known property of prime implicates [35, p. 59].

LEMMA 1. *If $\Gamma \models \Gamma'$ is the case, then for every prime implicate π' of Γ' there exists a prime implicate π of Γ such that $\pi \models \pi'$.*

Using the concept of prime implicates, we now extend the definition of simplicity from single formulas to pairs of formulas.

DEFINITION 2. *A pair of formulas $(\Gamma, \Gamma') \in \mathcal{L}(\Phi)^2$ is **simple**, if for every nonsimple prime implicate π' of Γ' there exists a simple prime implicate π of Γ such that $\pi \models \pi'$.*

In particular, if Γ' is simple (and thus has no nonsimple prime implicates), then (Γ, Γ') is simple for every $\Gamma \in \mathcal{L}(\Phi)$.

EXAMPLE 3. Consider the agenda $\Phi = \{\varphi_1, \varphi_2, \varphi_3, \varphi_4, \varphi_5\}$ and the following two constraints:

$$\begin{aligned}\Gamma &= (\varphi_1 \vee \varphi_2) \wedge (\varphi_3 \vee \varphi_4 \vee \varphi_5) \\ \Gamma' &= (\varphi_1 \vee \varphi_2 \vee \varphi_3)\end{aligned}$$

Then neither Γ nor Γ' are simple, because they involve 3-clauses that cannot be eliminated. Nevertheless, (Γ, Γ') is simple, because the 3-clause $\varphi_1 \vee \varphi_2 \vee \varphi_3$ is entailed by the 2-clause $\varphi_1 \vee \varphi_2$. \triangle

We are now ready to state a characterisation theorem that identifies the precise conditions under which the majority rule is safe to use.

THEOREM 2. *The majority rule guarantees Γ' -feasible outcomes on Γ -rational profiles if and only if $\Gamma \models \Gamma'$ and (Γ, Γ') is simple.*

PROOF. We first show that $\Gamma \models \Gamma'$ is a necessary condition for F_{Maj} guaranteeing Γ' -feasible outcomes on Γ -rational profiles. Suppose $\Gamma \not\models \Gamma'$ and thus $\text{Mod}(\Gamma) \not\subseteq \text{Mod}(\Gamma')$. This means that there exists a judgment $J \in \text{Mod}(\Gamma) \setminus \text{Mod}(\Gamma')$. Now consider the unanimous profile $\mathbf{J} = (J, \dots, J)$. Then, even though \mathbf{J} satisfies the rationality constraint Γ , the majority judgment $\text{Maj}(\mathbf{J}) = J$ returned by F_{Maj} fails to satisfy the feasibility constraint Γ' .

So, from now on assume $\Gamma \models \Gamma'$. It remains to be shown that, given this assumption, F_{Maj} guarantees Γ' -feasible outcomes on Γ -rational profiles if and only if (Γ, Γ') is simple.

(\Rightarrow) We prove the contrapositive. So suppose (Γ, Γ') is not simple. Thus, there exists a nonsimple prime implicate π' of Γ' such that for no simple prime implicate π of Γ it is the case that $\pi \models \pi'$. However, by Lemma 1, we know that there must exist *some* prime implicate π of Γ such that $\pi \models \pi'$. Hence, this π must be nonsimple. As they are nonsimple, both π and π' must be clauses with at least three literals, and as $\pi \models \pi'$ they must share at least three literals. W.l.o.g., we may assume that these shared literals are all positive (if not, given that the majority rule treats acceptance and rejection of

agenda items symmetrically, we can switch the roles of, say, φ_1 and $\neg\varphi_1$ everywhere), i.e., $\pi = \varphi_1 \vee \varphi_2 \vee \varphi_3 \vee \psi$ and $\pi' = \varphi_1 \vee \varphi_2 \vee \varphi_3 \vee \psi'$, with the φ_i being agenda items and ψ and ψ' being (possibly empty) clauses with $\psi \models \psi'$. Now consider a profile $\mathbf{J} \in \text{Mod}(\Gamma)^n$ in which the first $\lfloor \frac{n}{3} \rfloor$ agents accept φ_1 but neither φ_2 nor φ_3 , the next $\lfloor \frac{n}{3} \rfloor$ agents accept φ_2 but neither φ_1 nor φ_3 , and the remaining agents accept φ_3 but neither φ_1 nor φ_2 . Furthermore, $J_i \not\models \psi'$ (and thus $J_i \not\models \psi$) for all agents $i \in N$. Note that $J_i \models \pi$ for all $i \in N$ and thus it is indeed possible to extend this to a Γ -rational profile. However, $\text{Maj}(\mathbf{J})$ rejects all of φ_1 , φ_2 , and φ_3 , and it also does not satisfy ψ' . Hence, $\text{Maj}(\mathbf{J}) \not\models \pi'$, which implies $\text{Maj}(\mathbf{J}) \not\models \Gamma'$, i.e., the majority judgment is not Γ' -feasible.

(\Leftarrow) Suppose (Γ, Γ') is simple. For the sake of contradiction, assume there exists a profile $\mathbf{J} \in \text{Mod}(\Gamma)^n$ with $\text{Maj}(\mathbf{J}) \not\models \Gamma'$. Then we must have $\text{Maj}(\mathbf{J}) \not\models \pi'$ for one of the prime implicates π' of Γ' . There must be a simple prime implicate π of Γ with $\pi \models \pi'$: if π' is nonsimple, then this follows from the simplicity of (Γ, Γ') ; if π' is simple, then this follows from Lemma 1 and the fact that simplicity of π' together with $\pi \models \pi'$ implies simplicity of π . Due to $\text{Maj}(\mathbf{J}) \not\models \pi'$, we must have also $\text{Maj}(\mathbf{J}) \not\models \pi$. W.l.o.g., we may assume that π is of the form $\varphi_1 \vee \varphi_2$.⁵ As \mathbf{J} is Γ -rational, we must have $J_i(\varphi_1) = 1$ or $J_i(\varphi_2) = 1$ for all $i \in N$. Hence, as n is odd, there must be a strict majority accepting at least one of φ_1 and φ_2 . W.l.o.g., assume this is so for φ_1 , i.e., $\text{Maj}(\mathbf{J}) \models \varphi_1$. But then we get $\text{Maj}(\mathbf{J}) \models \pi$ and thus $\text{Maj}(\mathbf{J}) \models \pi'$ and $\text{Maj}(\mathbf{J}) \models \Gamma'$, which contradicts our original assumption. \square

So the majority rule is not well-behaved in all cases. In the remainder of this section, I therefore define several alternative aggregation rules that guarantee feasible outcomes by definition, whilst preserving some of the attractiveness of the majority rule.

3.2 Rules Based on Simple Majorities

Let Γ' be a feasibility constraint. The first rule goes through all judgments that satisfy Γ' and selects those for which the set of agenda items on which there is agreement with the majority judgment cannot be extended further.

DEFINITION 3. *In the context of the feasibility constraint Γ' , the **max-set rule** maps any given profile \mathbf{J} to the following outcome:*

$$\text{max-set}(\mathbf{J}, \Gamma') = \underset{J \in \text{Mod}(\Gamma')}{\text{argsetmax}} \text{Eq}(\mathbf{J}, \text{Maj}(\mathbf{J}))$$

Lang et al. [30] call this the *maximal Condorcet rule*, while Nehring et al. [40] call its outcome $\text{max-set}(\mathbf{J}, \Gamma')$ the *Condorcet set*.

Strictly speaking, max-set is a *family of aggregation rules*, one for each feasibility constraint Γ' (the same is true for all other rules to be introduced in this section). For a family of aggregation rules F such as max-set, we are going to write $F(\cdot, \Gamma')$ for the aggregation rule induced by feasibility constraint Γ' .

Adopting the terminology of Nehring et al. [40], let us call an aggregation rule F that guarantees Γ' -feasible outcomes *majoritarian*, if it never unnecessarily rejects a majority opinion, i.e., if $F(\mathbf{J}) \subseteq \text{max-set}(\mathbf{J}, \Gamma')$ for all profiles \mathbf{J} . Observe that a majoritarian rule F will return *only* the majority judgment whenever doing so is feasible: $\text{Maj}(\mathbf{J}) \models \Gamma'$ implies $F(\mathbf{J}) = \{\text{Maj}(\mathbf{J})\}$.

⁵In case π is not even a proper 2-clause but merely a single literal, the remainder of the proof is analogous (but simpler).

⁴The right-to-left part of this statement follows from the fact that Γ is equivalent to the conjunction of all its prime implicates. The left-to-right parts follows from the completeness of the resolution calculus (i.e., given a representation of Γ in terms of 2-clauses, we can infer every prime implicate of Γ) and the fact that applying the resolution rule to two 2-clauses can never yield a k -clause with $k > 2$.

The rule defined next seeks to maximise the *number* of agenda items on which the majority opinion is being respected.

DEFINITION 4. *In the context of the feasibility constraint Γ' , the **max-num rule** maps any given profile J to the following outcome:*

$$\text{max-num}(J, \Gamma') = \operatorname{argmax}_{J \in \text{Mod}(\Gamma')} |\text{Eq}(J, \text{Maj}(J))|$$

This has been called the *endpoint rule* by Miller and Osherson [37] and the *maximum-cardinality subagenda rule* by Lang et al. [29]. It also has been called the (generalised) *Slater rule* [13, 40], because—as we are going to see in Section 4—it is closely related to the well-known Slater rule in preference aggregation.

3.3 Rules Based on Weighted Majorities

If we also take into account the *strengths* of the majorities to be respected, then we naturally arrive at a rule that maximises the overall number of opinions that are being respected, summing both over individual agents and over agenda items.

DEFINITION 5. *In the context of the feasibility constraint Γ' , the **max-sum rule** maps any given profile J to the following outcome:*

$$\text{max-sum}(J, \Gamma') = \operatorname{argmax}_{J \in \text{Mod}(\Gamma')} \sum_{i \in N} |\text{Eq}(J, J_i)|$$

This has been called the *prototype rule* by Miller and Osherson [37] and the *maximum-weight subagenda rule* by Lang et al. [29]. More common names are “the” *distance-based rule* [15, 42], the *median rule* [30, 40], and the (generalised) *Kemeny rule* [9, 14]. The latter name is due to the close connection between max-sum and the Kemeny rule in preference aggregation (see Section 4).

Let me conclude this section by pointing out that there are further options for defining majoritarian rules that take the strengths of majorities into account. I briefly sketch two of them here. One option is to return judgments that are maximally majoritarian in a lexicographic sense, i.e., we first try to satisfy as many of the strongest majorities as we can, then as many of the second-strongest majorities as we can, and so forth. Such a *lexi-max rule* has been discussed by Nehring and Pivato [39] and Everaere et al. [16]. Another option is to implement this kind of sequential rule in a greedy fashion, i.e., to adopt majority decisions in order of their strengths but now without enforcing lexicographic optimality.⁶ This is closely related to Tideman’s *ranked-pairs rule* in preference aggregation [48]. In judgment aggregation such a *greedy-max rule* has been discussed by Lang et al. [29] and by Porello and Endriss [45], under the names of *ranked-agenda rule* and *support-based rule*, respectively.

4 EMBEDDING COMMON VOTING RULES

It is well known that judgment aggregation is a generalisation of the standard model of preference aggregation studied in social choice theory [13, 32]. Exploring this connection has been useful, for instance, to clarify the nature of impossibility theorems in both areas [10] and to obtain complexity results in judgment aggregation [15]. Nevertheless, attempts to exploit this connection to generalise voting rules (for preference aggregation) to judgment aggregation have only succeeded to a limited extent [29, 31, 40]. In

⁶This makes a difference only in case there are several agenda items with the exact same majority strength (in favour or against).

particular, there is no consensus in the literature as to what kind of judgment aggregation rule would correspond to the Borda rule, one of the most important voting rules [9, 11].

In this section, I am going to demonstrate that the clear distinction between rationality and feasibility constraints facilitates the embedding of a range of voting rules into judgment aggregation. In particular, we are going to see that—contrary to what earlier attempts may have suggested [31]—it is possible to simulate the Borda rule in judgment aggregation using the standard embedding of preferences into judgment aggregation. In addition, I provide the first simulation of the Uncovered Set within judgment aggregation and a new embedding of the Copeland rule. All of the voting rules discussed in this section will turn out to correspond to one of the three basic majoritarian aggregation rules introduced earlier.

This section is strongly inspired by work of Lang and Slavkovik [31], although I believe that my new model provides a better framework in which to carry out this kind of research agenda.

4.1 Voting Theory Preliminaries

We briefly recall the relevant fundamentals from the theory of voting [49]. Let X be a finite set of *alternatives* with $m = |X|$. We now think of $N = \{1, \dots, n\}$ as a set of *voters*. Again, we assume that n is odd. Every voter $i \in N$ is endowed with a *preference order* \succsim_i , which is a total order on X . We use $X!$ to denote the set of all possible preference orders. We write \succ_i for the strict part of \succsim_i . The *majority relation* \succ_M is defined on X as follows:

$$x \succ_M y \text{ if and only if } |\{i \in N \mid x \succ_i y\}| > \frac{n}{2}$$

Given the stated preferences of the voters, the question arises how one should select the “best” alternative. A *voting rule* (or *social choice function*) encodes a possible answer to this question. It is a function $F : X!^n \rightarrow \mathcal{P}_+(X)$, mapping any given profile of preferences to a set of (tied) election winners. Here are six examples:

- *Borda*. Under the Borda rule, we select the alternatives that maximise the score $B(x) = \sum_{i \in N} |\{y \mid x \succ_i y\}|$. Thus, an alternative x receives as many points from a voter i as there are other alternatives that i ranks below x [3].
- *Copeland*. Under the Copeland rule, we select the alternatives maximising the score $C(x) = |\{y \mid x \succ_M y\}| - |\{y \mid y \succ_M x\}|$. Thus, we select the alternatives that maximise the difference between won and lost pairwise majority contests [6].
- *Uncovered Set*. For any $x \in X$, let $D(x) = \{y \mid x \succ_M y\}$ denote the set of alternatives dominated by x . We say that x *covers* y if $D(x) \supseteq D(y)$. The Uncovered Set is the set of alternatives that are not covered by any other alternatives [20, 38].
- *Top Cycle*. The Top Cycle is the smallest nonempty subset S of X such that every alternative in S wins all majority contests against all alternatives not in S [5].
- *Slater*. Under the Slater rule, we select the top alternatives of those total orders on X that minimise disagreement with \succ_M . Here, disagreement is measured in terms of the number of pairs of alternatives ranked differently [46].
- *Kemeny*. Under the Kemeny rule, we select the top alternatives of those total orders on X that minimise the sum of the disagreements with the individual preference orders \succ_i [25].

An alternative x is called a *Condorcet winner* in case $x \succ_M y$ for all alternatives $y \neq x$. While there may be no Condorcet winner for a given profile, if there is one, then it is unique and—with the exception of the Borda rule—all of the solution concepts listed above will select that Condorcet winner and no other alternatives.

4.2 Simulating Preference Aggregation

We adapt the standard approach of embedding preference aggregation into judgment aggregation [10, 13] to our model. Given a set of alternatives X , we define the *preference agenda* Φ_{\succ}^X as follows:

$$\Phi_{\succ}^X = \{p_{x \succ y} \mid x, y \in X\}$$

Thus, for every pair of alternatives (x, y) we introduce an agenda item $p_{x \succ y}$, acceptance of which intuitively signifies that the agent in question likes x at least as much as y . When formulating constraints, we use $p_{x > y}$ as a shorthand for $p_{x \succ y} \wedge \neg p_{y \succ x}$.

We can now use formulas, with propositional variables ranging over the preference agenda, to describe properties of a binary relation \succ . Examples include the familiar properties of completeness, transitivity, and antisymmetry (which together define total orders):

$$\begin{aligned} \text{COMPLETE} &= \bigwedge_{x, y} (p_{x \succ y} \vee p_{y \succ x}) \\ \text{ANTI-SYM} &= \bigwedge_{x \neq y} \neg(p_{x \succ y} \wedge p_{y \succ x}) \\ \text{TRANSITIVE} &= \bigwedge_{x, y, z} (p_{x \succ y} \wedge p_{y \succ z} \rightarrow p_{x \succ z}) \end{aligned}$$

In fact, we need not restrict attention to properties of binary relations that are usually associated with preferences:

$$\begin{aligned} \text{NOCHAIN} &= \bigwedge_{x, y, z} \neg(p_{x > y} \wedge p_{y > z}) \\ \text{ROOTED} &= \bigvee_x \bigwedge_{y \neq x} p_{x > y} \end{aligned}$$

The NOCHAIN property expresses that there exists no chain of strictly ranked alternatives of length 3 or more. A judgment satisfies ROOTED when there is exactly one top alternative that is ranked strictly above all others.⁷ We can now use the basic properties defined earlier to describe more complex properties:

$$\begin{aligned} \text{WEAKORDER} &= \text{COMPLETE} \wedge \text{TRANSITIVE} \\ \text{RANKING} &= \text{WEAKORDER} \wedge \text{ANTI-SYM} \\ \text{DICHOTOMOUS} &= \text{WEAKORDER} \wedge \text{NOCHAIN} \\ \text{WINNER} &= \text{ROOTED} \wedge \text{DICHOTOMOUS} \end{aligned}$$

Thus, we have $J \models \text{RANKING}$ if and only if the binary relation \succ described by the agenda items accepted by J is a total order (i.e., a strict ranking of the alternatives). Note that the difference between ROOTED and WINNER is that the latter fixes the relationship between the alternatives that are not at the top (intuitively, WINNER declares indifference between all of them), while the former permits any kind of structure with a single top element.

⁷While ROOTED explicitly postulates the existence of *at least* one such top alternative, due to the semantics of $>$ there cannot be more than one.

EXAMPLE 4. Let $X = \{x, y, z\}$. The corresponding preference agenda has nine elements. Consider this profile for three agents:

	$p_{x \succ x}$	$p_{x \succ y}$	$p_{x \succ z}$	$p_{y \succ x}$	$p_{y \succ y}$	$p_{y \succ z}$	$p_{z \succ x}$	$p_{z \succ y}$	$p_{z \succ z}$
Agent 1	1	1	1	0	1	1	0	0	1
Agent 2	1	0	0	1	1	1	1	0	1
Agent 3	1	1	0	0	1	0	1	1	1

Thus, agent 1 has preference order $x \succ_1 y \succ_1 z$, agent 2 has $y \succ_2 z \succ_2 x$, and agent 3 has $z \succ_3 x \succ_3 y$. If we aggregate this profile using the majority rule, then we obtain the judgment $(1, 1, 0, 0, 1, 1, 1, 0, 1)$, which corresponds to the binary relation \succ with $x \succ y \succ z \succ x$. Thus, we obtain a cycle! This is an instance of the classic *Condorcet paradox* [49]. In the rendering in judgment aggregation given here, this paradox manifests itself in the fact that each of the three individual judgments satisfies RANKING, but the majority judgment does not. We can also make a subtly stronger statement: the majority judgment does not satisfy ROOTED, which means that there is no Condorcet winner in this profile. \triangle

We are going to impose RANKING as a rationality constraint to model the fact that each voter supplies us with a ranking of the alternatives.⁸ Feasibility constraints of interest are those that entail ROOTED, as this allows us to associate a single winning alternative with every single judgment in the outcome. Observe that RANKING, in particular, entails ROOTED, as every total order on a finite set has exactly one top element. We are going to focus on RANKING and WINNER for feasibility constraints, but also comment on the possibility of using ROOTED instead.

Thus, to simulate a voting rule in judgment aggregation, we first translate the given preference profile into a profile of judgments (which should all satisfy RANKING). We then apply a judgment aggregation rule. Finally, for each of the judgments J in the outcome (for which we must assume that they all satisfy ROOTED), we extract that one alternative x for which $p_{x > y}$ is satisfied by J for all other alternatives y . The following definition formalises this approach.

DEFINITION 6. Fix a set of alternatives X . Let F be an aggregation rule for the corresponding preference agenda Φ_{\succ}^X that guarantees ROOTED-feasible outcomes on RANKING-rational profiles. Then we say that F , when restricted to RANKING-rational profiles, *simulates* the voting rule F' if, for every preference profile $(\succ_1, \dots, \succ_n)$ and corresponding profile $J = (J_1, \dots, J_n)$, the following holds:

$$F'(\succ_1, \dots, \succ_n) = \bigcup_{J \in F(J)} \{x \in X \mid J \models p_{x > y} \text{ for all } y \neq x\}$$

As an aside, let us count how many distinct models the constraints we use to represent inputs and outputs have. There are m models of the constraint WINNER (as each of the m alternatives could be the top alternative) and $m!$ models of RANKING (one for each permutation of the m alternatives). For comparison, the number of models of ROOTED is $m \cdot 2^{(m-1)(m-2)}$: there are m possibilities for choosing a top alternative x , and then each of the $(m-1)(m-2)$ propositions $p_{y \succ z}$ with $x \neq y \neq z \neq x$ can be made either true or false.

⁸While most voting rules, including the six defined in Section 4.1, presuppose that ballots take the form of such rankings, there are exceptions. The most important is *approval voting*, where you vote by approving a subset of the alternatives [4]. For approval voting, DICHOTOMOUS would be an appropriate rationality constraint.

Rationality	Feasibility	max-set	max-num	max-sum
RANKED	RANKED	Top Cycle	Slater	Kemeny
RANKED	WINNER	Uncovered Set	Copeland	Borda

Table 1: Embedding Common Voting Rules

4.3 Results

Applying our three majoritarian aggregation rules (max-set, max-num, and max-sum) to profiles of judgments that all satisfy RANKING and using either RANKING or WINNER as the feasibility constraint yields—potentially—up to six different voting rules. Interestingly, they all turn out to be very well-known rules. These results are summarised in Table 1. In fact, the results in the first row (which can be expressed in the standard model of judgment aggregation, given that $\Gamma = \Gamma'$) were previously known. In particular, those for the Kemeny and the Slater rule are immediate from the relevant definitions and best classified as folklore. The former is implicit, for instance, in the work of Endriss et al. [15] and the latter, for instance, in that of Nehring et al. [40].

The fact that, when restricted to RANKING-rational profiles, max-set(\cdot , RANKING) simulates the Top Cycle was first noted by Lang and Slavkovik [31].⁹ This may be seen as follows. For a given preference profile, suppose we are looking for a ranking \succ that maximises agreement on the relative ordering of pairs of alternatives with the majority relation (with respect to set-inclusion). Then it cannot be the case that, in \succ , an element of the Top Cycle is ranked below an alternative not in the Top Cycle—because in that case swapping an adjacent pair of this kind would lead to strictly greater agreement with the majority relation. Thus, max-set(\cdot , RANKING) returns all rankings in which Top Cycle alternatives are ranked above non-Top Cycle alternatives. And the maximal elements in those rankings form the Top Cycle.

In the remainder of this section, I am going to prove and comment on the results in the second row of Table 1.

THEOREM 3. *When restricted to RANKING-rational profiles, max-sum(\cdot , WINNER) simulates the **Borda rule**.*

PROOF. For any given alternative $x \in X$, let J^x denote the judgment with $J^x(p_{y \succ z}) = 1$ if and only if $y = x$ or $z \neq x$. Observe that $\text{Mod}(\text{WINNER}) = \{J^x \mid x \in X\}$. Now, for any given profile $J \in \text{Mod}(\text{RANKING})^n$, we obtain:

$$\begin{aligned} \text{max-sum}(J, \text{WINNER}) &= \operatorname{argmax}_{J \in \text{Mod}(\text{WINNER})} \sum_{i \in N} |\text{Eq}(J, J_i)| \\ &= \operatorname{argmax}_{J^x \mid x \in X} \sum_{i \in N} |\text{Eq}(J^x, J_i)| \\ &= \operatorname{argmax}_{J^x \mid x \in X} \sum_{i \in N} |\{\varphi \in \Phi_{\neq}^X \mid J^x(\varphi) = J_i(\varphi)\}| \end{aligned}$$

Fixing x , we can decompose the preference agenda like this:

$$\begin{aligned} \Phi_{\neq}^X &= \{p_{y \succ z} \mid y = z\} \cup \{p_{y \succ z} \mid x \neq y \neq z \neq x\} \cup \\ &\quad \{p_{x \succ y} \mid y \neq x\} \cup \{p_{y \succ x} \mid y \neq x\} \end{aligned}$$

First, J^x and J_i agree on $\{p_{y \succ z} \mid y = z\}$ for all $J_i \in \text{Mod}(\text{RANKING})$: they both accept all m agenda items in this set. Second, while J^x also accepts all of $\{p_{y \succ z} \mid x \neq y \neq z \neq x\}$, any $J_i \in \text{Mod}(\text{RANKING})$

⁹To be precise, their result subtly differs from the result stated here.

accepts exactly half of them, i.e., $(m-1)(m-2)/2$ agenda items. Note that these two figures do not depend on the choice of x , so we can ignore them as we evaluate argmax . Finally, J^x accepts all of $\{p_{x \succ y} \mid y \neq x\}$ and rejects all of $\{p_{y \succ x} \mid y \neq x\}$. J_i agrees with these judgments exactly for those alternatives y that are ranked below x by agent i , i.e., the number of agreements for this part of the agenda is $2 \cdot |\{y \mid x \succ_i y\}|$. Note that the factor of 2 is irrelevant for the evaluation of argmax , so we can drop it and obtain:

$$\text{max-sum}(J, \text{WINNER}) = \operatorname{argmax}_{J^x \mid x \in X} \sum_{i \in N} |\{y \mid x \succ_i y\}|$$

But this is a direct simulation of the Borda rule, so we are done. \square

So the Borda and the Kemeny rule reduce to the same aggregation rule when viewed through the lense of judgment aggregation. This may come as a surprise for a number of reasons, one of them being the difference in computational complexity of the two rules. While the Kemeny rule is highly intractable [24], the Borda rule, of course, can be computed in polynomial time. The explanation is simple. While both rules require us to maximise the sum of relevant agreements, the key difference is the number of potential outcomes we have to inspect: for WINNER the number of models is polynomial, while for RANKING it is super-exponential.

Still, the fact that there are connections between the Borda and the Kemeny rule is not entirely new either. First, both can be defined in terms of *majority margins* [19]. Second, using the *distance-based approach* to the rationalisation of voting rules [12], the Kemeny rule is naturally rationalised using the *swap-distance* and, by a result of Farkas and Nitzan [18], the same is possible for Borda. Third, Lang et al. [28] observe similarities between Borda and Kemeny when interpreted in the context of graph aggregation.

As previously mentioned, how to model the Borda rule in judgment aggregation has long been an open problem. There now are alternative proposals to the one I have given here, but the question of what is the “right” approach certainly has not been settled. One approach is due to Dietrich [9]. Using my terminology, he defines a *scoring function* as a function that associates each pair (φ, J_i) of an agenda item and a (rational) judgment with a score. The corresponding aggregation rule is the rule that returns the (feasible) judgments J that maximise the sum of the scores we obtain when we apply the scoring function to all pairs of agenda items φ in J and judgments J_i in the profile. Then, when using the preference agenda and a specific scoring function (the *reversal scoring function*, scoring how many agenda items need to be flipped in J_i to preserve rationality if φ is flipped), we obtain a counterpart of the Borda rule. This is an attractive approach, in particular as it also permits modelling other *positional scoring rules* [49]. On the downside, the definition is very complex and, arguably, does not well reflect the algorithmic simplicity of the original Borda rule. Duddy et al. [11] propose an aggregation rule that is similar in spirit to the Borda rule, but which does not amount to a full simulation of Borda in judgment aggregation. Lang and Slavkovik [31] attempt a similar kind of embedding as I have performed here, using ROOTED instead of WINNER. They show that this attempt fails, although it does result in a voting rule that, again, is similar in spirit.

THEOREM 4. *When restricted to RANKING-rational profiles, max-num(\cdot , WINNER) simulates the **Copeland rule**.*

PROOF. Define J^x as in the proof of Theorem 3. Consider any profile $J \in \text{Mod}(\text{RANKING})^n$. We obtain:

$$\begin{aligned} \text{max-num}(J, \text{WINNER}) &= \operatorname{argmax}_{J \in \text{Mod}(\text{WINNER})} |\text{Eq}(J, \text{Maj}(J))| \\ &= \operatorname{argmax}_{J^x | x \in X} |\text{Eq}(J^x, \text{Maj}(J))| \\ &= \operatorname{argmax}_{J^x | x \in X} |\{\varphi \in \Phi_{\neq}^X \mid J^x(\varphi) = \text{Maj}(J)(\varphi)\}| \end{aligned}$$

Now consider the decomposition of Φ_{\neq}^X also considered in the proof of Theorem 3. J^x and $\text{Maj}(J)$ agree on all m propositions of the form $p_{y \succ z}$ with $y = z$ and on exactly $(m-1)(m-2)/2$ propositions of the form $p_{y \succ z}$ with $x \neq y \neq z \neq x$. Furthermore, they agree on $p_{x \succ y}$ if and only if they agree on $p_{y \succ x}$. Thus, to evaluate argmax , we can simply count the number of times J^x and $\text{Maj}(J)$ agree on propositions of the form $p_{x \succ y}$. But recall that $J^x(p_{x \succ y}) = 1$ for all alternatives y . Thus, we obtain:

$$\text{max-num}(J, \text{WINNER}) = \operatorname{argmax}_{J^x | x \in X} |\{p_{x \succ y} \mid \text{Maj}(J)(p_{x \succ y}) = 1\}|$$

Hence, x is selected if it maximises the number of alternatives y it beats in pairwise majority contests. Given that the number n of individuals is odd, which excludes the possibility of tied majority contests, this formulation is equivalent to the Copeland rule. \square

If we substitute WINNER with ROOTED , we still obtain the Copeland rule. This was shown by Lang and Slavkovik [31]. It is easy to adapt our proof of Theorem 4 to obtain this result: simply observe that, for any model of ROOTED that can possibly maximise the cardinality of the equaliser, there will be full agreement on all propositions of the form $p_{y \succ z}$ with $x \neq y \neq z \neq x$. But, arguably, the simulation in terms of WINNER is the more useful one, at least from an algorithmic point of view, given that the number of models of WINNER is linear while that of ROOTED is exponential in the number m of alternatives.

THEOREM 5. *When restricted to RANKING-rational profiles, $\text{max-set}(\cdot, \text{WINNER})$ simulates the rule returning the **Uncovered Set**.*

PROOF. Define J^x as in the proof of Theorem 3. Consider any profile $J \in \text{Mod}(\text{RANKING})^n$. We obtain:

$$\begin{aligned} \text{max-set}(J, \text{WINNER}) &= \operatorname{argsetmax}_{J \in \text{Mod}(\text{WINNER})} \text{Eq}(J, \text{Maj}(J)) \\ &= \operatorname{argsetmax}_{J^x | x \in X} \text{Eq}(J^x, \text{Maj}(J)) \\ &= \operatorname{argsetmax}_{J^x | x \in X} \{\varphi \in \Phi_{\neq}^X \mid J^x(\varphi) = \text{Maj}(J)(\varphi)\} \end{aligned}$$

Now consider once more the decomposition of Φ_{\neq}^X given in the proof of Theorem 3. First, J^x and $\text{Maj}(J)$ agree on all $p_{y \succ z}$ with $y = z$, so we can disregard propositions of this form when evaluating $\operatorname{argsetmax}$. Second, they agree on both $p_{x \succ y}$ and $p_{y \succ x}$ if $\text{Maj}(J)(p_{x \succ y}) = 1$; otherwise they disagree on both propositions. So we can focus on $p_{x \succ y}$ and ignore $p_{y \succ x}$ when evaluating $\operatorname{argsetmax}$. Third, they agree on $p_{y \succ z}$ with $x \neq y \neq z \neq x$ if and

only if $\text{Maj}(J)(p_{y \succ z}) = 1$, given that J^x accepts all such propositions. Thus, we can simplify and obtain:

$$\begin{aligned} \text{max-set}(J, \text{WINNER}) &= \operatorname{argsetmax}_{x \in X} \{p_{y \succ z} \mid y \neq z \text{ and } \text{Maj}(J)(p_{y \succ z}) = 1 \text{ and } z \neq x\} \\ &= \operatorname{argsetmax}_{x \in X} \{(y, z) \mid y \succ_M z\} \setminus \{(y, x) \mid y \succ_M x\} \\ &= \operatorname{argsetmax}_{x \in X} \{z \in X \mid x \succ_M z\} \end{aligned}$$

The final step above is sanctioned by the fact that $\{(y, x) \mid y \succ_M x\}$ is minimal if and only if $\{(x, z) \mid x \succ_M z\}$ is maximal. Now, given that $D(x) = \{z \in X \mid x \succ_M z\}$, we have in fact obtained the Uncovered Set and are done. \square

To the best of my knowledge, this is the first characterisation of the Uncovered Set in judgment aggregation.

5 CONCLUSION

I have introduced a new model of judgment aggregation that clearly separates the constraints to be satisfied by individuals from those to be satisfied by the collective decision arrived at by these individuals. I have argued that the natural distinction between rationality and feasibility at the conceptual level should be reflected at the technical level as well. The new model allows for this distinction and thus can represent a wider range of application domains in a natural manner. The technical results in this paper support this view. First, comparing the characterisation theorem for the majority rule proved here with known results for the standard model of judgment aggregation shows that there are pairs of integrity constraints that each would permit paradoxical outcomes in the standard model but that will avoid such paradoxes when one is used as the rationality constraint and the other as the feasibility constraint. Second, the results regarding the simulation of voting rules show that the new model greatly simplifies the task of finding counterparts of common voting rules in judgment aggregation.

The new model should be investigated further and this paper opens up several concrete avenues for future work. This includes, in particular, the research agenda concerned with finding natural counterparts of common voting rules in judgment aggregation. To this end, the lexi-max and greedy-max rules should be explored, as should be alternative feasibility constraints, such as a variant of the WINNER -constraint under which incomparability rather than indifference is stipulated amongst nonwinning alternatives.

ACKNOWLEDGMENTS

An early version of this paper was previously circulated under the name ‘‘A Note on Embedding Voting Rules into Judgment Aggregation’’ and presented at the Workshop on Game Theory and Social Choice held in Budapest in 2015. I thank the participants of that workshop as well as Umberto Grandi, Ronald de Haan, Jérôme Lang, Marija Slavkovik, Zoi Terzopoulou, Bill Zwicker, and the anonymous AAMAS reviewers for the helpful feedback received.

REFERENCES

- [1] Dorothea Baumeister, Gábor Erdélyi, and Jörg Rothe. 2016. Judgment Aggregation. In *Economics and Computation*, Jörg Rothe (Ed.). Springer.
- [2] Farah Benamara, Souhila Kaci, and Gabriella Pigozzi. 2010. Individual Opinions-Based Judgment Aggregation Procedures. In *Proceedings of the 7th International Conference on Modeling Decisions for Artificial Intelligence (MDAI-2010)*. Springer.

- [3] Jean-Charles de Borda. 1781. *Mémoire sur les Élections au Scrutin*. Histoire de l'Académie Royale des Sciences, Paris.
- [4] Steven J. Brams and Peter C. Fishburn. 1978. Approval Voting. *The American Political Science Review* 72, 3 (1978), 831–847.
- [5] Felix Brandt, Markus Brill, and Paul Harrenstein. 2016. Tournament Solutions. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press.
- [6] Arthur H. Copeland. 1951. A "Reasonable" Social Welfare Function. Mimeographed Notes from a Seminar on Applications of Mathematics to the Social Sciences. University of Michigan.
- [7] Marco Costantini, Carla Groenland, and Ulle Endriss. 2016. Judgment Aggregation under Issue Dependencies. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI-2016)*.
- [8] Dirk van Dalen. 2013. *Logic and Structure* (5th ed.). Springer.
- [9] Franz Dietrich. 2014. Scoring Rules for Judgment Aggregation. *Social Choice and Welfare* 42, 4 (2014), 873–911.
- [10] Franz Dietrich and Christian List. 2007. Arrow's Theorem in Judgment Aggregation. *Social Choice and Welfare* 29, 1 (2007), 19–33.
- [11] Conal Duddy, Ashley Piggins, and William S. Zwicker. 2016. Aggregation of Binary Evaluations: A Borda-like Approach. *Social Choice and Welfare* 46, 2 (2016), 301–333.
- [12] Edith Elkind and Arkadii Slinko. 2016. Rationalizations of Voting Rules. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press.
- [13] Ulle Endriss. 2016. Judgment Aggregation. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press.
- [14] 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-2015)*. IFAAMAS.
- [15] Ulle Endriss, Umberto Grandi, and Daniele Porello. 2012. Complexity of Judgment Aggregation. *Journal of Artificial Intelligence Research* 45 (2012), 481–514.
- [16] Patricia Everaere, Sébastien Konieczny, and Pierre Marquis. 2014. Counting Votes for Aggregating Judgments. In *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2014)*. IFAAMAS.
- [17] Patricia Everaere, Sébastien Konieczny, and Pierre Marquis. 2017. An Introduction to Belief Merging and its Links with Judgment Aggregation. In *Trends in Computational Social Choice*, Ulle Endriss (Ed.). AI Access.
- [18] Daniel Farkas and Shmuel Nitzan. 1979. The Borda Rule and Pareto Stability: A Comment. *Econometrica* 47, 5 (1979), 1305–1306.
- [19] Felix Fischer, Olivier Hudry, and Rolf Niedermeier. 2016. Weighted Tournament Solutions. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press.
- [20] Peter C. Fishburn. 1977. Condorcet Social Choice Functions. *SIAM J. Appl. Math.* 33, 3 (1977), 469–489.
- [21] Umberto Grandi. 2012. *Binary Aggregation with Integrity Constraints*. Ph.D. Dissertation. ILLC, University of Amsterdam.
- [22] Umberto Grandi and Ulle Endriss. 2013. Lifting Integrity Constraints in Binary Aggregation. *Artificial Intelligence* 199–200 (2013), 45–66.
- [23] Davide Grossi and Gabriella Pigozzi. 2014. *Judgment Aggregation: A Primer*. Morgan & Claypool Publishers.
- [24] Edith Hemaspaandra, Holger Spakowski, and Jörg Vogel. 2005. The Complexity of Kemeny Elections. *Theoretical Computer Science* 349, 3 (2005), 382–391.
- [25] John G. Kemeny. 1959. Mathematics without Numbers. *Daedalus* 88, 4 (1959), 577–591.
- [26] Sébastien Konieczny and Ramón Pino Pérez. 2002. Merging Information under Constraints: A Logical Framework. *Journal of Logic and Computation* 12, 5 (2002), 773–808.
- [27] Lewis A. Kornhauser and Lawrence G. Sager. 1993. The One and the Many: Adjudication in Collegial Courts. *California Law Review* 81, 1 (1993), 1–59.
- [28] Jérôme Lang, Jérôme Monnot, Arkadii Slinko, and William S. Zwicker. 2017. Beyond Electing and Ranking: Collective Dominating Chains, Dominating Subsets and Dichotomies. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2017)*. IFAAMAS.
- [29] Jérôme Lang, Gabriella Pigozzi, Marija Slavkovic, 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-2011)*.
- [30] Jérôme Lang, Gabriella Pigozzi, Marija Slavkovic, Leendert van der Torre, and Srđjan Vesic. 2017. A Partial Taxonomy of Judgment Aggregation Rules and their Properties. *Social Choice and Welfare* 48, 2 (2017), 327–356.
- [31] Jérôme Lang and Marija Slavkovic. 2013. Judgment Aggregation Rules and Voting Rules. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT-2013)*. Springer.
- [32] Christian List. 2012. The Theory of Judgment Aggregation: An Introductory Review. *Synthese* 187, 1 (2012), 179–207.
- [33] Christian List and Philippe Pettit. 2002. Aggregating Sets of Judgments: An Impossibility Result. *Economics and Philosophy* 18, 1 (2002), 89–110.
- [34] Christian List and Clemens Puppe. 2009. Judgment Aggregation: A Survey. In *Handbook of Rational and Social Choice*, Paul Anand, Prasanta K. Pattanaik, and Clemens Puppe (Eds.). Oxford University Press.
- [35] Pierre Marquis. 2000. Consequence Finding Algorithms. In *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, Dov M. Gabbay and Philippe Smets (Eds.). Vol. 5. Springer.
- [36] Michael K. Miller. 2008. Judgment Aggregation and Subjective Decision-Making. *Economics and Philosophy* 24, 2 (2008), 205–231.
- [37] Michael K. Miller and Daniel Osherson. 2009. Methods for Distance-Based Judgment Aggregation. *Social Choice and Welfare* 32, 4 (2009), 575–601.
- [38] Nicholas R. Miller. 1980. A New Solution Set for Tournaments and Majority Voting: Further Graph-theoretical Approaches to the Theory of Voting. *American Journal of Political Science* 24, 1 (1980), 68–96.
- [39] Klaus Nehring and Marcus Pivato. 2013. *Majority Rule in the Absence of a Majority*. MPRA Paper 46721. Munich Personal RePEc Archive.
- [40] Klaus Nehring, Marcus Pivato, and Clemens Puppe. 2014. The Condorcet Set: Majority Voting over Interconnected Propositions. *Journal of Economic Theory* 151 (2014), 268–303.
- [41] Klaus Nehring and Clemens Puppe. 2007. The Structure of Strategy-Proof Social Choice. Part I: General Characterization and Possibility Results on Median Spaces. *Journal of Economic Theory* 135, 1 (2007), 269–305.
- [42] Gabriella Pigozzi. 2006. Belief Merging and the Discursive Dilemma: An Argument-Based Account of Paradoxes of Judgment Aggregation. *Synthese* 152, 2 (2006), 285–298.
- [43] Daniele Porello. 2013. A Proof-Theoretical View of Collective Rationality. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI-2013)*.
- [44] Daniele Porello. 2014. Logics for Collective Reasoning. In *Proceedings of the European Conference on Social Intelligence (ECSI-2014)*. CEUR.
- [45] Daniele Porello and Ulle Endriss. 2014. Ontology Merging as Social Choice: Judgment Aggregation under the Open World Assumption. *Journal of Logic and Computation* 24, 6 (2014), 1229–1249.
- [46] Patrick Slater. 1961. Inconsistencies in a Schedule of Paired Comparisons. *Biometrika* 48, 3–4 (1961), 303–312.
- [47] Marija Slavkovic. 2012. *Judgment Aggregation for Multiagent Systems*. Ph.D. Dissertation. University of Luxembourg.
- [48] T. Nicolaus Tideman. 1987. Independence of Clones as a Criterion for Voting Rules. *Social Choice and Welfare* 4, 3 (1987), 185–206.
- [49] William S. Zwicker. 2016. Introduction to the Theory of Voting. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press.