

Aggregating Incomplete Pairwise Preferences by Weight

Zoi Terzopoulou* , Ulle Endriss

Institute for Logic, Language and Computation, University of Amsterdam
z.terzopoulou@uva.nl, u.endriss@uva.nl

Abstract

We develop a model for the aggregation of preferences that do not need to be either complete or transitive. Our focus is on the normative characterisation of aggregation rules under which each agent has a weight that depends only on the size of her ballot, i.e., on the number of pairs of alternatives for which she chooses to report a relative ranking. We show that for rules that satisfy a restricted form of majoritarianism these weights in fact must be constant, while for rules that are invariant under agents with compatible preferences forming pre-election pacts it must be the case that an agent’s weight is inversely proportional to the size of her ballot.

1 Introduction

Modelling and aggregating the preferences of agents is relevant to various important applications of AI [Domshlak *et al.*, 2011], ranging from online recommendation systems, to decision-theoretic planning engines, to the emerging topic of incorporating ethical principles into computer-supported decision making [Conitzer *et al.*, 2017]. While it is acknowledged that decision makers cannot be expected to be perfectly rational and to always report transitive preferences regarding the full set of alternatives under consideration [Darmann *et al.*, 2017; Hansson and Grüne-Yanoff, 2018], much existing work in computational social choice is still based on the classical model of preferences that presupposes completeness and transitivity [Brandt *et al.*, 2016]. In this paper we propose a model of preference aggregation that does not make such assumptions. Instead, preferences are represented as merely acyclic sets of pairwise comparisons between alternatives.

To illustrate the need for systems that can handle such non-standard preferences, consider the following example.

Example 1. *Suppose you are asked in a poll to rank several apps in order of preference. You strictly prefer the New York Times app to Facebook for getting informed about the news and Facebook to Gmail when it comes to communicating with friends. But NYT and Gmail serve two completely distinct purposes for you, and you simply cannot rank them. The poll also includes apps you have never used, you are not*

aware of their features, and hence do not hold any preference about them. So, you presumably would greatly appreciate the possibility of (i) expressing your preferences in a pairwise manner rather than as a list, and of (ii) leaving some of the apps unranked altogether.

We focus on a family of aggregation rules under which the *weight* assigned to agent *i*’s ranking of two alternatives depends on how many other pairs of alternatives she ranks as well. For example, we may want to give an agent who only reports her preferences on a selected few pairs of alternatives more weight than an agent who expresses a view on all pairs. We consider different variants of these aggregation rules, concerning the *type* of output required: we may use a rule to obtain a consensus regarding the *ranking* of the alternatives, we may look for a single *winner* amongst the alternatives, or we may need to select a *set of winners* of a given size. Our new model is thus very flexible both with respect to the individual preferences and the corresponding collective outcomes.

Our main technical results concern the rules characterised by two normative principles (also known as *axioms*):¹

- *Majoritarianism:* Try to respect the will of the majority regarding the relative ranking of pairs of alternatives. We consider a highly restricted form of this principle that avoids majority cycles [de Condorcet, 1785].
- *Splitting:* If several agents have mutually compatible preferences, it should be possible for them to form a pre-election pact and all report the union of their individual preference sets—without this change affecting the outcome. We consider different variants of this idea.

We concentrate our analysis on (restricted) majoritarianism and splitting because these axioms—contrary to other well-known ones—are specifically relevant to scenarios with incomplete preferences. We find that restricted majoritarianism characterises the rule with *constant weights* and that (certain) splitting axioms characterise the rule where an agent’s weight is inversely proportional to the size of her reported preference set (we call this the *even-and-equal-weight rule*). These results are robust across all types of output requirements. We also simulate prominent rules of the standard voting framework with complete preferences and suggest new ones that

¹Similar axioms have been discussed by Terzopoulou *et al.* [2018] in the context of judgment aggregation.

*Contact Author

(despite being very natural) have not been considered in the literature before. These proposals for new rules for standard voting are highlighted as Remarks 1–4. Finally, we spell out basic complexity results for all the rules introduced.

The dominant approach to handling incomplete preferences in computational social choice so far has been the *possible winner* solution concept [Konczak and Lang, 2005; Boutilier and Rosenschein, 2016]. Given a profile of incomplete preferences, the possible winners of a voting rule F are all the alternatives that win under F for *some* completion of the preferences. This is a reasonable solution concept when in fact the preferences of the agents are complete but we have not yet fully elicited them. Along similar lines, one could also assume that there is an underlying true ranking of the alternatives and that the agents’ incomplete preferences capture their partial knowledge about that ranking [Caragiannis *et al.*, 2017]. Here, instead, we are interested in the aggregation of *bona fide* incomplete preferences, i.e., in scenarios where it is the preferences themselves that are incomplete rather than our information about them. Of course, which approach is appropriate depends on context. Like us, Pini *et al.* [2008] and Endriss *et al.* [2009] have studied the aggregation of *bona fide* incomplete preferences, but their proposals are less radical (they do not dispense with transitivity) and, in some sense, less flexible (they do not have the notion of outcome type, which is prominent in our model). While Pini *et al.* extend a number of classical impossibility theorems in social choice theory to the case of incomplete preferences and while Endriss *et al.* investigate how forcing voters to report preferences that are either more or less complete than their true ones affects strategic behaviour, our focus is on the design and axiomatic characterisation of new aggregation rules.

The main body of this paper is organised as follows. We introduce our model in Section 2, and each of the next three sections is dedicated to results for a specific type of output of aggregation rules: a collective ranking (Section 3), a single winner (Section 4), and a set of winners (Section 5).

2 The Model

In this section we present our model of pairwise-preference aggregation and establish some of its elementary properties.

2.1 Basic Notation and Terminology

The *superpopulation* \mathcal{N} is an unbounded set of all *potential agents* that may participate in a preference aggregation scenario, and \mathcal{A} is an unbounded set of all *potential alternatives* that agents in \mathcal{N} may form preferences on.² Any specific preference aggregation scenario concerns a finite *group of agents* $N = \{1, \dots, n\} \subseteq \mathcal{N}$ and a finite subset $A \subseteq \mathcal{A}$, which we call *the set of alternatives*.

Every agent $i \in N$ performs *pairwise comparisons* over the alternatives in A and submits her *ballot* R_i , which is a set of strictly ordered pairs of alternatives:

$$R_i = \{(a, b) \in A \times A \mid a \text{ is ranked above } b \text{ by agent } i\}$$

²The notions of superpopulation and potential alternatives make explicit what is often left implicit in related models: that the number of agents and alternatives may vary across applications. In the proofs, they allow us to move between different scenarios.

We write ab as an abbreviation for (a, b) and use the terms “ballot” and “(pairwise-)preference (set)” interchangeably. We assume that for every agent i the set R_i is *acyclic*: if $a_1 a_2, a_2 a_3, \dots, a_{k-1} a_k \in R_i$, then $a_k a_1 \notin R_i$, for every $a_1, a_2, \dots, a_k \in A$.³ To talk about the transitive closure of a ballot $\{a_1 a_2, a_2 a_3, \dots, a_{k-1} a_k\}$, we sometimes write $a_1 \succ a_2 \succ \dots \succ a_k$. Lastly, let us denote by $\mathcal{R}(A)$ the set of all acyclic pairwise preferences over a set of alternatives A .

A *profile* $\mathbf{R} = (R_1, \dots, R_n) \in \mathcal{R}(A)^n$ captures the ballots of all agents in N . $N_{ab}^{\mathbf{R}}$ denotes the set of agents i for which $ab \in R_i$ in \mathbf{R} . A *preference aggregation rule* F is a function that maps any profile $\mathbf{R} \in \mathcal{R}(A)^n$ for any set of alternatives A and group N to a nonempty set of preferences on the same alternatives, i.e., to a nonempty subset of $2^{\mathcal{R}(A)}$. Thus, the set $F(\mathbf{R})$ may contain several, tied, outcomes.

2.2 Weight Rules and Types of Preferences

In this paper we introduce a specific family of aggregation rules. The primary constituent of a *weight rule* for aggregating pairwise preferences is a *weight function* $w : \mathbb{N} \rightarrow \mathbb{R}^+$. A weight function assigns to each natural number, representing a ballot size $|R_i|$, a positive real number $w_{|R_i|} = w(|R_i|)$, expressing how much any ordered pair $ab \in R_i$ weighs, given the size of R_i . We may also think of w as an infinite vector. For example, $w = (1, \frac{1}{2}, \frac{1}{3}, \dots)$ means that we distribute a total weight of 1 equally across all pairs in R_i , while $w = (1, \frac{1}{M}, \frac{1}{M^2}, \dots)$, for a sufficiently large number M , weighs smaller ballots “infinitely” more than larger ones.

The concept of a weight function is essential for a number of applications. For example, in the context of peer-grading, suppose one grader submits a ranking of only three out of the twenty submissions, while another grader ranks all twenty of them. Should the same weight be assigned to those two rankings? Should the small ranking weigh less, or more? By the end of this paper we will be able answer these questions.

We shall assume that any weight w_λ for $\lambda \in \mathbb{N}$ is computable in polynomial time (polynomial in the length of the description of λ).⁴ We use the following notation to refer to the cumulative weight assigned to all the pairs in a preference set $R \subseteq A \times A$ from the perspective of holding R_i :

$$w_{R_i}(R) = \sum_{ab \in R} w_{|R_i|} \cdot \mathbb{1}_{ab \in R_i} = w_{|R_i|} \cdot |R \cap R_i|$$

A second component of the definition of a weight rule is a *type* \mathcal{T} , which is a function that maps every set of alternatives $A \subseteq \mathcal{A}$ to a subset $\mathcal{T}(A) \subseteq \mathcal{R}(A)$ of preference sets. We focus on four specific types of immediate practical relevance. First, in the literature on preference aggregation both individual preferences and outcomes traditionally are *complete rankings* over the alternatives. So we consider the type of *linear orders* \mathcal{L} , where $\mathcal{L}(A)$ contains all asymmetric, complete, and transitive relations over A . On the other hand, in voting theory we are mainly concerned with finding

³Acyclicity implies *asymmetry* (i.e., $ab \in R_i \Rightarrow ba \notin R_i$ for every $a, b \in A$). However, we stress that we do not assume that R_i needs to be transitive or complete; that is, there may exist alternatives $a, b, c \in A$ such that $ab, bc \in R_i$, but $ac \notin R_i$.

⁴This is important for the (upper bounds in) complexity results.

a (single) *winner*, that is, an alternative ranked higher than any other alternative [Zwicker, 2016]. Hence, we also define the type \mathcal{W} , where $\mathcal{W}(A)$ includes all pairwise preferences of the form $R = \{aa' \mid a' \in A \setminus \{a\}\}$, for some $a \in A$. In multiwinner voting the outcome should be a given number of winning alternatives, constituting, for instance, a *committee* [Faliszewski *et al.*, 2017]. So we define the type \mathcal{W}_k , where $\mathcal{W}_k(A)$ consists of all pairwise preferences that designate a set of k winners (for some fixed $k > 1$), i.e., that are of the form $R = \{(a_1, a'), \dots, (a_k, a') \mid a' \in A \setminus \{a_1, \dots, a_k\}\}$, for some $a_1, \dots, a_k \in A$ with $a_i \neq a_j$ for $i \neq j$. Finally, we define the type \mathcal{C} of *chains*, which are linear orders over subsets of A , to model situations where agents rank part of the alternatives. Formally, $\mathcal{C}(A)$ consists of all pairwise preferences of the form $R = (a_1 \succ a_2 \succ \dots \succ a_k)$, for some $k \in \mathbb{N}$ and $a_1, \dots, a_k \in A$ with $a_i \neq a_j$ for $i \neq j$.⁵

In line with our existing notation $\mathcal{R}(A)$, we use the letter \mathcal{R} to denote the type of *all* preference sets.

Given a weight function w and a profile $\mathbf{R} = (R_1, \dots, R_n)$ for a specific set of alternatives $A \subseteq \mathcal{A}$ and group of agents $N \subseteq \mathcal{N}$, a *weight rule of type \mathcal{T}* , denoted by $F_w^{\mathcal{T}}$, decides the outcome by selecting those preference sets in $\mathcal{T}(A)$ that maximise the total weight across all agents:

$$F_w^{\mathcal{T}}(\mathbf{R}) = \operatorname{argmax}_{R \in \mathcal{T}(A)} \sum_{i \in N} w_{R_i}(R)$$

Note that for every weight rule $F_w^{\mathcal{T}}$ there are infinitely many weight functions that induce it. For example, if we multiply all weights with some positive constant, then the rule being induced does not change. The following two weight rules will play a major role in this paper:

- Take any weight function w with $w_\lambda = w_{\lambda'}$ for all $\lambda, \lambda' \in \mathbb{N}$ and any type \mathcal{T} . We call the weight rule $F_c^{\mathcal{T}}$ induced by w the *constant-weight rule* of type \mathcal{T} .
- Take any weight function w with $w_\lambda = \frac{w_1}{\lambda}$ for all $\lambda \in \mathbb{N}$ and any type \mathcal{T} . We call the weight rule $F_{ev}^{\mathcal{T}}$ induced by w the *equal-and-even-weight rule* of type \mathcal{T} .

The latter rule borrows its name from its counterpart for voting with approval ballots, the *equal-and-even cumulative voting rule* [Glasser, 1959; Alcalde-Unzu and Vorsatz, 2009].

2.3 Type Restrictions

The *restriction* of a set of preferences $S \subseteq \mathcal{R}(A)$ to a given type \mathcal{T} is formally defined as $S|_{\mathcal{T}} = S \cap \mathcal{T}(A)$.

Recall that, to compute $F_w^{\mathcal{T}}(\mathbf{R})$, we have to go through all preferences R of type \mathcal{T} and select those that maximise total weight. An alternative definition would be to first select all preferences R of *any* type that maximise total weight and to only then restrict the resulting set to \mathcal{T} , i.e., to compute $(\operatorname{argmax}_{R \in \mathcal{R}(A)} \sum_{i \in N} w_{R_i}(R))|_{\mathcal{T}}$, which we can also write as $F_w^{\mathcal{R}}(\mathbf{R})|_{\mathcal{T}}$. Now the question arises: When do these two alternative approaches result in the same outcomes?

⁵Lang *et al.* [2017] define certain aggregation outcomes that include ours. For instance, their *plain dominating 1-chain* and *plain dominating k-subset* correspond to our single-winner and k -winner sets, respectively. But the model and the results they provide are restricted to complete input profiles.

To answer this question, we require some further terminology. Given two types \mathcal{T} and \mathcal{T}' we say that \mathcal{T} *refines* \mathcal{T}' if $\mathcal{T}(A) \subseteq \mathcal{T}'(A)$ for all $A \subseteq \mathcal{A}$. We furthermore say that \mathcal{T} *extends* \mathcal{T}' if for all $A \subseteq \mathcal{A}$ and preferences $R' \in \mathcal{T}'(A)$ there exists a preference $R \in \mathcal{T}(A)$ with $R \supseteq R'$.

Proposition 1. *If \mathcal{T} simultaneously refines and extends \mathcal{T}' , then $F_w^{\mathcal{T}}(\mathbf{R}) = F_w^{\mathcal{T}'}(\mathbf{R})|_{\mathcal{T}}$ for any w and \mathbf{R} .*

Proof. Since $\mathcal{T}(A) \subseteq \mathcal{T}'(A)$, which holds as \mathcal{T} refines \mathcal{T}' , we have that $F_w^{\mathcal{T}}(\mathbf{R}) = \operatorname{argmax}_{R \in \mathcal{T}(A)} \sum_{i \in N} w_{R_i}(R) \supseteq \mathcal{T}(A) \cap \operatorname{argmax}_{R \in \mathcal{T}'(A)} \sum_{i \in N} w_{R_i}(R) = F_w^{\mathcal{T}'}(\mathbf{R})|_{\mathcal{T}}$.

For the other direction, we first show that there exists at least one $R \in \mathcal{T}(A)$ with $R \in F_w^{\mathcal{T}'}(\mathbf{R})$. Take any $R' \in F_w^{\mathcal{T}'}(\mathbf{R})$. Since \mathcal{T} extends \mathcal{T}' , there exists an $R \in \mathcal{T}(A)$ with $R \supseteq R'$. Observe that $w_{|R_i|} \cdot |R_i \cap R| \geq w_{|R_i|} \cdot |R_i \cap R'|$ for all $i \in N$ and thus $\sum_{i \in N} w_{R_i}(R) \geq \sum_{i \in N} w_{R_i}(R')$. Hence, we indeed get $R \in F_w^{\mathcal{T}'}(\mathbf{R})$ as claimed.

Now take any $R^* \in F_w^{\mathcal{T}}(\mathbf{R})$. By definition of the weight rule, this entails $\sum_{i \in N} w_{R_i}(R^*) \geq \sum_{i \in N} w_{R_i}(R)$ for all $R \in \mathcal{T}(A)$, including the specific R we know to exist in $F_w^{\mathcal{T}'}(\mathbf{R})$. As furthermore $R^* \in \mathcal{T}'(A)$, which is due to $\mathcal{T}(A) \subseteq \mathcal{T}'(A)$, we get $R^* \in F_w^{\mathcal{T}'}(\mathbf{R})$ and thus also $R^* \in F_w^{\mathcal{T}'}(\mathbf{R})|_{\mathcal{T}}$. We conclude that $F_w^{\mathcal{T}}(\mathbf{R}) \subseteq F_w^{\mathcal{T}'}(\mathbf{R})|_{\mathcal{T}}$. \square

Every type \mathcal{T} refines \mathcal{R} . But which types \mathcal{T} also extend \mathcal{R} ? Certainly \mathcal{L} extends \mathcal{R} , as every pairwise-preference set can be extended to a linear order, and therefore so does every type \mathcal{T} that includes all linear orders, answering our question.

Corollary 2. *For any weight function w and type \mathcal{T} that is refined by \mathcal{L} , it holds that $F_w^{\mathcal{T}}(\mathbf{R}) = F_w^{\mathcal{R}}(\mathbf{R})|_{\mathcal{T}}$.*

3 Obtaining a Collective Ranking

In this section we study the family of weight rules for outcomes that are complete rankings, i.e., outcomes of type \mathcal{L} .

Agents may still report incomplete preferences. But for the special case of complete profiles $\mathbf{L} \in \mathcal{L}(A)^n$, all ballots are associated with the same weight (since their size is the same and their weight only depends on their size) and the choice of the weight function w is irrelevant. Hence, the induced weight rule for outcomes in $\mathcal{L}(A)$ for this particular case coincides with the well-known rule of Kemeny [1959].^{6,7}

Proposition 3. *For any profile $\mathbf{L} \in \mathcal{L}(A)^n$ and weight function w , it holds that $F_w^{\mathcal{L}}(\mathbf{L}) = \text{Kemeny}(\mathbf{L})$.*

We now state the *relative ranking problem* for a weight rule $F_w^{\mathcal{L}}$: “Given a profile $\mathbf{R} \in \mathcal{R}(A)^n$ and two distinguished alternatives $a, b \in A$, is a ranked above b in at least one ranking in $F_w^{\mathcal{L}}(\mathbf{R})$?” Proposition 4 clarifies that determining relative rankings under any weight rule of type \mathcal{L} is complete

⁶In the interest of space, we omit a formal definition of this rule and refer to Brandt *et al.* [2016] for a modern exposition.

⁷It might be tempting to confuse the possible outcomes (as defined by Konczak and Lang 2005) of the Kemeny rule with the outcomes induced by the constant-weight rule $F_c^{\mathcal{L}}$. But in different aggregation situations possible outcomes may coincide with, include, or be included in the set of outcomes of $F_c^{\mathcal{L}}$. For an example of the latter case, consider the profile $(\{ab, bc, ac\}, \{ab, ca\}, \{ab, ca\})$.

for $P_{||}^{\text{NP}}$, i.e., for the class of problems solvable in polynomial time via parallel access to an NP oracle. We note that computing the outcome under such a rule is at least as hard.

Proposition 4. *The relative ranking determination problem for $F_w^{\mathcal{L}}$ is $P_{||}^{\text{NP}}$ -complete, for any weight function w .*

Proof sketch. By Proposition 3, $P_{||}^{\text{NP}}$ -hardness follows from the corresponding result for the Kemeny rule [Hemaspaandra *et al.*, 2005]. Proving membership in $P_{||}^{\text{NP}}$ is routine. \square

Next, we wonder: what weights induce “good” rules? To answer, we adopt an axiomatic perspective [Arrow *et al.*, 2002; Brandt *et al.*, 2016]. For the sake of our definitions, we fix a set of alternatives $A \subseteq \mathcal{A}$, a group of agents $N \subseteq \mathcal{N}$, and an aggregation rule F of type \mathcal{L} , and say that F satisfies an axiom if the relevant definition holds for every A and N .

3.1 Majoritarianism

Generating a collective ranking based on the pairwise preferences of the majority is commonly considered a desirable attribute of an aggregation rule. However, blindly following the majority’s preferences easily leads to problematic cases of preference cycles [de Condorcet, 1785]. But there clearly are situations where some pairs of alternatives cannot cause a Condorcet-type paradox on a profile. Such is the case when the alternatives a, b are *independent* of a profile \mathbf{R} , i.e., when for each ballot R_i in \mathbf{R} , a and b might be compared to each other, but neither a nor b are compared to any other c (formally, $ac, ca, bc, cb \notin R_i$ for all $c \in A \setminus \{a, b\}$ and $i \in N$).

The axiom of *restricted majoritarianism* expresses a fundamental normative principle (namely that more agents ranking a above b than b above a should imply collective preference of a to b), but limits its scope to cases of independence. As its name suggests, restricted majoritarianism is indeed very restrictive and this makes our findings all the more remarkable: Even by using only this weak (and hence indisputable) axiom, we get a characterisation of a very natural rule (Theorem 5).

For a profile $\mathbf{R} \in \mathcal{R}(A)^n$, let us define the *simple-majority set* as $m(\mathbf{R}) = \{(a, b) \in A \times A \mid |N_{ab}^{\mathbf{R}}| > |N_{ba}^{\mathbf{R}}|\}$.

Axiom 1 (Restricted majoritarianism). *A rule F is said to satisfy restricted majoritarianism if for all alternatives $a, b \in A$ independent of a profile $\mathbf{R} \in \mathcal{R}(A)^n$, it holds that $ab \in m(\mathbf{R})$ implies that $ab \in R$ for all $R \in F(\mathbf{R})$.*

Theorem 5. *The only weight rule of type \mathcal{L} satisfying restricted majoritarianism is the constant-weight rule $F_c^{\mathcal{L}}$.*

Proof. That the constant-weight rule $F_c^{\mathcal{L}}$ satisfies restricted majoritarianism is easy to verify. For the other direction, consider an arbitrary weight function w that induces a weight rule $F_w^{\mathcal{L}}$ and suppose that $F_w^{\mathcal{L}}$ satisfies restricted majoritarianism. To prove that $F_w^{\mathcal{L}}$ is the constant weight rule it suffices to show that $w_\lambda = w_{\lambda'}$ for all $\lambda, \lambda' \in \mathbb{N}$.

We first show that $F_w^{\mathcal{L}}$ satisfying restricted majoritarianism for a group $N \subseteq \mathcal{N}$ with odd $n > 1$ implies that $\frac{w_\lambda}{w_{\lambda'}} < 1 + \frac{2}{n-1}$ for all λ, λ' . Indeed, consider two arbitrary $\lambda > \lambda'$ and a set of alternatives $A = \{a_1, b_1, a'_1, b'_1, \dots, a_\lambda, b_\lambda, a'_\lambda, b'_\lambda\} \subseteq \mathcal{A}$ (which is possible because \mathcal{A} is unbounded). Then, take $R = \{a_1 b_1, a_2 b_2, \dots, a_\lambda b_\lambda\} \in \mathcal{R}(A)$ and $R' =$

$\{b_1 a_1, a'_2 b'_2, \dots, a'_\lambda b'_\lambda\} \in \mathcal{R}(A)$. We construct the following profile $\mathbf{R} \in \mathcal{R}(A)^n$ for an odd $n > 1$:

$$\mathbf{R} = \underbrace{(R, \dots, R)}_{(n-1)/2}, \underbrace{(R', \dots, R')}_{(n+1)/2}$$

Clearly, a_1, b_1 are independent of \mathbf{R} , and $b_1 a_1 \in m(\mathbf{R})$. Hence, by restricted majoritarianism it must be the case that $b_1 a_1 \in R''$ for all $R'' \in F_w^{\mathcal{L}}(\mathbf{R})$, implying that $\frac{n-1}{2} w_{|R|} < \frac{n+1}{2} w_{|R'|}$ and thus $\frac{w_\lambda}{w_{\lambda'}} < 1 + \frac{2}{n-1}$.

To conclude, we know that if $F_w^{\mathcal{L}}$ satisfies restricted majoritarianism, then it does so for every group $N \subseteq \mathcal{N}$ of odd cardinality. Thus, we must have that $\frac{w_\lambda}{w_{\lambda'}} < 1 + \frac{2}{n-1}$ for every odd $n > 1$ and for every $w_\lambda, w_{\lambda'} \in \mathbb{R}^+$. Letting n go to infinity, this implies that $w_\lambda = w_{\lambda'}$ for all λ, λ' . \square

3.2 Splitting

We now define a novel axiom tailored to the aggregation of *bona fide* incomplete preferences. The basic idea is that when several agents have compatible and disjoint preference sets, they should be able to simply all report the union of their individual sets without affecting the outcome. This is useful in practice: If some members of a council meet before a vote and realise that they do not disagree on any of the relative rankings between alternatives, they could simply send a messenger to the main meeting to vote on their behalf. Alternatively, in a strategic setting, we may wish to disincentivise pre-election pacts between agents who care about different issues—*vote trading* [Eldar, 2008] is a special case of this.

Axiom 2 (Splitting). *A rule F is said to satisfy (arbitrary) splitting if for every profile $\mathbf{R} \in \mathcal{R}(A)^n$ and subgroup $N' \subseteq N$ of agents with pairwise disjoint ballots, it is the case that $\bigcup_{i \in N'} R_i \in \mathcal{R}(A)$ implies $F(\mathbf{R}) = F(\mathbf{R}')$, where \mathbf{R}' arises from \mathbf{R} by replacing the ballot of each member of N' by the union $\bigcup_{i \in N'} R_i$. F furthermore is said to satisfy equal splitting if the ballots of all agents in N' are of equal size, and single splitting if they are all singletons.*

We find that the single-splitting axiom characterises the equal-and-even-weight rule, and so does the equal-splitting axiom. But interestingly, the arbitrary-splitting axiom proves to be too demanding: it is not satisfied by any weight rule.

Theorem 6. *The only weight rule of type \mathcal{L} satisfying the single-splitting axiom is the equal-and-even-weight rule $F_{ee}^{\mathcal{L}}$.*

Proof. We consider an arbitrary instance of the single-splitting axiom and show that $F_{ee}^{\mathcal{L}}$ satisfies it. Take two profiles \mathbf{R}, \mathbf{R}' , where \mathbf{R}' arises from \mathbf{R} by replacing the preference set of each member of N' by the union $\bigcup_{i \in N'} R_i$. By the definition of the weights for the rule $F_{ee}^{\mathcal{L}}$, we know that every pair of alternatives in $\bigcup_{i \in N'} R_i$ weighs exactly the same in \mathbf{R} and in \mathbf{R}' , because $|R_i| = 1$ for every $i \in N'$, and $w_\lambda = \frac{w_1}{\lambda}$ for every λ . The same holds for all pairs that are not in $\bigcup_{i \in N'} R_i$ too, since they trivially appear in exactly the same preference sets in both profiles \mathbf{R} and \mathbf{R}' . This means that $F_{ee}^{\mathcal{L}}(\mathbf{R}) = F_{ee}^{\mathcal{L}}(\mathbf{R}')$.

For the other direction, consider an arbitrary $\lambda \in \mathbb{N}$, a set of alternatives $A = \{a_1, b_1, \dots, a_\lambda, b_\lambda\} \subseteq \mathcal{A}$, and

two profiles $\mathbf{R} = (\{a_1b_1\}, \{b_1a_1\}, \{a_2b_2\}, \dots, \{a_\lambda b_\lambda\}) \in \mathcal{R}(A)^n$ and $\mathbf{R}' = (\{a_1b_1\}, \{b_1a_1, a_2b_2, \dots, a_\lambda b_\lambda\}, \dots, \{b_1a_1, a_2b_2, \dots, a_\lambda b_\lambda\}) \in \mathcal{R}(A)^n$ (this can be done for a sufficiently large group $N \subseteq \mathcal{N}$). First, for any weight rule $F_w^{\mathcal{L}}$ we must have $L, L' \in F_w^{\mathcal{L}}(\mathbf{R})$ for some $\{a_1b_1, a_2b_2, \dots, a_\lambda b_\lambda\} \subseteq L$ and $\{b_1a_1, a_2b_2, \dots, a_\lambda b_\lambda\} \subseteq L'$. Moreover, if $F_w^{\mathcal{L}}$ satisfies single splitting, it must be the case that $F_w^{\mathcal{L}}(\mathbf{R}) = F_w^{\mathcal{L}}(\mathbf{R}')$. So the weights of a_1b_1 and b_1a_1 should be the same in the profile \mathbf{R}' , that is, $w_1 = \lambda \cdot w_\lambda$. \square

Theorem 7. *The only weight rule of type \mathcal{L} satisfying the equal-splitting axiom is the equal-and-even-weight rule $F_{ee}^{\mathcal{L}}$.*

Proof. If a weight rule satisfies equal splitting, then it satisfies single splitting too, so Theorem 6 implies that it must be the equal-and-even-scoring rule. The proof of the other direction, showing that $F_{ee}^{\mathcal{L}}$ satisfies equal splitting, proceeds exactly as the first part of the proof of Theorem 6. \square

Proposition 8. *No weight rule of type \mathcal{L} satisfies the arbitrary-splitting axiom.*

Proof. Every weight rule of type \mathcal{L} that satisfies arbitrary splitting also satisfies single splitting, so by Theorem 6, the only candidate rule is $F_{ee}^{\mathcal{L}}$. Then it remains to show that $F_{ee}^{\mathcal{L}}$ fails arbitrary splitting on at least one counterexample. Let $A = \{a_1, b_1, a_2, b_2, a_3, b_3\}$ and consider the profiles $\mathbf{R} = (\{b_1a_1\}, \{a_1b_1\}, \{a_2b_2, a_3b_3\})$ and $\mathbf{R}' = (\{b_1a_1\}, \{a_1b_1, a_2b_2, a_3b_3\}, \{a_1b_1, a_2b_2, a_3b_3\})$. Let w be the weight function associated with $F_{ee}^{\mathcal{L}}$. Since both a_1b_1 and b_1a_1 weigh w_1 in profile \mathbf{R} , there will exist two collective preference sets in $F_{ee}^{\mathcal{L}}(\mathbf{R})$ that each contain one of them. Now, note that $w_{\{a_1b_1, a_2b_2, a_3b_3\}} = \frac{w_1}{3}$. So, a_1b_1 weighs $\frac{2w_1}{3} < w_1$ in \mathbf{R}' , while b_1a_1 still weighs w_1 . This means that for every preference set in $F_{ee}^{\mathcal{L}}(\mathbf{R}')$, only b_1a_1 but not a_1b_1 will belong to that set. Hence, $F_{ee}^{\mathcal{L}}(\mathbf{R}) \neq F_{ee}^{\mathcal{L}}(\mathbf{R}')$, which violates the arbitrary-splitting axiom. \square

4 Obtaining a Winner

In this section we study weight rules for outcomes of type \mathcal{W} , indicating a single winner. To refer to a winning alternative of a preference set R , we define the set of the *top alternatives* in R : $\text{top}(R) = \{a \in A \mid (b, a) \in R \text{ for no } b \in A\}$. For any set $S \subseteq \mathcal{R}(A)$, we define $\text{top}(S) = \bigcup_{R \in S} \text{top}(R)$.

For complete input profiles, the winners of any weight rule are exactly the winners according to the rule of Borda (1784): Given a linear order $a_1 \succ a_2 \succ \dots \succ a_m$, an alternative gets $m - 1$ points for the first place, $m - 2$ points for the second place, and so forth. A winning alternative of Borda maximises the added score from all agents. Proposition 9 follows from Observation 4 of Lang *et al.* [2017].⁸

Proposition 9. *For any profile $\mathbf{L} \in \mathcal{L}(A)^n$ and weight function w , it holds that $\text{top}(F_w^{\mathcal{W}}(\mathbf{L})) = \text{Borda}(\mathbf{L})$.*

Moreover, for any weight rule of type \mathcal{W} , the outcome can be computed in polynomial time, generalising the well-known

⁸A similar result is also given by Endriss [2018], who simulates voting in the framework of judgment aggregation.

fact that the Borda winner can be found in polynomial time (the simple proof is omitted due to space constraints).⁹

Proposition 10. *For any profile $\mathbf{R} \in \mathcal{R}(A)^n$ and weight function w , $F_w^{\mathcal{W}}(\mathbf{R})$ can be computed in polynomial time.*

We next wonder what kind of weights make sense when the outcome is of type \mathcal{W} but some of the agents submit strictly incomplete preferences (as opposed to the complete input profiles we just examined).

4.1 Majoritarianism

We observe that the axiom of restricted majoritarianism, as formulated earlier, does not apply to outcomes of type \mathcal{W} , because these are incomplete. We therefore slightly modify it (and for complete outcomes the two versions coincide).

Axiom 3 (Weak restricted majoritarianism). *A rule F satisfies weak restricted majoritarianism if for every profile $\mathbf{R} \in \mathcal{R}(A)^n$ and alternatives $a, b \in A$ that are independent of \mathbf{R} , it holds that: if $ab \in m(\mathbf{R})$, then $ba \notin R$ for all $R \in F(\mathbf{R})$.*

We obtain a similar result as for rules of type \mathcal{L} (Theorem 5).

Theorem 11. *The only weight rule of type \mathcal{W} satisfying weak restricted majoritarianism is the constant-weight rule $F_c^{\mathcal{W}}$.*

Proof. The proof matches that of Theorem 5, so we only explain here the part that differs, after the construction of the profile \mathbf{R} . Let R^a denote the preference set in $\mathcal{W}(A)$ corresponding to a being the winning alternative. We have $\sum_{i \in N} w_{R_i}(R^{a_1}) = \sum_{i \in N} w_{R_i}(R^{a_2}) = \dots = \sum_{i \in N} w_{R_i}(R^{a_\lambda}) = \frac{n-1}{2}w_\lambda$ and $\sum_{i \in N} w_{R_i}(R^{b_1}) = \sum_{i \in N} w_{R_i}(R^{a_2}) = \dots = \sum_{i \in N} w_{R_i}(R^{a_{\lambda'}}) = \frac{n+1}{2}w_{\lambda'}$. Thus, either each of $a_1, a_2, \dots, a_\lambda$ win in some profile or none of them do, and the same is true for $b_1, a_2', \dots, a_{\lambda}'$. This means that it cannot be the case that both $a_1b_1 \notin R''$ and $b_1a_1 \notin R''$ for all $R'' \in F_w^{\mathcal{W}}(\mathbf{R})$, because then $\emptyset \in F_w^{\mathcal{W}}(\mathbf{R})$, which would be contrary to the definition of a weight rule. But a_1, b_1 are independent of \mathbf{R} , and $b_1a_1 \in m(\mathbf{R})$. Hence, by weak restricted majoritarianism, it must hold that $b_1a_1 \in R''$ for some $R'' \in F_w^{\mathcal{W}}(\mathbf{R})$. The remainder of the proof proceeds exactly as for Theorem 5. \square

Remark 1. *To demonstrate how the constant-weight rule of winner type works, we inspect it on profiles of chains (i.e., on profiles $\mathbf{C} \in \mathcal{C}(A)^n$), by first defining a variant of the Borda rule: For a chain $a_1 \succ a_2 \succ \dots \succ a_k$, a_1 gets $k-1$ points, a_2 gets $k-2$ points, and so forth, with a_k and all alternatives not appearing in the chain getting 0 points. Short-sighted Borda (sBorda) then maximises the collected points in a profile.¹⁰*

⁹So, there is a clear difference in complexity between using weight rules and computing possible winners: $F_w^{\mathcal{W}}$, which generalises the Borda rule, can be evaluated in polynomial time—yet, the possible-winner determination problem for Borda is NP-complete [Xia and Conitzer, 2008].

¹⁰Short-sighted Borda differs from *optimistic* (or *modified*) and *pessimistic* Borda, which implicitly take alternatives not included in a chain to be less preferred than those included [Baumeister *et al.*, 2012]. Optimistic Borda (used also by Caragiannis *et al.* 2015) assigns k points to the most preferred alternative in a chain with k elements, $k-1$ points to the second most preferred, and so on, while

Proposition 12. For any profile $C \in \mathcal{C}(A)^n$, it holds that $\text{top}(F_c^{\mathcal{W}}(C)) = \text{sBorda}(C)$.

Proof sketch. For simplicity take $w_\lambda = 1$ for all $\lambda \in \mathbb{N}$. Then, for an outcome R^a corresponding to alternative a winning, we have $w_{C_i}(R^a) = |C_i \cap R^a| = |\{b \in A \mid ab \in C_i\}|$, which means that R^a gains as many points from the ballot of agent i as there are alternatives that i explicitly ranks below a . This is precisely how the sBorda score of a is defined. \square

4.2 Splitting

The splitting axioms for rules of type \mathcal{W} can be (and thus are) defined and operate exactly as for rules of type \mathcal{L} . The proof of our next result is similar to the proofs of Theorems 6 and 7.

Theorem 13. The only weight rule of type \mathcal{W} satisfying the single-splitting (or the equal-splitting) axiom is the equal-and-even-weight rule $F_{ee}^{\mathcal{W}}$.

Remark 2. When applied to profiles of chains, $F_{ee}^{\mathcal{W}}$ is reminiscent of a positional scoring rule, suggesting the definition of a new family of voting rules (the exploration of which is beyond the scope of this paper) that integrates positional scoring rules [Zwicker, 2016] and cumulative voting rules [Glasser, 1959]: Consider the rule that associates a chain $a_1 \succ a_2 \succ \dots \succ a_{k-1} \succ a_k$ of length k with the scoring vector $(\frac{k-1}{k(k-1)/2}, \frac{k-2}{k(k-1)/2}, \dots, \frac{1}{k(k-1)/2}, 0)$. Under this rule, every agent distributes a total weight of 1 across the alternatives in her chain in a Borda-like fashion.

5 Obtaining a Set of Winners

This section is concerned with weight rules generating outcomes of type \mathcal{W}_k , which designate k -sized sets of winners.

Whenever the input profile consists of complete preferences only, any weight rule simulates the k -Borda rule, returning the k alternatives with the highest Borda scores [Faliszewski *et al.*, 2017]. Proposition 14 can be seen as a corollary of Observation 5 of Lang *et al.* [2017].

Proposition 14. For any profile $\mathbf{L} \in \mathcal{L}(A)^n$ and weight function w , it holds that $\text{top}(F_w^{\mathcal{W}_k}(\mathbf{L})) = k\text{-Borda}(\mathbf{L})$.

In analogy to Proposition 10, for any weight rule of type \mathcal{W}_k , the set of winners is easy to compute.

Proposition 15. For any profile $\mathbf{R} \in \mathcal{R}(A)^n$ and weight function w , $F_w^{\mathcal{W}_k}(\mathbf{R})$ can be computed in polynomial time.

Moving on to the investigation of multiwinner weight rules for incomplete ballots, we note that the definitions of weak restricted majoritarianism and splitting apply and the proofs of Theorems 11 and 13 carry over to the case of k winners.

Theorem 16. The only weight rule of type \mathcal{W}_k satisfying weak restr. majoritarianism is the constant-weight rule $F_c^{\mathcal{W}_k}$.

Theorem 17. The only weight rule of type \mathcal{W}_k satisfying the single-splitting (or the equal-splitting) axiom is the equal-and-even-weight rule $F_{ee}^{\mathcal{W}_k}$.

the least preferred alternative in the chain gets 1 point and all others get 0 points. For an example where the two rules give different outcomes, consider the profile $(a \succ b \succ c, c \succ d)$. However, sBorda coincides with the rule assigning as score to an alternative a based on the length of the *longest path* below a [Endriss *et al.*, 2009].

Remark 3. For profiles of chains, the constant-weight k -winner rule is not k -sBorda,¹¹ but it gives rise to a neat original definition of a multiwinner rule. Conceptually, when deciding for a set of k winners, one could regard the potential outcome (for instance, the committee) as a whole and compare the power of its members (based on the voters' explicitly declared preferences) against its nonmembers: Consider the multiwinner voting rule that makes a set $S \subseteq A$ of fixed size win when the alternatives in S are preferred over the alternatives outside S the maximum number of times.

Remark 4. The equal-and-even-weight k -winner rule closely follows the constant-weight rule of the same type; however, in this case, the power of a committee member over an alternative outside the committee hinges also on the sizes of the ballots that rank the member above the nonmember.

Overall, our results in Sections 3, 4, and 5 suggest that majoritarianism and splitting can never be satisfied together, by any weight rule; we thus formulate a principal impossibility:

Corollary 18. For any type $\mathcal{T} \in \{\mathcal{L}, \mathcal{W}, \mathcal{W}_k\}$, there is no weight rule $F_w^{\mathcal{T}}$ that satisfies weak restricted majoritarianism together with single, equal, or arbitrary splitting.

Corollary 18 sheds light on a fundamental theoretical fact that also has significant practical implications. For instance, we now know that for an online recommendation system based on weight rules it is impossible to respect the opinion of the majority of the recommenders and at the same time disincentivise them to manipulate the outcome in groups.

6 Conclusion

We have introduced a model for the representation and aggregation of incomplete preferences, designing a family of rules based on the weight a pairwise ranking of two alternatives receives when reported within a preference set. Our model is relevant to multiagent scenarios where agents may satisfy minimal rationality requirements and violate not only completeness but also transitivity. We have examined two original normative principles—restricted majoritarianism and splitting—that tie in directly with the incompleteness of preferences and characterise two weight rules with particularly simple weight functions. These results are robust across aggregation scenarios with varying output types (full rankings, single winning alternatives, and sets of winners). Besides providing a basic complexity analysis, we have also shown that the rules thus characterised reduce to well-known rules for the special case of complete input profiles and suggest the definition of new ones in the standard voting framework.

Our work opens up several opportunities for future research: not only to identify further axioms and the weight rules they characterise, but also to explore our model of aggregation rules for incomplete pairwise preference sets more broadly, beyond the class of weight rules we have focussed on in this first paper on the topic. Finally, an experimental study to compare different aggregation rules in addition to our theoretical analysis could also be an interesting next step.

¹¹For a counterexample, consider a profile with three ballots that are chains: $a \succ b \succ c \succ d$, $c \succ d \succ a$, and $d \succ a$.

References

- [Alcalde-Unzu and Vorsatz, 2009] Jorge Alcalde-Unzu and Marc Vorsatz. Size approval voting. *Journal of Economic Theory*, 144(3):1187–1210, 2009.
- [Arrow *et al.*, 2002] Kenneth Arrow, Amartya Sen, and Kōtarō Suzumura, editors. *Handbook of Social Choice and Welfare*. North-Holland, 2002.
- [Baumeister *et al.*, 2012] Dorothea Baumeister, Piotr Faliszewski, Jérôme Lang, and Jörg Rothe. Campaigns for lazy voters: Truncated ballots. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2012.
- [de Borda, 1784] Jean-Charles de Borda. Memoire sur les elections au scrutin. *Histoire de l'Academie Royale des Sciences*, 1784.
- [Boutilier and Rosenschein, 2016] Craig Boutilier and Jeffrey S. Rosenschein. Incomplete information and communication in voting. In *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- [Brandt *et al.*, 2016] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- [Caragiannis *et al.*, 2015] Ioannis Caragiannis, George Krimpas, and Alexandros Voudouris. Aggregating partial rankings with applications to peer grading in massive online open courses. In *Proceedings of the 14th Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2015.
- [Caragiannis *et al.*, 2017] Ioannis Caragiannis, Xenophon Chatzigeorgiou, George Krimpas, and Alexandros Voudouris. Optimizing positional scoring rules for rank aggregation. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, 2017.
- [de Condorcet, 1785] Nicolas de Condorcet. Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix, 1785.
- [Conitzer *et al.*, 2017] Vincent Conitzer, Walter Sinnott-Armstrong, Jana Schaich Borg, Yuan Deng, and Max Kramer. Moral decision making frameworks for artificial intelligence. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 2017.
- [Darmann *et al.*, 2017] Andreas Darmann, Julia Grundner, and Christian Klamlner. Election outcomes under different ways to announce preferences: An analysis of the 2015 parliament election in the Austrian federal state of Styria. *Public Choice*, 173(1–2):201–216, 2017.
- [Domshlak *et al.*, 2011] Carmel Domshlak, Eyke Hüllermeier, Souhila Kaci, and Henri Prade. Preferences in AI: An overview. *Artificial Intelligence*, 175(7–8):1037–1052, 2011.
- [Eldar, 2008] Ofer Eldar. Vote-trading in international institutions. *European Journal of International Law*, 19(1):3–41, 2008.
- [Endriss *et al.*, 2009] Ulle Endriss, Maria Silvia Pini, Francesca Rossi, and K. Brent Venable. Preference aggregation over restricted ballot languages: Sincerity and strategy-proofness. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, 2009.
- [Endriss, 2018] Ulle Endriss. Judgment aggregation with rationality and feasibility constraints. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2018.
- [Faliszewski *et al.*, 2017] Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. Multiwinner voting: A new challenge for social choice theory. In *Trends in Computational Social Choice*. AI Access, 2017.
- [Glasser, 1959] Gerald Glasser. Game theory and cumulative voting for corporate directors. *Management Science*, 5(2):151–156, 1959.
- [Hansson and Grüne-Yanoff, 2018] Sven Ove Hansson and Till Grüne-Yanoff. Preferences. In *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, 2018.
- [Hemaspaandra *et al.*, 2005] Edith Hemaspaandra, Holger Spakowski, and Jörg Vogel. The complexity of Kemeny elections. *Theoretical Computer Science*, 349(3):382–391, 2005.
- [Kemeny, 1959] John Kemeny. Mathematics without numbers. *Daedalus*, 88:577–591, 1959.
- [Konczak and Lang, 2005] Kathrin Konczak and Jérôme Lang. Voting procedures with incomplete preferences. In *Proceedings of the IJCAI Multidisciplinary Workshop on Advances in Preference Handling (MPREF)*, 2005.
- [Lang *et al.*, 2017] Jérôme Lang, Jérôme Monnot, Arkadii Slinko, and William S. Zwicker. Beyond electing and ranking: Collective dominating chains, dominating subsets and dichotomies. In *Proceedings of the 16th Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2017.
- [Pini *et al.*, 2008] Maria Silvia Pini, Francesca Rossi, K. Brent Venable, and Toby Walsh. Aggregating partially ordered preferences. *Journal of Logic and Computation*, 19(3):475–502, 2008.
- [Terzopoulou *et al.*, 2018] Zoi Terzopoulou, Ulle Endriss, and Ronald de Haan. Aggregating incomplete judgments: Axiomatisations for scoring rules. In *Proceedings of the 7th International Workshop on Computational Social Choice (COMSOC)*, 2018.
- [Xia and Conitzer, 2008] Lirong Xia and Vincent Conitzer. Determining possible and necessary winners under common voting rules given partial orders. In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence (AAAI)*, 2008.
- [Zwicker, 2016] William S. Zwicker. Introduction to the theory of voting. In *Handbook of Computational Social Choice*. Cambridge University Press, 2016.