

Binary Aggregation by Selection of the Most Representative Voter

Ulle Endriss

ILLIC, University of Amsterdam
ulle.endriss@uva.nl

Umberto Grandi

University of Padova
umberto.uni@gmail.com

Abstract

In a binary aggregation problem, a group of voters each express yes/no choices regarding a number of possibly correlated issues and we are asked to decide on a collective choice that accurately reflects the views of this group. A good collective choice will minimise the distance to each of the individual choices, but using such a distance-based aggregation rule is computationally intractable. Instead, we explore a class of aggregation rules that select the most representative voter in any given situation and return that voter’s choice as the collective outcome. Two such rules, the average-voter rule and the majority-voter rule, are particularly attractive. We analyse their social choice-theoretic properties, their algorithmic efficiency, and the extent to which they are able to approximate the ideal defined by the distance-based rule. We also discuss the relevance of our results for the related framework of preference aggregation.

1 Introduction

Multiple AI applications now make use of collective decision making technologies. Examples range from multiagent planning, to crowdsourcing and human computation, to collaborative filtering for recommender systems, to rank aggregation for search engines, to coordination and resource allocation in multiagent systems. Several formal frameworks have been proposed in the literature on computational social choice [Chevalleyre *et al.*, 2007; Brandt *et al.*, 2013] to study these problems. The best known such frameworks are *voting theory*, in which a choice is made from a set of alternatives given the preferences of a group of agents, and *preference aggregation*, in which several preferences are aggregated into a single collective preference order [Arrow *et al.*, 2002]. Related frameworks dealing with information other than preferences are *belief merging* [Konieczny and Pino Pérez, 2002] and *judgment aggregation* [List and Puppe, 2009].

Here we focus on a setting in which individuals make yes/no choices on several binary issues and we need to aggregate this information into a collective view. This framework is known to be general enough to subsume both preference

aggregation and judgment aggregation [Grandi and Endriss, 2011]. Consider this example with three issues and 41 voters:

Issue:	1	2	3
20 voters:	0	1	1
10 voters:	1	0	1
11 voters:	1	1	0

What would be a good collective choice? A natural approach is to minimise a notion of distance from the individual choices to get the best compromise. This idea has been used in preference aggregation [Kemeny, 1959], belief merging [Konieczny and Pino Pérez, 2002], and judgment aggregation [Pigozzi, 2006; Miller and Osherson, 2009; Lang *et al.*, 2011]. In our example, the distance-based rule would suggest the combination (1, 1, 1), as it minimises the (Hamming) distance to the individual ballots: there are 41 disagreements (each voter disagrees on exactly one issue). But now suppose that (1, 1, 1) is not a *feasible* outcome (maybe, due to a budget constraint, we can accept at most two out of three proposals). If some outcomes are excluded, then distance-based aggregation quickly becomes highly intractable [Hemaspaandra *et al.*, 2005; Endriss *et al.*, 2012].

To tackle this problem, we propose to hold on to the idea of minimisation, but to restrict the space of outcomes considered during minimisation to the individual choices provided. That is, we propose to look for *the most representative voter* and to return that voter’s ballot as the outcome. In our example, a natural choice would be any of the voters voting (0, 1, 1). The distance of this choice to the individual ballots is 42 (21 voters disagree on 2 issues each), i.e., this solution is only marginally worse than the solution returned by the full distance-based rule, and it is optimal if (1, 1, 1) is infeasible.

We focus on two natural selection methods: the *average-voter rule* (selecting the voter who is closest to the “vector of averages”) and the *majority-voter rule* (selecting the voter who is closest to the outcome of the simple majority rule). Despite their simplistic definition, these rules turn out to be surprisingly attractive aggregation methods. They have very low computational complexity, they have interesting social choice-theoretic properties, they are guaranteed to never produce an inconsistent outcome, they can easily be explained to voters, and they are good approximations of the ideal defined by the much more complex distance-based rule.

In Section 2 we introduce our formal model and Section 3

is a comparison of the average-voter rule and the majority-voter rule. We then focus on the extent to which these rules can approximate the ideal of the distance-based rule in Section 4, and we compare our findings to known results in preference aggregation in Section 5. Section 6 concludes.

2 The Model

In this section we recall the framework of *binary aggregation with integrity constraints* [Grandi and Endriss, 2010; 2011], which we shall be working with. It is a variant of both binary aggregation with explicitly specified feasible sets [Dokow and Holzman, 2010] and judgment aggregation [List and Puppe, 2009], and all of our results can easily be translated into these other frameworks as well. Besides introducing the basic framework, we also define several aggregation rules and state some of their fundamental properties.

2.1 Basic Definitions

Let $\mathcal{I} = \{1, \dots, m\}$ be a finite set of *issues*. We want to model collective decision making problems where a group of voters have to jointly decide for which issues to choose “yes” and for which to choose “no”. A *ballot* B is an element of $\{0, 1\}^m$, which associates each issue with either a 1 (“yes”) or a 0 (“no”). We write b_j for the j th element of ballot B .

In general, not every element of $\{0, 1\}^m$ might be a *feasible* or *rational* choice. For instance, if the issues are projects that may or may not get funded, then a budget constraint might mean that no outcome with more than, say, five 1’s is feasible. We shall assume that the same constraints apply both to the individual ballots and to the outcomes of aggregation. The range of rational ballots (and thus of feasible outcomes) can be specified in different ways. In most work on binary aggregation they are given explicitly (see, e.g., Dokow and Holzman [2010]). In (formula-based) judgment aggregation constraints are implicit in the notion of consistency of a judgment set [List and Puppe, 2009]. Here we represent the set of rational ballots in a compact way by means of a formula in a propositional language [Grandi and Endriss, 2010; 2011]. Formally, let $PS = \{p_1, \dots, p_m\}$ be a set of propositional symbols, one for each issue in \mathcal{I} . An *integrity constraint* is a formula $IC \in \mathcal{L}_{PS}$, where \mathcal{L}_{PS} is obtained from PS by closing under the standard propositional connectives ($\neg, \wedge, \vee, \rightarrow, \leftrightarrow$). Let $\text{Mod}(IC) \subseteq \{0, 1\}^m$ denote the set of models of IC , i.e., the set of rational ballots satisfying IC .

Let $\mathcal{N} = \{1, \dots, n\}$ be a finite set of *voters* (we shall assume $n \geq 2$ throughout). A *profile* is a vector of rational ballots $\mathbf{B} = (B_1, \dots, B_n) \in \text{Mod}(IC)^n$, one for each voter. We write $b_{i,j}$ for the j th element of ballot B_i , the i th element of profile \mathbf{B} . The *support* of a profile $\mathbf{B} = (B_1, \dots, B_n)$ is the set of all ballots that occur at least once within \mathbf{B} :

$$\text{SUPP}(\mathbf{B}) = \{B_1\} \cup \dots \cup \{B_n\}.$$

An (irresolute) *aggregation rule* $F : \{0, 1\}^{m \times n} \rightarrow 2^{\{0, 1\}^m}$ is a function that associates with every profile \mathbf{B} a non-empty set of collective ballots $F(\mathbf{B})$. That is, the outcome of aggregation is a set of elements of the same type as our ballots, which is why we use the term ballot also to refer to outcomes.

An example of an aggregation rule is the *majority rule*, which accepts an issue if a majority of the voters accept it.

There are two possible definitions for this rule: the *weak majority rule* defined as $\text{W-Maj}(\mathbf{B})_j = 1$ iff $|\{i \in \mathcal{N} \mid b_{i,j} = 1\}| \geq \lceil \frac{n}{2} \rceil$, and the *strict majority rule* defined as $\text{S-Maj}(\mathbf{B})_j = 1$ iff $|\{i \in \mathcal{N} \mid b_{i,j} = 1\}| \geq \lceil \frac{n+1}{2} \rceil$. Observe that both rules are *resolute*, i.e., the collective outcome is always a single binary ballot. We define the majority rule Maj as the irresolute aggregation rule that outputs the union of the strict and the weak majority outcome: $\text{Maj}(\mathbf{B}) = \{\text{W-Maj}(\mathbf{B})\} \cup \{\text{S-Maj}(\mathbf{B})\}$.

In the presence of an integrity constraint, a rule may sometimes output an irrational ballot from a rational profile. Consider for instance the following example, in which the integrity constraint $IC = p_C \leftrightarrow p_A \wedge p_B$ forces individuals to accept issue C if and only if the first two issues are accepted:¹

Issue:	A	B	C
1 voter:	0	1	0
1 voter:	1	0	0
1 voter:	1	1	1
Maj:	1	1	0

We call an aggregation rule *collectively rational* wrt. an integrity constraint IC if all ballots in $F(\mathbf{B})$ satisfy IC whenever \mathbf{B} is composed of rational ballots, i.e., whenever all B_i satisfy IC . Most paradoxes studied in social choice theory can be viewed as failures of collective rationality wrt. a suitable integrity constraint [Grandi, 2012].

2.2 The Distance-Based Rule

The Hamming distance between *two ballots* $B = (b_1, \dots, b_m)$ and $B' = (b'_1, \dots, b'_m)$ is defined as the number of issues on which they differ:

$$H(B, B') = |\{j \in \mathcal{I} \mid b_j \neq b'_j\}|$$

For example, $H((1, 0, 0), (1, 1, 1)) = 2$. The Hamming distance between a *ballot* B and a *profile* \mathbf{B} is the sum of the Hamming distances between B and the ballots in \mathbf{B} :

$$\mathcal{H}(B, \mathbf{B}) = \sum_{i \in \mathcal{N}} H(B, B_i)$$

Let S and S' be sets of ballots. By a slight abuse of notation, we write $\mathcal{H}(S, \mathbf{B})$ for $\{\mathcal{H}(B, \mathbf{B}) \mid B \in S\}$, and $\mathcal{H}(S, \mathbf{B}) \leq \mathcal{H}(S', \mathbf{B})$ iff $\max(\mathcal{H}(S, \mathbf{B})) \leq \min(\mathcal{H}(S', \mathbf{B}))$.

Definition 1. Given an integrity constraint IC , the *distance-based rule* DBR^{IC} is the following function:

$$\text{DBR}^{IC}(\mathbf{B}) = \underset{B \in \text{Mod}(IC)}{\text{argmin}} \sum_{i \in \mathcal{N}} H(B, B_i)$$

Thus, winning ballots under the DBR^{IC} are rational ballots that minimise disagreement with the individual ballots. Note that the DBR^{IC} is collectively rational by definition (outcomes are chosen from $\text{Mod}(IC)$). Also note that the definition of the distance-based rule is dependent on the IC .

Fact 1. If $IC = \top$, then $\text{DBR}^{IC} = \text{Maj}$.

¹In the literature on judgment aggregation, this example is known as the *discursive dilemma* [List and Puppe, 2009].

That is, if the IC does not restrict the set of ballots, the outcome of the DBR coincides with that of the majority rule.

The DBR^{IC} has good social choice-theoretic properties, and—in its preference aggregation version known as the Kemeny rule [Kemeny, 1959]—is one of the most studied aggregation rules. However, it has a prohibitively high computational complexity: winner determination is Θ_2^P -complete [Hemaspaandra *et al.*, 2005; Endriss *et al.*, 2012].

2.3 Rules based on Representative Voters

A simple idea to reconcile distance minimisation with algorithmic efficiency is to restrict the search for a representative collective view to the set of ballots submitted by the individuals. This gives rise to a class of aggregation rules known as *generalised dictatorships* [Grandi and Endriss, 2010]. Rules in this class are collectively rational for *every possible* IC, and no rule outside this class has this desirable property. Furthermore, they satisfy certain desirable social choice-theoretic axioms, notably *unanimity* and *neutrality* [Grandi and Endriss, 2011]. Still, not all such rules are “good” rules: a (proper) dictatorship that chooses as collective outcome the ballot of the *same* voter in all profiles is certainly not a desirable rule. The problem of selecting the most representative voter is thus crucial to obtaining interesting rules in this class.

How should we select this “most representative voter” for a given profile? There arguably are two natural choices:

Definition 2. *The average-voter rule is the aggregation rule that selects those individual ballots that minimise the Hamming distance to the profile:*

$$\text{AVR}(\mathbf{B}) = \operatorname{argmin}_{B \in \text{SUPP}(\mathbf{B})} \mathcal{H}(B, \mathbf{B})$$

Definition 3. *The majority-voter rule is the aggregation rule that selects those individual ballots that minimise the Hamming distance to one of the majority outcomes:*

$$\text{MVR}(\mathbf{B}) = \operatorname{argmin}_{B \in \text{SUPP}(\mathbf{B})} \min\{H(B, B') \mid B' \in \text{Maj}(\mathbf{B})\}$$

Both of these rules are generalised dictatorships. They combine the idea of selecting a most representative voter with the basic principles at the heart of the Kemeny rule [1959] and the Slater rule [1961], respectively, in preference aggregation. The judgment aggregation rules corresponding to Kemeny and Slater, respectively, have been called *prototype* and *endpoint* by Miller and Osherson [2009].

3 Comparing the AVR and the MVR

In this section we compare our two rules based on representative voters, the AVR and the MVR. While their definitions are very similar, they can result in radically different outcomes.

Example 1. *Consider a scenario with 5 issues and 23 voters:*

Issue:	1	2	3	4	5
1 voter:	0	1	1	1	1
2 voters:	1	0	0	0	0
10 voters:	0	1	1	0	0
10 voters:	0	0	0	1	1
Maj:	0	0	0	0	0
MVR:	1	0	0	0	0
AVR:	0	1	1	0	0
AVR:	0	0	0	1	1

That is, there are two AVR-winners and each of them differs with the unique MVR-winner on a majority of the issues. It is interesting to compare these outcomes with the vector of “average votes” $(\frac{2}{23}, \frac{11}{23}, \frac{11}{23}, \frac{11}{23}, \frac{11}{23})$, showing for each issue the proportion of voters who chose 1 rather than 0. This demonstrates that—with $\frac{11}{23}$ being close to $\frac{1}{2}$ —the choices made for issues 2–5 are relatively uncritical, while the choice made for issue 1 is not. Also note that the distance to the profile is 48 for each of the AVR-winners and 65 for the MVR-winner.

In Example 1 the AVR produced outcomes that were closer to the profile than the outcome produced by the MVR. This is not a coincidence, but true for *every* profile \mathbf{B} :

Fact 2. $\mathcal{H}(\text{AVR}(\mathbf{B}), \mathbf{B}) \leq \mathcal{H}(\text{MVR}(\mathbf{B}), \mathbf{B})$ for all \mathbf{B} .

Recall that this means that the worst AVR-winner is at least as close to the profile as the best MVR-winner (in fact, all AVR-winners are equally close). Fact 2 follows immediately from the definition of AVR-winners as the set of those ballots that minimise the distance to the profile—together with the fact that the MVR also selects from the set of individual ballots.

So, if we are interested in minimising the distance to the input profile, then the AVR is superior to the MVR. On the other hand, the computational complexity of computing winners is lower for the MVR than for the AVR:

Fact 3. *Winner determination for the MVR is in $O(mn)$.*

Fact 4. *Winner determination for the AVR is in $O(mn \log n)$.*

For the MVR, we can compute the majority outcome in $O(mn)$. We need a further $O(mn)$ steps to compare each of the n ballots on each of the m issues with that majority outcome. For the AVR, we first compute, for each issue, the number of voters choosing 1 in $O(mn)$. Then for each of the n individual ballots B , we check how far the vector $n \cdot B$ is from that vector of sums, on each of the m issues. The additional complexity in the case of the AVR is due to the fact that, for each issue, we have to work with numbers that require up to $O(\log n)$ bits to be represented.

A third way of comparing two rules is to use normative arguments. Next we identify a normatively appealing property (i.e., an *axiom* in the language of social choice theory) that is satisfied by the AVR but not by the MVR. The axiom in question is closely related to the *reinforcement* axiom (often, if somewhat untowardly, referred to as *consistency*) introduced by Young in his work on the characterisation of the positional scoring rules in classical voting theory [Young, 1975].

Suppose two electorates $\mathcal{N} = \{1, \dots, n\}$ and $\mathcal{N}' = \{1, \dots, n'\}$ each vote on the same set of issues \mathcal{I} , resulting in the profiles \mathbf{B} and \mathbf{B}' . Further suppose we use an aggregation rule that is well-defined for any number of voters. Then, if a particular ballot wins both under \mathbf{B} and \mathbf{B}' , we should expect it to also win under $\mathbf{B} \oplus \mathbf{B}' = (B_1, \dots, B_n, B'_1, \dots, B'_{n'})$, i.e., when the two electorates vote together in the same election. Let us make this intuitive idea precise:

Definition 4. *An aggregation rule F satisfies reinforcement if for any two profiles \mathbf{B} and \mathbf{B}' with $\text{SUPP}(\mathbf{B}) = \text{SUPP}(\mathbf{B}')$ and $F(\mathbf{B}) \cap F(\mathbf{B}') \neq \emptyset$ we have $F(\mathbf{B} \oplus \mathbf{B}') = F(\mathbf{B}) \cap F(\mathbf{B}')$.*

Reinforcement is certainly a desirable property: if two groups independently agree that a certain outcome is best, we would expect them to uphold this choice when choosing together.

Proposition 5. *The AVR satisfies reinforcement.*

Proof. We shall make use of the fact that for any ballot B and any two profiles \mathbf{B} and \mathbf{B}' , the following holds:

$$\mathcal{H}(B, \mathbf{B}) + \mathcal{H}(B, \mathbf{B}') = \mathcal{H}(B, \mathbf{B} \oplus \mathbf{B}') \quad (1)$$

Take any two profiles \mathbf{B} and \mathbf{B}' with $\text{SUPP}(\mathbf{B}) = \text{SUPP}(\mathbf{B}')$ and $\text{AVR}(\mathbf{B}) \cap \text{AVR}(\mathbf{B}') \neq \emptyset$. We need to show that $\text{AVR}(\mathbf{B} \oplus \mathbf{B}') = \text{AVR}(\mathbf{B}) \cap \text{AVR}(\mathbf{B}')$.

For the first direction, let $B^* \in \text{AVR}(\mathbf{B}) \cap \text{AVR}(\mathbf{B}')$. By Equation (1), as B^* minimises both $\mathcal{H}(B, \mathbf{B})$ and $\mathcal{H}(B, \mathbf{B}')$ amongst all B , it must also minimise $\mathcal{H}(B, \mathbf{B} \oplus \mathbf{B}')$. In other words, $B^* \in \text{AVR}(\mathbf{B} \oplus \mathbf{B}')$.

For the other direction, let $B^* \in \text{AVR}(\mathbf{B} \oplus \mathbf{B}')$. For the sake of contradiction, suppose $B^* \notin \text{AVR}(\mathbf{B}) \cap \text{AVR}(\mathbf{B}')$. W.l.o.g., let $B^* \notin \text{AVR}(\mathbf{B})$. Choose any $B \in \text{AVR}(\mathbf{B}) \cap \text{AVR}(\mathbf{B}')$. As B^* is in the support of both \mathbf{B} and \mathbf{B}' , i.e., as B did beat (did draw with) B^* for \mathbf{B} (for \mathbf{B}') we get:

$$\begin{aligned} \mathcal{H}(B, \mathbf{B}) &< \mathcal{H}(B^*, \mathbf{B}) \\ \mathcal{H}(B, \mathbf{B}') &\leq \mathcal{H}(B^*, \mathbf{B}') \end{aligned}$$

But together with Equation (1), applied first to B and then to B^* , this yields $\mathcal{H}(B, \mathbf{B} \oplus \mathbf{B}') < \mathcal{H}(B^*, \mathbf{B} \oplus \mathbf{B}')$. Thus, B beats B^* for profile $\mathbf{B} \oplus \mathbf{B}'$ under the AVR, i.e., B^* does not win and we have obtained a the required contradiction. \square

Proposition 6. *The MVR violates reinforcement.*

Proof. We construct a counterexample. Consider the following two elections on 5 issues, with 12 voters each. Each of them has a single majority winner:

Issue:	1	2	3	4	5	Issue:	1	2	3	4	5
1 voter:	1	1	1	1	1	1 voter:	1	1	1	1	1
2 voters:	1	1	0	0	0	2 voters:	1	1	0	0	0
3 voters:	0	1	1	1	0	1 voter:	0	1	1	1	0
1 voter:	0	1	1	0	1	3 voters:	0	1	1	0	1
3 voters:	1	0	1	1	0	2 voters:	1	0	1	1	0
2 voters:	1	0	1	0	1	3 voters:	1	0	1	0	1
Maj:	1	1	1	1	0	Maj:	1	1	1	0	1

For the first election, the individual ballots that are closest to the majority outcome, i.e., the MVR-winners, are $(1, 1, 1, 1, 1)$, $(0, 1, 1, 1, 0)$ and $(1, 0, 1, 1, 0)$. They all have distance 1 to the majority outcome. For the second election the MVR-winners are $(1, 1, 1, 1, 1)$, $(0, 1, 1, 0, 1)$ and $(1, 0, 1, 0, 1)$, again with distance 1. That is, the intersection of the two sets of winners is nonempty and includes only $(1, 1, 1, 1, 1)$. Therefore, if the MVR were to satisfy reinforcement, the only winner of the election we obtain if we join the two electorates should also be $(1, 1, 1, 1, 1)$. This is an election with 24 voters:

Issue:	1	2	3	4	5
2 voters:	1	1	1	1	1
4 voters:	1	1	0	0	0
4 voters:	0	1	1	1	0
4 voters:	0	1	1	0	1
5 voters:	1	0	1	1	0
5 voters:	1	0	1	0	1
Maj:	1	1	1	0	0

The distance between the majority outcome and $(1, 1, 1, 1, 1)$ is 2, while the distance between the majority outcome and $(1, 1, 0, 0, 0)$ is only 1. Hence, $(1, 1, 1, 1, 1)$ cannot be an MVR-winner; thus, the MVR violates reinforcement. \square

Hence, whenever we consider reinforcement an important property, we should prefer the AVR over the MVR.

4 Approximation Results

If we consider the distance to the profile a crucial parameter when assessing quality of an election outcome, then the DBR is an optimal aggregation rule. But due to its high complexity, it may not be a viable choice in practice. In this section we analyse to what extent aggregation rules based on the selection of a representative voter can approximate the DBR.

Definition 5. *Let F and F' be aggregation rules. Then F is said to be an α -approximation of F' if $\mathcal{H}(F(\mathbf{B}), \mathbf{B}) \leq \alpha \cdot \mathcal{H}(F'(\mathbf{B}), \mathbf{B})$ for every profile \mathbf{B} .*

Recall that $F(\mathbf{B})$ and $F'(\mathbf{B})$ are sets of winning outcomes that need not be singletons. Given our conventions on notation, above inequality means that the *worst* F -winner has a distance to the profile that is at most α times the distance from the *best* F' -winner to the profile. F is a *strict* α -approximation of F' if the above inequality is strict (except when the second distance is zero).

Our main goal in this section will be to show that both the AVR and the MVR are strict 2-approximations of the DBR^{IC}, for any integrity constraint IC. We will also see that this bound cannot be improved further, neither for these two rules nor for any other generalised dictatorship that we might wish to use to approximate the DBR. To put our main result in context, let us first establish a very basic bound.

Proposition 7. *Every generalised dictatorship F is (at least) an $(n-1)$ -approximation of every other aggregation rule F' .*

Proof (sketch). It is easy to see that the worst case is one where $n-1$ voters submit the same ballot B , 1 voter submits a ballot \bar{B} that differs from B on every single issue, the generalised dictatorship F returns \bar{B} , and F' returns B . As in this case $\mathcal{H}(B, \mathbf{B}) = m$ and $\mathcal{H}(\bar{B}, \mathbf{B}) = m \cdot (n-1)$, we obtain the bound claimed. \square

Given that the above approximation ratio is linear in the number of voters, this is not a very attractive result. Fortunately, we can do much better if we choose the most suitable generalised dictatorship instead. The technical core of our argument is the following result, establishing an approximation of the majority rule by means of the MVR.²

Proposition 8. *The MVR is a strict 2-approximation of Maj.*

Proof. Let \mathbf{B} be any profile. The MVR selects the individual ballots that are closest to one of the outcomes in $\text{Maj}(\mathbf{B})$. Fix B^{MVR} to be one of the worst ballots in $\text{MVR}(\mathbf{B})$, i.e., an MVR-winner that is most distant from \mathbf{B} . Note that the

²This result can also be derived using a somewhat simpler proof based on the triangular inequality, similarly to Footnote 3. The proof we present here arguably has the advantage of being more easily adapted to prove other results such as our Propositions 12 and 13.

majority outcomes in $\text{Maj}(\mathbf{B})$ are all equally distant from \mathbf{B} . Fix B^{Maj} to be one of those that is closest to B^{MVR} . We have to show that $\mathcal{H}(B^{\text{MVR}}, \mathbf{B}) < 2 \cdot \mathcal{H}(B^{\text{Maj}}, \mathbf{B})$.

In case the majority outcome happens to be represented in the profile (i.e., in case $B^{\text{Maj}} \in \text{SUPP}(\mathbf{B})$ and thus $B^{\text{Maj}} = B^{\text{MVR}}$), we are done, as in this case the two distances are the same (i.e., the approximation ratio is 1). So from now on assume $B^{\text{MVR}} \neq B^{\text{Maj}}$.

Let B be any of the ballots in \mathbf{B} corresponding to one of the voters that is different from the voter that provided B^{MVR} . We will show that B disagrees with B^{MVR} at most twice as often as it disagrees with B^{Maj} , i.e., that $H(B^{\text{MVR}}, B) \leq 2 \cdot H(B^{\text{Maj}}, B)$. By summing up for all individual ballots in the profile we will then obtain the required approximation bound. We partition the set of issues \mathcal{I} into four sets:

- (1) $\mathcal{I}^{++} = \{j \in \mathcal{I} \mid b_j = b_j^{\text{Maj}} \text{ and } b_j = b_j^{\text{MVR}}\}$
- (2) $\mathcal{I}^{+-} = \{j \in \mathcal{I} \mid b_j = b_j^{\text{Maj}} \text{ and } b_j \neq b_j^{\text{MVR}}\}$
- (3) $\mathcal{I}^{-+} = \{j \in \mathcal{I} \mid b_j \neq b_j^{\text{Maj}} \text{ and } b_j = b_j^{\text{MVR}}\}$
- (4) $\mathcal{I}^{--} = \{j \in \mathcal{I} \mid b_j \neq b_j^{\text{Maj}} \text{ and } b_j \neq b_j^{\text{MVR}}\}$

That is, \mathcal{I}^{+-} , for instance, is the set of issues on which B agrees with the majority outcome but disagrees with the worst MVR-winner. By definition of these sets, we can now express one of our distances of interest as follows:

$$H(B^{\text{MVR}}, B) = |\mathcal{I}^{+-}| + |\mathcal{I}^{--}| \quad (2)$$

Let $k \geq 1$ be the number of issues on which B^{Maj} disagrees with B^{MVR} . By definition of the MVR, B must disagree with B^{Maj} on at least k issues:

$$k \leq H(B^{\text{Maj}}, B) \quad (3)$$

Observe that the set of issues on which B^{Maj} disagrees with B^{MVR} is the union of the set of issues \mathcal{I}^{+-} that B^{Maj} agrees on with B but B disagrees on with B^{MVR} and the set of issues \mathcal{I}^{-+} that B^{Maj} disagrees on with B but B agrees on with B^{MVR} . Hence, $k = |\mathcal{I}^{+-} \cup \mathcal{I}^{-+}|$. As an immediate consequence, we obtain that $k \geq |\mathcal{I}^{+-}|$. Due to the fact that B cannot disagree with B^{Maj} on more issues than B^{MVR} does (i.e., on more than k issues), we also have $k \geq |\mathcal{I}^{-+} \cup \mathcal{I}^{--}|$. Thus, we also get $k \geq |\mathcal{I}^{--}|$. We can now put together Equation (2) with the two inequalities just obtained:

$$H(B^{\text{MVR}}, B) \leq 2k \quad (4)$$

Equations (3) and (4) together show that B will disagree with B^{MVR} at most twice as many times as with B^{Maj} . This is true for every individual ballot B other than the MVR-winner itself. For the latter, the disagreement with B^{MVR} (i.e., with itself) is zero, while it is non-zero with B^{Maj} (recall that we assumed $B^{\text{MVR}} \neq B^{\text{Maj}}$). Thus, overall we obtain an approximation ratio that is *strictly* better than 2. \square

Corollary 9. *Both the AVR and the MVR are strict 2-approximations of the DBR^{IC} for $\text{IC} = \top$.*

Proof. Immediate from Proposition 8 together with the fact that the majority rule is equivalent to the DBR^{\top} (Fact 1) and the fact that the AVR-winner is always at least as close to the profile as the MVR-winner (Fact 2). \square

What about the DBR for other integrity constraints? As it turns out, $\text{IC} = \top$ is in fact the worst case and we can do at least as well for any other integrity constraint. This will follow from our next result, which shows that the stronger the integrity constraint, the more distant the outcome of the DBR will be from the profile. This result compares the distance to the profile achieved by the DBR for two different integrity constraints. Observe that IC logically entails IC' iff $\text{Mod}(\text{IC}) \subseteq \text{Mod}(\text{IC}')$. That is, any profile $\mathbf{B} \in \text{Mod}(\text{IC})^n$ that is admissible for DBR^{IC} will also be admissible for $\text{DBR}^{\text{IC}'}$. In other words, both the DBR^{IC} and the $\text{DBR}^{\text{IC}'}$ are well-defined on any such \mathbf{B} .

Lemma 10. *If IC entails IC' , then $\mathcal{H}(\text{DBR}^{\text{IC}}(\mathbf{B}), \mathbf{B}) \geq \mathcal{H}(\text{DBR}^{\text{IC}'}(\mathbf{B}), \mathbf{B})$ for every profile $\mathbf{B} \in \text{Mod}(\text{IC})^n$.*

Proof. To prove the claim, it suffices to observe that both the DBR^{IC} and the $\text{DBR}^{\text{IC}'}$ aim at minimising the same objective function (namely the Hamming distance between the profile and the winning ballot), while the $\text{DBR}^{\text{IC}'}$ can select from a larger set of candidate ballots. \square

We can now state our main approximation result, which shows that the average-voter rule and the majority-voter rule, besides having good axiomatic properties and being easy to compute, are both good approximations of the much more complex distance-based rule, independently from the integrity constraint used to delimit the set of feasible outcomes.

Theorem 11. *Both the AVR and the MVR are strict 2-approximations of the DBR^{IC} for any integrity constraint IC .*

Proof. By Corollary 9, both $\mathcal{H}(\text{AVR}(\mathbf{B}), \mathbf{B})$ and $\mathcal{H}(\text{MVR}(\mathbf{B}), \mathbf{B})$ are strictly less than $2 \cdot \mathcal{H}(\text{DBR}^{\top}(\mathbf{B}), \mathbf{B})$. We can now use Lemma 10, together with the fact that any formula IC logically entails \top , to conclude that this last figure is smaller than (or equal to) $2 \cdot \mathcal{H}(\text{DBR}^{\text{IC}}(\mathbf{B}), \mathbf{B})$, obtaining the inequalities required for a strict 2-approximation. \square

Note that the AVR always provides a better bound than the MVR, which, however, is easier to compute. A direct proof of the fact that the AVR is a (*non-strict*) 2-approximation of the DBR^{IC} can also be obtained using the triangle inequality.³

Is this the best we can do? Yes, as the following example will demonstrate, neither the AVR nor any other generalised dictatorship can guarantee a better approximation ratio than 2.

Example 2. *Let $n = m$, i.e., there are as many voters as there are issues. Suppose each voter $i \in \mathcal{N}$ approves only of issue i , i.e., we are considering the profile \mathbf{B} with $b_{i,i} = 1$ and $b_{i,j} = 0$ for $i \neq j$. Then the DBR^{\top} will return the outcome $B_0 = (0, \dots, 0)$. Here is an illustration for $n = 5$:*

³This is a folk theorem in preference aggregation [Ailon *et al.*, 2008]. The proof can be sketched as follows: $\mathcal{H}(B^{\text{AVR}}, \mathbf{B}) = \sum_{i=1}^n H(B^{\text{AVR}}, B_i) \leq \sum_{i=1}^n \frac{1}{n} \sum_{k=1}^n H(B_k, B_i) \leq \sum_{i=1}^n \frac{1}{n} \cdot \sum_{k=1}^n [H(B_k, B^{\text{DBR}}) + H(B^{\text{DBR}}, B_i)] = \frac{1}{n} \cdot [n \cdot \mathcal{H}(B^{\text{DBR}}, \mathbf{B}) + n \cdot \mathcal{H}(B^{\text{DBR}}, \mathbf{B})] = 2 \cdot \mathcal{H}(B^{\text{DBR}}, \mathbf{B})$.

Issue:	1	2	3	4	5
1 voter:	1	0	0	0	0
1 voter:	0	1	0	0	0
1 voter:	0	0	1	0	0
1 voter:	0	0	0	1	0
1 voter:	0	0	0	0	1
DBR [⊤] :	0	0	0	0	0

The distance between B_0 and the profile is $\mathcal{H}(B_0, \mathbf{B}) = n$ (one disagreement per voter). On the other hand, if we limit ourselves to selecting one of the individual ballots as the outcome, then the distance to the profile will be $\mathcal{H}(B_i, \mathbf{B}) = (n-1) + (n-1)$, whichever $i \in \mathcal{N}$ we pick. That is, the approximation ratio for this scenario is $2 \cdot \frac{n-1}{n}$. Hence, by increasing n and m , we can go arbitrarily close to 2.

To summarise, at this point we know that the AVR is a strict 2-approximation for all incarnations of the DBR and we cannot do better than that—at least for the weakest possible integrity constraint $\text{IC} = \top$. However, Lemma 10 opens up the possibility that better approximation ratios might be achievable for stronger integrity constraints, and Example 2 suggests that the worst case occurs when $n = m$. Indeed, under additional assumptions improved bounds are possible. For lack of space, we state the following two results without proof (the same technique as for Proposition 8 may be used). The main intuition underlying both results may be obtained by adapting Example 2 to the situations covered by these results.

Recall that we have seen that the precise approximation ratio is $2 \cdot \frac{n-1}{n}$. It is possible to show that the approximation ratio can alternatively be expressed in terms of m alone. In most real-world applications n will be much larger than m , which makes the following result attractive:

Proposition 12. *If $n > m$, then the AVR and the MVR are α -approximations of the DBR^{IC} with $\alpha = 2 \cdot \frac{m-1}{m}$ for any IC.*

For integrity constraints that are equivalent to *cubes*, i.e., conjunctions of literals (without repeated or complementary literals), we obtain an even better approximation ratio:

Proposition 13. *If $n > m$ and if IC is a cube of length k , then the AVR and the MVR are α -approximations of the DBR^{IC} with $\alpha = 2 \cdot \frac{m-k-1}{m-k}$.*

5 An Application to Preference Aggregation

Next we show how our approach yields new approximation results for the Kemeny rule in preference aggregation. This problem has been the subject of several publications: Dwork *et al.* [2001] present a 2-approximation, Ailon *et al.* [2008] use a randomised process to obtain an $\frac{11}{7}$ -approximation, and Kenyon-Mathieu and Schudy [2007] provide a PTAS for this optimisation problem, i.e., a polynomial algorithm which, given an instance of the problem and a positive number ϵ , returns a $(1 + \epsilon)$ -approximation of the optimum.

Preference aggregation can be viewed as an instance of binary aggregation by devising a suitable integrity constraint: issues are propositions of the form $p_{a \succ b}$, and IC encodes the properties of linear orders [Grandi and Endriss, 2011].⁴ Our

⁴For instance, to model transitivity, IC has to include conjuncts of the form $p_{a \succ b} \wedge p_{b \succ c} \rightarrow p_{a \succ c}$ for all triples of alternatives a, b, c .

approximation results thus transfer to the framework of preference aggregation, showing that the AVR and the MVR both are strict 2-approximations of the Kemeny rule. While a 2-approximation result using (the equivalent of) the AVR was previously known [Ailon *et al.*, 2008], we have strengthened this result by showing that the approximation is strict. The MVR, to the best of our knowledge, has never been considered for preference aggregation. It is a strict 2-approximation of the Kemeny rule that is computable in linear time.

Thus, while there are sharper approximation results in the literature specific to the problem of preference aggregation, both the AVR and the MVR are attractive as they have very low computational complexity and as they have a natural interpretation as rules based on the selection of the most representative voter, while the procedures used to obtain a PTAS are algorithmically interesting but do not lead to the definition of normatively appealing aggregation rules. The AVR and the MVR arguably are also easy to explain, which will be relevant for elections involving human voters.

6 Conclusion

We have argued that simple aggregation rules that return the proposal made by the most representative voter as the outcome of a collective decision making process have surprisingly attractive properties. We have developed our results in the framework of binary aggregation with integrity constraints, but they immediately extend to other binary aggregation frameworks and to judgment aggregation as well.

We have focussed on two representative-voter rules: the average-voter rule, which may be considered the result of combining the representative-voter idea with the principle underlying the Kemeny rule familiar from preference aggregation; and the majority-voter rule, which does the corresponding thing for the Slater rule. Besides Kemeny and Slater, there is a third preference aggregation rule that naturally extends to binary aggregation, namely Tideman’s method of *Ranked Pairs* [Tideman, 1987; Zavist and Tideman, 1989]: Under this rule, we accept choices for the issues in the order of the strength of the majorities supporting them, unless doing so would violate the IC. In judgment aggregation, this rule has been proposed under the names of *ranked-agenda rule* [Lang *et al.*, 2011] and *support-based procedure* [Porello and Endriss, 2011]. Winner determination for Tideman’s rule is intractable in the general case [Brill and Fischer, 2012] and it is thus interesting to study the approximation ratio of its representative-voter version.

More generally speaking, for any aggregation rule that selects from the set of all feasible outcomes according to some notion of optimality, we may define a corresponding representative-voter rule by restricting the search space to the individual ballots. Such rules represent good compromises between algorithmic considerations and the need for principled methods of aggregation. The extent to which these rules can approximate more complex ones opens up a wide range of interesting questions for future work. These questions may be studied for binary aggregation and for restricted domains such as preference aggregation, where intractability results for optimality-based rules are widespread [Hudry, 2012].

References

- [Ailon *et al.*, 2008] N. Ailon, M. Charikar, and A. Newman. Aggregating inconsistent information: Ranking and clustering. *Journal of the ACM*, 55(5), 2008.
- [Arrow *et al.*, 2002] K. J. Arrow, A. K. Sen, and K. Suzumura, editors. *Handbook of Social Choice and Welfare*. North-Holland, 2002.
- [Brandt *et al.*, 2013] F. Brandt, V. Conitzer, and U. Endriss. Computational social choice. In G. Weiss, editor, *Multiagent Systems*. MIT Press, 2013.
- [Brill and Fischer, 2012] M. Brill and F. A. Fischer. The price of neutrality for the ranked pairs method. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI-2012)*, 2012.
- [Chevalleyre *et al.*, 2007] Y. Chevalleyre, U. Endriss, J. Lang, and N. Maudet. A short introduction to computational social choice. In *Proceedings of the 33rd Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM-2007)*. Springer-Verlag, 2007.
- [Dokow and Holzman, 2010] E. Dokow and R. Holzman. Aggregation of binary evaluations. *Journal of Economic Theory*, 145(2):495–511, 2010.
- [Dwork *et al.*, 2001] C. Dwork, R. Kumar, M. Naor, and D. Sivakumar. Rank aggregation methods for the web. In *Proceedings of the 10th International Conference on World Wide Web (WWW-2001)*, 2001.
- [Endriss *et al.*, 2012] U. Endriss, U. Grandi, and D. Porello. Complexity of judgment aggregation. *Journal of Artificial Intelligence Research*, 45:481–514, 2012.
- [Grandi and Endriss, 2010] U. Grandi and U. Endriss. Lifting rationality assumptions in binary aggregation. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI-2010)*, 2010.
- [Grandi and Endriss, 2011] U. Grandi and U. Endriss. Binary aggregation with integrity constraints. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI-2011)*, 2011.
- [Grandi, 2012] U. Grandi. The common structure of paradoxes in aggregation theory. In *Proceedings of the 4th International Workshop on Computational Social Choice (COMSOC-2012)*. AGH University of Science and Technology, Kraków, 2012.
- [Hemaspaandra *et al.*, 2005] E. Hemaspaandra, H. Spakowski, and J. Vogel. The complexity of Kemeny elections. *Theoretical Computer Science*, 349(3):382–391, 2005.
- [Hudry, 2012] Olivier Hudry. On the computation of median linear orders, of median complete preorders and of median weak orders. *Mathematical Social Sciences*, 64(1):2–10, 2012.
- [Kemeny, 1959] J. Kemeny. Mathematics without numbers. *Daedalus*, 88:577–591, 1959.
- [Kenyon-Mathieu and Schudy, 2007] C. Kenyon-Mathieu and W. Schudy. How to rank with few errors. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing (STOC-2007)*, 2007.
- [Konieczny and Pino Pérez, 2002] S. Konieczny and R. Pino Pérez. Merging information under constraints: A logical framework. *Journal of Logic and Computation*, 12(5):773–808, 2002.
- [Lang *et al.*, 2011] J. Lang, G. Pigozzi, M. Slavkovik, and L. van der Torre. Judgment aggregation rules based on minimization. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-XIII)*, 2011.
- [List and Puppe, 2009] C. List and C. Puppe. Judgment aggregation: A survey. In *Handbook of Rational and Social Choice*. Oxford University Press, 2009.
- [Miller and Osherson, 2009] M. K. Miller and D. Osherson. Methods for distance-based judgment aggregation. *Social Choice and Welfare*, 32(4):575–601, 2009.
- [Pigozzi, 2006] G. Pigozzi. Belief merging and the discursive dilemma: An argument-based account of paradoxes of judgment aggregation. *Synthese*, 152(2):285–298, 2006.
- [Porello and Endriss, 2011] D. Porello and U. Endriss. Ontology merging as social choice. In *Proceedings of the 12th International Workshop on Computational Logic in Multiagent Systems (CLIMA-2011)*. Springer-Verlag, 2011.
- [Slater, 1961] P. Slater. Inconsistencies in a schedule of paired comparisons. *Biometrika*, 48(3-4):303–312, 1961.
- [Tideman, 1987] T. N. Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4(3):185–206, 1987.
- [Young, 1975] H. P. Young. Social choice scoring functions. *SIAM Journal on Applied Mathematics*, 28(4):824–838, 1975.
- [Zavist and Tideman, 1989] T. M. Zavist and T. N. Tideman. Complete independence of clones in the ranked pairs rule. *Social Choice and Welfare*, 6(2):167–173, 1989.