Binary Aggregation by Selection of the Most Representative Voter

Ulle Endriss
Institute for Logic, Language and Computation
University of Amsterdam
The Netherlands
ulle.endriss@uva.nl

Umberto Grandi
Department of Mathematics
University of Padova
Italy
umberto.uni@gmail.com

Abstract

In binary aggregation, each member of a group expresses yes/no choices regarding several correlated issues and we need to decide on a collective choice that accurately reflects the views of the 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 low-complexity aggregation rules that select the most representative voter in any given situation and return that voter's choice as the outcome.

1 Introduction

Many 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 frameworks have been proposed in the literature on computational social choice (Chevaleyre et al. 2007; Brandt, Conitzer, and Endriss 2013) to study these problems. The best known 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, Sen, and Suzumura 2002). Related frameworks for information other than preferences are *belief merging* (Konieczny and Pino Pérez 2011) 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 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 the distance from the individual choices. This idea has been used in preference aggregation (Kemeny 1959), belief merging (Konieczny and Pino Pérez 2002), and judgment aggregation (Miller and Osherson 2009; Lang et al. 2011). In our example, the distance-based rule returns $(1, 1, 1)$, as that 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 proposals). If some outcomes are excluded, then distance-based aggregation quickly becomes intractable (Endriss, Grandi, and Porello 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 look for *the most representative voter* and return her 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 distance-based rule—and it is optimal in case $(1, 1, 1)$ is infeasible.

We focus on two natural selection methods: the *average-voter rule* (selecting the voter closest to the "vector of averages") and the *majority-voter rule* (selecting the voter closest to the outcome of the simple majority rule). Despite their simplistic definitions, these rules turn out to be surprisingly attractive. They are of low computational complexity, they have good social choice-theoretic properties, they are guaranteed to never produce an infeasible 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. Such approximation results form the technical core of the paper. We also introduce a third representative-voter rule, the *ranked-voter rule*, which—although based on similarly natural basic principles—is considerably less attractive from the point of view of approximation.

In Section 2 we introduce our formal model. Section 3 presents our approximation results, while Section 4 compares the two most attractive rules further, in view of their complexity and w.r.t. a basic axiom.

2 The Model

In this section we recall the framework of *binary aggregation with integrity constraints* (Grandi and Endriss 2011;

2013). 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. We also define several concrete aggregation rules.

2.1 Basic Definitions

Let $\mathcal{I} = \{1, \dots, m\}$ be a finite set of *issues*. We want to model 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 jth element of ballot B.

Not every element of $\{0,1\}^m$ might be a *feasible* or *rational* choice. For instance, if the issues are funding decisions, 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 to both individual ballots and outcomes. Let *PS* $= \{p_1, \dots, p_m\}$ be a set of propositional variables, one for each issue in $\mathcal{I}$. An *integrity constraint* is a consistent formula IC in the language obtained from *PS* by closing under the standard propositional connectives ($\neg$, $\wedge$, $\vee$, $\rightarrow$, $\leftrightarrow$). Let Mod(IC) $\subseteq \{0,1\}^m$ be 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* (with $n \geqslant 2$). A *profile* is a vector of rational ballots $\boldsymbol{B} = (B_1, \dots, B_n) \in \mathrm{Mod(IC)}^n$, one for each voter. We write $b_{i,j}$ for the jth element of ballot B_i, the ith element of profile $\boldsymbol{B}$. We write $N^{\boldsymbol{B}}_{j:x} = \{i \in \mathcal{N} \mid b_{i,j} = x\}$ for the set of voters choosing value $x \in \{0,1\}$ for issue j in profile $\boldsymbol{B}$. The *support* of a profile $\boldsymbol{B} = (B_1, \dots, B_n)$ is the set of all ballots that occur at least once within $\boldsymbol{B}$: $\textsc{Supp}(\boldsymbol{B}) = \{B_1, \dots, B_n\}$. An (irresolute) *aggregation rule* $F : \{0,1\}^{m \times n} \rightarrow 2^{\{0,1\}^m}$ is a function that maps every profile $\boldsymbol{B}$ to a non-empty subset of $\{0,1\}^m$. That is, F returns a set of ballots.

An example of a rule is the *majority rule*, which accepts an issue if a majority of the voters do. In fact, there are two possible definitions of "majority": we may accept an issue if *at least* half of the voters do (weak majority) or only if *more than* half of them do (strict majority). We define Maj as the irresolute aggregation rule that returns the set including both the weak and the strict majority winner (which may or may not coincide): $\mathrm{Maj}(\boldsymbol{B}) = \{B^w, B^s\}$ with $b^w_j = 1$ iff $|N^{\boldsymbol{B}}_{j:1}| \geqslant \lceil \frac{n}{2} \rceil$ and $b^s_j = 1$ iff $|N^{\boldsymbol{B}}_{j:1}| \geqslant \lceil \frac{n+1}{2} \rceil$.

In the presence of an integrity constraint, a rule may sometimes output an irrational ballot given a rational profile. An example is the following scenario, in which the integrity constraint IC $= p_C \leftrightarrow p_A \wedge p_B$ forces voters to accept issue C if and only if the first two issues are accepted:[1]

Issue:	A	B	C
Voter 1:	0	1	0
Voter 2:	1	0	0
Voter 3:	1	1	1
Maj:	1	1	0

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

An aggregation rule is called *collectively rational* w.r.t. an integrity constraint IC if all elements of $F(\boldsymbol{B})$ satisfy IC whenever $\boldsymbol{B}$ is rational, i.e., whenever all B_i satisfy IC.

2.2 The Distance-Based Rule

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

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

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

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

For $S \subseteq \{0,1\}^m$, by a slight abuse of notation, we will write $\mathcal{H}(S, \boldsymbol{B})$ as a shorthand for $\{\mathcal{H}(B, \boldsymbol{B})) \mid B \in S\}$.

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

$$\mathrm{DBR}^{\mathrm{IC}}(\boldsymbol{B}) \quad = \quad \underset{B \in \mathrm{Mod(IC)}}{\mathrm{argmin}} \ \mathcal{H}(B, \boldsymbol{B})$$

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

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

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

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). The preference aggregation version of the DBR$^{\mathrm{IC}}$ is known as the Kemeny rule (1959), and is one of the most studied aggregation rules. However, it has a prohibitively high computational complexity: winner determination is Θ^p_2-complete (Hemaspaandra, Spakowski, and Vogel 2005).

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 2013). Rules in this class are collectively rational for *every possible* IC, and no rule outside this class has this desirable property. Still, not all such rules are "good" rules: a *proper dictatorship* $F_i : \boldsymbol{B} \mapsto B_i$ that chooses as collective outcome the ballot of the *same* voter i in all profiles is certainly not attractive. 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 are several natural choices. Here we define three of them, each inspired by an existing rule:

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

$$\text{AVR}(\boldsymbol{B}) \quad = \quad \underset{B \in \text{Supp}(\boldsymbol{B})}{\text{argmin}} \ \mathcal{H}(B, \boldsymbol{B})$$

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:*

$$\text{MVR}(\boldsymbol{B}) \quad = \quad \underset{B \in \text{Supp}(\boldsymbol{B})}{\text{argmin}} \ \min\{H(B, B') \mid B' \in \text{Maj}(\boldsymbol{B})\}$$

We need some further notation for the third definition. For a given profile $\boldsymbol{B}$, define the *majority strength* of issue j as $\text{MS}^{\boldsymbol{B}}(j) = \max\{|N^{\boldsymbol{B}}_{j:0}|, |N^{\boldsymbol{B}}_{j:1}|\}$, inducing an ordering $\succ^{\boldsymbol{B}}_{\tau}$ on issues, with ties broken using a permutation $\tau : \mathcal{I} \to \mathcal{I}$. For a *partial* function $\ell : \mathcal{I} \to \{0, 1\}$ and a ballot B, we write $\ell \subseteq B$ in case ℓ agrees with B on those issues on which it is defined, i.e., if $\ell(j) = x$ implies $b_j = x$. Now, for a given profile $\boldsymbol{B}$ and permutation $\tau : \mathcal{I} \to \mathcal{I}$, we define the (total) function $\ell^{\boldsymbol{B}}_{\tau}$ via the following procedure:

for $j \in \mathcal{I}$, following order $\succ^{\boldsymbol{B}}_{\tau}$ **do**
  $\ell^{\boldsymbol{B}}_{\tau}(j) := \text{Maj}(\boldsymbol{B})_j$ if $\exists\, i \in \mathcal{N}$ such that $\ell^{\boldsymbol{B}}_{\tau} \subseteq B_i$
  $\ell^{\boldsymbol{B}}_{\tau}(j) := 1 - \text{Maj}(\boldsymbol{B})_j$ otherwise

Definition 4. *The **ranked-voter rule** is the aggregation rule that selects the individual ballots returned by the above procedure for some tie-breaking rule τ:*

$$\text{RVR}(\boldsymbol{B}) \quad = \quad \{\ell^{\boldsymbol{B}}_{\tau} \mid \tau \textit{ is a permutation on } \mathcal{I}\}$$

That is, for the RVR we take over the majority choices in the order of their relative strength, subject to the availability of a representative voter agreeing with the choices made.

All three rules are generalised dictatorships. They combine the idea of selecting a most representative voter with the basic principles at the heart of three well-known rules in preference aggregation: the *Kemeny rule* (1959), the *Slater rule* (1961), and the *ranked-pairs* rule (Tideman 1987).

Example 1. *Suppose there are 6 issues and 21 voters:*

Issue:	**1**	**2**	**3**	**4**	**5**	**6**
1 voter:	1	0	0	0	0	0
10 voters:	0	1	1	0	0	0
10 voters:	0	0	0	1	1	1
Maj:	0	0	0	0	0	0
MVR:	1	0	0	0	0	0
AVR:	0	1	1	0	0	0

That is, the majority rule, the MVR*, and the* AVR *can all return distinct winners. Refer to the proof of Theorem 5 for an example where the* RVR *differs from those three rules.*

3 Approximation Results

If we consider the distance to the input profile a crucial parameter when assessing the quality of an election outcome, then the DBR is the optimal aggregation rule. But due to its high complexity, using the DBR may not always be a viable choice in practice. In this section, we therefore analyse to what extent our three representative-voter rules can approximate the DBR, for varying integrity constraints.

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

Recall that $F(\boldsymbol{B})$ and $F'(\boldsymbol{B})$ are *sets* of outcomes. Above inequality is understood to apply to *all* elements of these sets, i.e., $\max[\mathcal{H}(F(\boldsymbol{B}), \boldsymbol{B})] \leqslant \alpha \cdot \min[\mathcal{H}(F'(\boldsymbol{B}), \boldsymbol{B})]$. Thus, 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 called a *strict α-approximation* of F', if the inequality is strict (when the second distance is not zero).

As a first result, we show that the stronger the integrity constraint, the easier it is to approximate the DBR. Observe that IC logically entails IC$'$ iff Mod(IC) $\subseteq$ Mod(IC$'$). That is, any profile $\boldsymbol{B} \in \text{Mod(IC)}^n$ that is admissible for DBR$^{\text{IC}}$ will also be admissible for DBR$^{\text{IC}'}$, i.e., both the DBR$^{\text{IC}}$ and the DBR$^{\text{IC}'}$ are well-defined on any such $\boldsymbol{B}$.

Lemma 2. *If* IC *entails* IC$'$*, then* $\mathcal{H}(\text{DBR}^{\text{IC}}(\boldsymbol{B}), \boldsymbol{B}) \geqslant \mathcal{H}(\text{DBR}^{\text{IC}'}(\boldsymbol{B}), \boldsymbol{B})$ *for every profile* $\boldsymbol{B} \in \text{Mod(IC)}^n$.

Proof. It suffices to observe that the DBR$^{\text{IC}}$ and the DBR$^{\text{IC}'}$ aim at minimising the same objective function (namely the Hamming distance between $\boldsymbol{B}$ and the winning ballot), while the DBR$^{\text{IC}'}$ can select from a larger set of ballots. □

In the sequel, we will prove several approximation results w.r.t. the the majority rule Maj, which by Fact 1 is equivalent to DBR$^{\top}$, and thus a good representative of the family of distance-based rules. Positive approximation results w.r.t. Maj are particularly interesting: $\top$ is the weakest possible integrity constraint, so by Lemma 2 any such result immediately generalises to all instances of DBR$^{\text{IC}}$.

We will see that the AVR and the MVR are good approximations of the DBR (with a constant approximation ratio of 2), while the RVR is not (with a linear approximation ratio).

3.1 Baseline Results: Dictatorships

To put our main results into context, let us first establish a very basic bound for arbitrary generalised dictatorships, which applies not only to approximations of the DBR, but to approximations of arbitrary reference rules.

Proposition 3. *Every generalised dictatorship F is an $O(n)$-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 $\overline{B}$ that differs from B on every single issue, the generalised dictatorship F returns $\overline{B}$, and F' returns B. As in this case $\mathcal{H}(B, \boldsymbol{B}) = m$ and $\mathcal{H}(\overline{B}, \boldsymbol{B}) = m \cdot (n-1)$, we obtain an approximation ratio of $n - 1 \in O(n)$. □

Given that the above approximation ratio is linear in the number of voters, this is not a very attractive result. At the same time, Proposition 3 only states an upper bound; specific generalised dictatorships may do much better when approximating specific reference rules. Let us now focus on the approximability of the DBR, and specifically the basic DBR with a tautological integrity constraint, which we have seen to be equivalent to the majority rule (DBR$^{\top}$ = Maj).

How well can a *proper* dictatorship approximate Maj? The scenario given in the proof of Proposition 3 occurs when the dictator has the opposite view of all other voters. Thus, now $n-1$ is not only an upper but also a lower bound on the approximation ratio, and we obtain a second baseline result:

Proposition 4. *Every proper dictatorship $F_i : \boldsymbol{B} \mapsto B_i$ is a $\Theta(n)$-approximation of the majority rule* Maj.

That is, unsurprisingly, proper dictatorships are amongst the worst approximators of all generalised dictatorships.

3.2 Negative Results: RVR

Much more surprising is the fact that the RVR is just as bad as a proper dictatorship when it comes to approximating the majority rule (and thus the basic DBR):

Theorem 5. *The* RVR *is a $\Theta(n)$-approximation of* Maj.

Proof. The fact that the RVR is an $O(n)$-approximation of Maj follows from Proposition 3. To prove the corresponding lower bound, it suffices to show that the RVR is an $\Omega(n)$-approximation of Maj for *some* family of profiles. Thus, we may assume $m > n$. Consider the following scenario:

$$
\begin{array}{ll}
 & \overbrace{0\,1\,1\,1\,1 \cdots 1\,1}^{n-2}\overbrace{1 \cdots\cdots\cdots 1}^{m-(n-2)} \\
\text{Voter 1:} & 0\,1\,1\,1\,1 \cdots 1\,1\,1 \cdots\cdots\cdots 1 \\
\text{Voter 2:} & 1\,0\,1\,1\,1 \cdots 1\,1\,1 \cdots\cdots\cdots 1 \\
\text{Voter 3:} & 1\,1\,0\,1\,1 \cdots 1\,1\,1 \cdots\cdots\cdots 1 \\
\vdots & \qquad \vdots \qquad\qquad \vdots \\
\text{Voter } n-2\text{:} & 1\,1\,1\,1\,1 \cdots 1\,0\,1 \cdots\cdots\cdots 1 \\
\text{Voter } n-1\text{:} & 1\,1\,1\,1\,1 \cdots 1\,1\,0 \cdots\cdots\cdots 0 \\
\text{Voter } n\text{:} & 1\,1\,1\,1\,1 \cdots 1\,1\,0 \cdots\cdots\cdots 0
\end{array}
$$

The majority winner for this profile is $(1,\ldots,1)$, with a distance of $(n-2)+(m-n+2)\cdot 2$ to the profile. The majority strength is $n-1$ for the first $n-2$ issues, but only $n-2$ for the remaining ones. Hence, the RVR will first lock in a 1 for the first $n-2$ issues. But at this point only the ballot of voter n (which is equal to that of voter $n-1$) is still consistent with these choices, i.e., the outcome for the RVR must be equal to the ballot of voter n. The distance of the RVR-winner to the profile is $(n-2)+(m-n+2)\cdot(n-2)$.

The ratio between the two distances of interest grows linearly in $m-n$. As we can choose $m-n$ to grow linearly with n, our claimed lower bound follows. $\square$

In practice, the RVR will often do better than a dictatorship. What our results show is that *in the worst case* both rules do worse than the majority rule by a factor that is linear in n.

Our proof of Theorem 5 relies on a worst-case scenario in which n, the number of voters, is smaller than m, the number of issues. However, in many cases of practical interest, m will be (much) smaller than n. Can we do better in case $m \leqslant n$? Not much—even then the approximation ratio is bounded from below by a function growing linearly in m:

Theorem 6. *For aggregation problems with $m \leqslant n$, the* RVR *is an $\Omega(m)$-approximation of the majority rule* Maj.

Proof (sketch). Let $m \leqslant n$. We only sketch the proof. The basic idea is to take several copies of the gadget used in the proof of Theorem 5. Let k be a divisor of $n-1$ that is smaller than m. For each $j \leqslant k-1$, create $\frac{n-1}{k}$ voters who reject issue j and accept all others. Additionally, create $\frac{n-1}{k}+1$ voters who accept the first $k-1$ issues and reject the other $m-k+1$ ones. Also for this profile, $(1,\ldots,1)$ is the majority outcome. The first $k-1$ issues each have majority strength $n-\frac{n-1}{k}$; the others only have majority strength $n-\frac{n-1}{k}-1$. Hence, under the RVR the first $k-1$ issues are accepted and the rest are rejected. The ratio of the distance from the RVR-winner to the profile and the distance from the majority outcome to the profile grows linearly in k. As we can choose k close to m, the claimed lower bound follows. $\square$

As the RVR is based on principles that have been found to be useful in classical voting (Tideman 1987), the results above raise the question of whether there are any representative-voter rules at all that offer a better approximation ratio. We are about to give a positive answer to this question.

3.3 Positive Results: AVR and MVR

We now want to analyse the ability of the AVR and the MVR to approximate the DBR. First, observe that the AVR provides at least as good an approximation of the DBR as any other generalised dictatorship, including the MVR. This follows immediately from the definition of AVR-winners as the set of those ballots that minimise the distance to the profile. We record this fact in the following lemma:

Lemma 7. $\mathcal{H}(\text{AVR}(\boldsymbol{B}), \boldsymbol{B}) \leqslant \mathcal{H}(F(\boldsymbol{B}), \boldsymbol{B})$ *for every profile $\boldsymbol{B}$ and every generalised dictatorship F.*

We get a very good approximation ratio for the MVR:

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

Proof. W.l.o.g., we may restrict attention to profiles $\boldsymbol{B}$ for which the majority winner closest to the profile is $(0,\ldots,0)$. If there exists a voter $i \in \mathcal{N}$ with $B_i = (0,\ldots,0)$, then the approximation ratio is 1 and we are done. So suppose every voter labels at least one issue with a 1, i.e., $m_i := |\{j \in \mathcal{I} \mid b_{i,j} = 1\}| > 0$ for every $i \in \mathcal{N}$. The distance between the majority winner and the profile is $\sum_{i\in\mathcal{N}} m_i$, i.e., it is equal to the number of 1's occurring anywhere in the profile. Observe that the MVR-winners are those voters labelling the fewest issues with 1. Let $i^* \in \text{argmin}_{i\in\mathcal{N}} m_i$ be one of them. We need to prove the following:

$$\sum_{i\in\mathcal{N}} H(B_{i^*}, B_i) \quad < \quad 2 \cdot \sum_{i\in\mathcal{N}} m_i$$

We are done if $H(B_{i^*}, B_i) \leqslant 2 \cdot m_i$ for all $i \neq i^*$ (strictness of the above inequality then follows from $H(B_{i^*}, B_{i^*}) = 0$ and $0 < m_{i^*}$). Observe that $H(B_{i^*}, B_i) \leqslant m_{i^*} + m_i$: equality holds when the issues labelled 1 by i^* are all distinct from those labelled 1 by i; in all other cases there is more agreement between the two ballots and the distance is strictly less. Hence, as $m_{i^*} \leqslant m_i$ by definition, we are done. $\square$

We are now ready to prove an important positive result, showing that the AVR and the MVR are both very good approximations of the much more complex DBR, independently of the IC used to delimit the set of feasible outcomes.

Theorem 9. *Both the* AVR *and the* MVR *are strict* 2*-approximations of the* DBR^{IC} *for any integrity constraint* IC.

Proof. For the MVR and IC $= \top$, the claim follows from Proposition 8 and Fact 1. As any formula IC logically entails $\top$, by Lemma 2 this generalises to all integrity constraints. By Lemma 7, finally, it also applies to the AVR. □

Is this the best we can do? Yes, at least for IC $= \top$, i.e., for DBR$^\top$, which (by Lemma 2) is the DBR that is hardest to approximate. As the following example shows, neither the AVR nor any other generalised dictatorship can guarantee a better approximation ratio for DBR$^\top$ (i.e., for Maj).

Example 2. *Let $n = m$. Consider a profile $\boldsymbol{B}$ with $b_{i,i} = 1$ and $b_{i,j} = 0$ for $i \neq j$. For this profile the majority outcome is $B_0 = (0, \dots, 0)$. Here is an illustration for $n = 5$:*

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

The distance between B_0 and the profile is $\mathcal{H}(B_0, \boldsymbol{B}) = n$ (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 $\mathcal{H}(B_i, \boldsymbol{B}) = (n-1) + (n-1)$, whichever $i \in \mathcal{N}$ we pick. That is, the approximation ratio for this scenario is $2\frac{n-1}{n}$. Hence, by increasing n (and m), we can go arbitrarily close to 2.

Arguably, this example only applies under conditions that may be deemed too generous. Indeed, it is often reasonable to assume that m is relatively small. Can we do better in this case? Let us first state a limiting result, showing that this restriction does not help as far as the MVR is concerned:

Proposition 10. *Even when m is constant, there exists no $\alpha < 2$ such that the* MVR *is an α-approximation of* Maj*.*

Proof (sketch). Choose n such that $m-1$ divides $n-1$. Consider a profile $\boldsymbol{B}$ in which each voter accepts exactly one issue: the last voter is the only one accepting the last issue, i.e., $b_{n,n} = 1$, and for all other voters $i \leqslant n-1$ let $b_{i,j} = 1$ iff $[i \equiv j \mod (m-1)]$ and $j \leqslant m-1$. The majority winner is $(0, \dots, 0)$. Its distance to $\boldsymbol{B}$ is n (each voter disagrees on one issue). Observe that *every* individual ballot is an MVR-winner. Take the worst MVR-winner B_n. Its distance to $\boldsymbol{B}$ is $(n-1) + (n-1)$. Thus, we get an approximation ratio of $2\frac{n-1}{n}$, i.e., a ratio that moves arbitrarily close to 2 as n grows and that does not depend on m. □

Note that B_n in the above proof is *not* an AVR-winner. In fact, for the AVR we can prove a strong positive result:

Theorem 11. *Let m be constant. Then the* AVR *is an α-approximation of the* DBR$^{\text{IC}}$ *with $\alpha = 2\frac{m-1}{m}$ for any integrity constraint* IC*.*

Proof. We will prove the claim for DBR$^\top$ = Maj; the full theorem then follows from Lemma 2. Recall that AVR-winners achieve the best approximation ratio amongst all individual ballots. So, for the sake of contradiction, assume that there exists a profile $\boldsymbol{B}$ such that *no* ballot achieves the claimed approximation ratio. W.l.o.g., we may assume that $B^{\text{Maj}} := (0, \dots, 0)$ is a majority winner. Thus, for all $i \in \mathcal{N}$:

$$\mathcal{H}(B_i, \boldsymbol{B}) \quad > \quad 2 \cdot \frac{m-1}{m} \cdot \mathcal{H}(B^{\text{Maj}}, \boldsymbol{B})$$

Let $m_i := |\{j \in \mathcal{I} \mid b_{i,j} = 1\}|$ denote the number of issues labelled 1 by voter i. Let $M := \sum_{i \in \mathcal{N}} m_i$ denote the overall number of 1's in the profile. Thus, $\mathcal{H}(B^{\text{Maj}}, \boldsymbol{B}) = M$. Recall that $\mathcal{H}(B_i, \boldsymbol{B}) = \sum_{i' \in \mathcal{N}} H(B_i, B_{i'})$. Now, summing up the n inequalities above (one for each $i \in \mathcal{N}$), we obtain:

$$\sum_{i,i' \in \mathcal{N}} H(B_i, B_{i'}) \quad > \quad 2 \cdot n \cdot \frac{m-1}{m} \cdot M \qquad (1)$$

Let $\sigma_{i,i'}$ be the number of issues labelled 1 by *both* voter i and i'. Thus, $H(B_i, B_{i'}) = m_i + m_{i'} - 2\sigma_{i,i'}$: we can obtain the distance between B_i and $B_{i'}$ by adding the number of 1's in the former to the number of 1's in the latter, and then twice subtracting the overlap to account for over-counting. Now let $n_j := |N^{\boldsymbol{B}}_{j:1}|$ be the number of voters labelling issue j as 1, and note that $\sum_{i,i' \in \mathcal{N}} \sigma_{i,i'} = \sum_{j \in \mathcal{I}} n_j^2$: if we go issue-by-issue, then n_j^2 is the number of pairs of voters that overlap at issue j. As $\sum_{i,i' \in \mathcal{N}} [m_i + m_{i'}] = 2 \cdot n \cdot M$, we get:

$$\sum_{i,i' \in \mathcal{N}} H(B_i, B_{i'}) \quad = \quad 2 \cdot n \cdot M - 2 \cdot \sum_{j \in \mathcal{I}} n_j^2$$

In combination with Equation (1), this yields:

$$\frac{n}{m} \cdot M \quad > \quad \sum_{j \in \mathcal{I}} n_j^2$$

Observe that $\sum_{j \in \mathcal{I}} n_j = M$. From the Cauchy-Schwartz inequality $(\sum_j x_j^2)(\sum_j y_j^2) \geqslant (\sum_j x_j y_j)^2$ with $x_j = 1$ and $y_j = n_j$ we obtain $m \cdot \sum_{j \in \mathcal{I}} n_j^2 \geqslant M^2$. Thus:

$$\frac{n}{m} \cdot M \quad > \quad \sum_{j \in \mathcal{I}} n_j^2 \quad \geqslant \quad \frac{1}{m} \cdot M^2$$

Hence, $n > M$. This means that not everyone of the n voters labels at least one issue as 1. But in that case, there exists a voter i^* with $B_{i^*} = (0, ..., 0)$ achieving an approximation ratio of 1, contradicting our original assumption. □

That is, for small values of m the superiority of the AVR over the MVR in terms of approximation really matters. Lemma 2 suggests that even sharper results are possible for logically strong integrity constraints. We state such a result here without proof (it may be obtained by adapting the proof of Theorem 11 and observing that any conjunct in IC that is a literal simply fixes the choice for the corresponding issue):

Proposition 12. *Let m be constant and let* IC *be a conjunction of k distinct literals. Then the* AVR *is an α-approximation of the* DBR$^{\text{IC}}$ *with $\alpha = 2\frac{m-k-1}{m-k}$.*

4 Additional Properties of AVR and MVR

Our analysis shows that the AVR and the MVR are strong contenders when we are looking for an easy-to-compute rule providing good approximation ratios w.r.t. the ideal defined by the DBR (while the RVR is significantly less attractive). The AVR is the better approximator of the two, although the difference is minor for large m. In this section, we want to briefly compare the two rules in terms of two other criteria: computational complexity and normative desiderata (i.e., *axioms* in the sense of social choice theory).

Both the AVR and the MVR are clearly of low complexity. The MVR is slightly superior to the AVR in this respect:

Fact 13. MVR *winners can be computed in* $O(mn)$.

Fact 14. AVR *winners can be computed in* $O(mn \log n)$.

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.

Yet another way of comparing two rules is in terms of axioms familiar from social choice theory (Gaertner 2006). Simply by virtue of being generalised dictatorships, both the AVR and the MVR satisfy the familiar axioms of *unanimity* and *neutrality* (Grandi and Endriss 2011). Both also are *anonymous*. Next we identify a normatively appealing axiom that is satisfied by the AVR but not the MVR. It 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 $\mathcal{N} = \{1, \ldots, n\}$ and $\mathcal{N}' = \{1, \ldots, n'\}$ each vote on the same set of issues $\mathcal{I}$, resulting in the profiles $\boldsymbol{B}$ and $\boldsymbol{B}'$. Further suppose we use a rule that is well-defined for any number of voters. Then, if a particular ballot wins both under $\boldsymbol{B}$ and $\boldsymbol{B}'$, we should expect it to also win under $\boldsymbol{B} \oplus \boldsymbol{B}' = (B_1, \ldots, B_n, B'_1, \ldots, B'_{n'})$, i.e., when the two electorates vote together in the same election. Let us make this intuitive idea precise:

Definition 6. F *satisfies* ***reinforcement*** *if for any two profiles* $\boldsymbol{B}$ *and* $\boldsymbol{B}'$ *with* $\text{SUPP}(\boldsymbol{B}) = \text{SUPP}(\boldsymbol{B}')$ *and* $F(\boldsymbol{B}) \cap F(\boldsymbol{B}') \neq \emptyset$ *we have* $F(\boldsymbol{B} \oplus \boldsymbol{B}') = F(\boldsymbol{B}) \cap F(\boldsymbol{B}')$.

Reinforcement is 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 15. *The* AVR *satisfies reinforcement.*

Proof. We will make use of the fact that $\mathcal{H}(B, \boldsymbol{B} \oplus \boldsymbol{B}') = \mathcal{H}(B, \boldsymbol{B}) + \mathcal{H}(B, \boldsymbol{B}')$ for any ballot B and any two profiles $\boldsymbol{B}$ and $\boldsymbol{B}'$. Take any two profiles $\boldsymbol{B}$ and $\boldsymbol{B}'$ with $\text{SUPP}(\boldsymbol{B}) = \text{SUPP}(\boldsymbol{B}')$ and $\text{AVR}(\boldsymbol{B}) \cap \text{AVR}(\boldsymbol{B}') \neq \emptyset$. We need to show that $\text{AVR}(\boldsymbol{B} \oplus \boldsymbol{B}') = \text{AVR}(\boldsymbol{B}) \cap \text{AVR}(\boldsymbol{B}')$.

For the first direction, let $B^* \in \text{AVR}(\boldsymbol{B}) \cap \text{AVR}(\boldsymbol{B}')$. As B^* minimises both $\mathcal{H}(B, \boldsymbol{B})$ and $\mathcal{H}(B, \boldsymbol{B}')$ amongst all B, it also minimises $\mathcal{H}(B, \boldsymbol{B} \oplus \boldsymbol{B}')$. Thus, $B^* \in \text{AVR}(\boldsymbol{B} \oplus \boldsymbol{B}')$.

For the other direction, let $B^* \in \text{AVR}(\boldsymbol{B} \oplus \boldsymbol{B}')$. For the sake of contradiction, assume $B^* \notin \text{AVR}(\boldsymbol{B}) \cap \text{AVR}(\boldsymbol{B}')$. W.l.o.g., let $B^* \notin \text{AVR}(\boldsymbol{B})$. Take any $B \in \text{AVR}(\boldsymbol{B}) \cap \text{AVR}(\boldsymbol{B}')$. B beat B^* for $\boldsymbol{B}$ and drew with B^* for $\boldsymbol{B}'$. Thus:

$$\mathcal{H}(B, \boldsymbol{B}) < \mathcal{H}(B^*, \boldsymbol{B}) \quad \text{and} \quad \mathcal{H}(B, \boldsymbol{B}') \leqslant \mathcal{H}(B^*, \boldsymbol{B}')$$

But by our initial observation, this yields $\mathcal{H}(B, \boldsymbol{B} \oplus \boldsymbol{B}') < \mathcal{H}(B^*, \boldsymbol{B} \oplus \boldsymbol{B}')$. Thus, B beats B^* for $\boldsymbol{B} \oplus \boldsymbol{B}'$ under the AVR, i.e., B^* does not win (the required contradiction). □

Proposition 16. *The* MVR *violates reinforcement.*

Proof. We construct a counterexample. Consider the following two elections on 5 issues, with 12 voters each:

	1	**2**	**3**	**4**	**5**
$1\times$	1	1	1	1	1
$2\times$	1	1	0	0	0
$3\times$	0	1	1	1	0
$1\times$	0	1	1	0	1
$3\times$	1	0	1	1	0
$2\times$	1	0	1	0	1
Maj:	1	1	1	1	0

	1	**2**	**3**	**4**	**5**
$1\times$	1	1	1	1	1
$2\times$	1	1	0	0	0
$1\times$	0	1	1	1	0
$3\times$	0	1	1	0	1
$2\times$	1	0	1	1	0
$3\times$	1	0	1	0	1
Maj:	1	1	1	0	1

The highlighted ballots are the MVR-winners. The intersection of the two winner sets includes only $(1, 1, 1, 1, 1)$. However, for the union of the two profiles, the distance between the majority outcome $(1, 1, 1, 0, 0)$ 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, i.e., the MVR violates reinforcement. □

Hence, whenever we consider reinforcement an important property, we should prefer the AVR over the MVR.[2]

5 Conclusion and Related Work

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 defined three such rules: the average-voter rule, the majority-vote ruler, and the ranked-voter rule, which combine the idea of a most representative voter with the principles of three important rules in preference aggregation, namely the Kemeny rule, the Slater rule, and Tideman's ranked-pairs rule, respectively. Our focus has been on approximation results, and we have seen that the first two rules are strong contenders when we want to select a winning outcome that is close to the original profile, while the third rule is less attractive in this respect.

Results in the literature have focused on approximating intractable rules such as the minimax solution in approval voting (LeGrand, Markakis, and Mehta 2007; Caragiannis, Kalaitzis, and Markakis 2010), and more importantly the Kemeny rule in preference (a.k.a. rank) aggregation: Dwork et al. (2001) present a 2-approximation; Ailon, Charikar, and Newman (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. Our approximation bounds are proven in a general setting that subsumes preference aggregation, thus our results transfer to that setting. While there are sharper approximation bounds 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. The AVR and the MVR arguably are also easy to explain, which will be relevant for elections involving human voters.

[2] But note that even the AVR does not satisfy a somewhat stronger notion of reinforcement omitting the precondition about the sets of support of the two subelectorates being the same.

Acknowledgements

This paper subsumes previous work presented at the 7th Multidisciplinary Workshop on Advances in Preference Handling (MPREF-2013). The authors acknowledge the support of the programme on Algorithmic Game Theory and Computational Social Choice at the National University of Singapore's Institute for Mathematical Sciences in 2013 and of COST Action IC1205 on Computational Social Choice.

References

Ailon, N.; Charikar, M.; and Newman, A. 2008. Aggregating inconsistent information: Ranking and clustering. *Journal of the ACM* 55(5).

Arrow, K. J.; Sen, A. K.; and Suzumura, K., eds. 2002. *Handbook of Social Choice and Welfare*. North-Holland.

Brandt, F.; Conitzer, V.; and Endriss, U. 2013. Computational social choice. In Weiss, G., ed., *Multiagent Systems*. MIT Press.

Caragiannis, I.; Kalaitzis, D.; and Markakis, E. 2010. Approximation algorithms and mechanism design for minimax approval voting. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI-2010)*.

Chevaleyre, Y.; Endriss, U.; Lang, J.; and Maudet, N. 2007. 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.

Dokow, E., and Holzman, R. 2010. Aggregation of binary evaluations. *Journal of Economic Theory* 145(2):495–511.

Dwork, C.; Kumar, R.; Naor, M.; and Sivakumar, D. 2001. Rank aggregation methods for the web. In *Proceedings of the 10th International Conference on World Wide Web (WWW-2001)*.

Endriss, U.; Grandi, U.; and Porello, D. 2012. Complexity of judgment aggregation. *Journal of Artificial Intelligence Research* 45:481–514.

Gaertner, W. 2006. *A Primer in Social Choice Theory*. LSE Perspectives in Economic Analysis. Oxford University Press.

Grandi, U., and Endriss, U. 2011. Binary aggregation with integrity constraints. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI-2011)*.

Grandi, U., and Endriss, U. 2013. Lifting integrity constraints in binary aggregation. *Artificial Intelligence* 199–200:45–66.

Hemaspaandra, E.; Spakowski, H.; and Vogel, J. 2005. The complexity of Kemeny elections. *Theoretical Computer Science* 349(3):382–391.

Kemeny, J. 1959. Mathematics without numbers. *Daedalus* 88:577–591.

Kenyon-Mathieu, C., and Schudy, W. 2007. How to rank with few errors. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing (STOC-2007)*.

Konieczny, S., and Pino Pérez, R. 2002. Merging information under constraints: A logical framework. *Journal of Logic and Computation* 12(5):773–808.

Konieczny, S., and Pino Pérez, R. 2011. Logic based merging. *Journal of Philosophical Logic* 40(2):239–270.

Lang, J.; Pigozzi, G.; Slavkovik, M.; and van der Torre, L. 2011. Judgment aggregation rules based on minimization. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-XIII)*.

LeGrand, R.; Markakis, E.; and Mehta, A. 2007. Some results on approximating the minimax solution in approval voting. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2007)*.

List, C., and Puppe, C. 2009. Judgment aggregation: A survey. In *Handbook of Rational and Social Choice*. Oxford University Press.

Miller, M. K., and Osherson, D. 2009. Methods for distance-based judgment aggregation. *Social Choice and Welfare* 32(4):575–601.

Slater, P. 1961. Inconsistencies in a schedule of paired comparisons. *Biometrika* 48(3-4):303–312.

Tideman, T. N. 1987. Independence of clones as a criterion for voting rules. *Social Choice and Welfare* 4(3):185–206.

Young, H. P. 1975. Social choice scoring functions. *SIAM*