Evaluating Approval-Based Multiwinner Voting in Terms of Robustness to Noise

Ioannis Caragiannis[1,2] , **Christos Kaklamanis**[1,2] , **Nikos Karanikolas**[1] and **George A. Krimpas**[1]

[1]Department of Computer Engineering and Informatics, University of Patras, 26504 Rion, Greece
[2]Computer Technology Institute "Diophantus", 26504 Rion, Greece

{caragian, kakl, nkaranik, krimpas}@ceid.upatras.gr

Abstract

Approval-based multiwinner voting rules have recently received much attention in the Computational Social Choice literature. Such rules aggregate approval ballots and determine a winning committee of alternatives. To assess effectiveness, we propose to employ new noise models that are specifically tailored for approval votes and committees. These models take as input a ground truth committee and return random approval votes to be thought of as noisy estimates of the ground truth. A minimum robustness requirement for an approval-based multiwinner voting rule is to return the ground truth when applied to profiles with sufficiently many noisy votes. Our results indicate that approval-based multiwinner voting can indeed be robust to reasonable noise. We further refine this finding by presenting a hierarchy of rules in terms of how robust to noise they are.

1 Introduction

Voting has received much attention by the AI community recently, mostly due to its suitability for simple and effective decision making. One popular line of research, that originates from Arrow [1951], has aimed to characterize voting rules in terms of the *social choice axioms* they satisfy. Another approach views voting rules as *estimators*. It assumes that there is an objectively correct choice, a *ground truth*, and votes are noisy estimates of it. Then, the main criterion for evaluating a voting rule is whether it can determine the ground truth as outcome when applied to noisy votes.

A typical scenario in studies that follow the second approach employs a hypothetical *noise model* that uses the ground truth as input and produces random votes. Then, a voting rule is applied on profiles of such random votes and is considered effective if it acts as a *maximum likelihood estimator* [Conitzer and Sandholm, 2005; Young, 1988] or if it has *low sample complexity* [Caragiannis *et al.*, 2016]. As such evaluations are heavily dependent on the specifics of the noise model, relaxed effectiveness requirements, such as the *accuracy in the limit*, sought in broad classes of noise models [Caragiannis *et al.*, 2016] can be more informative.

We restrict our attention to *approval voting*, where ballots are simply sets of alternatives that are approved by the voters [Laslier and Sanver, 2010]. Furthermore, we consider *multiwinner voting rules* [Faliszewski *et al.*, 2017], which determine committees of alternatives as outcomes [Kilgour, 2010; Aziz *et al.*, 2017]. In particular, we focus on approval-based counting choice rules (or, simply, *ABCC rules*), which were defined recently by Lackner and Skowron [2018]. A famous rule in this category is known as multiwinner approval voting (AV). Each alternative gets a point every time it appears in an approval vote and the outcome consists of a fixed number of alternatives with the highest scores.

We consider noise models that are particularly tailored for approval votes and committees. These models use a committee as ground truth and produce random sets of alternatives as votes. We construct broad classes of noise models that share a particular structure, parameterized by *distance metrics* defined over sets of alternatives. In this way, we adapt to approval-based multiwinner voting the approach of Caragiannis *et al.* [2016] for voting rules over rankings.

Figure 1 illustrates our evaluation framework. The noise model is depicted at the left. It takes as input the ground truth committee and its probability distribution over approval votes which is consistent to a distance metric d. Repeated executions of the noise model produce a profile of random approval votes. The ABCC rule (defined using a bivariate function f; see Section 2) is then applied on this profile and returns one or more winning committees. Our requirement for the ABCC rule is to be accurate in the limit (informally, on profiles with infinitely many votes, it must return the ground truth as the unique winning committee), not only for a single noise model, but for all models that belong to a sufficiently broad class. The breadth of this class quantifies the robustness of the ABCC rule to noise.

The details of our framework are presented in Section 2. Our results indicate that it indeed allows for a classification of ABCC rules in terms of their robustness to noise. In particular, we identify (in Section 3) the modal committee rule (MC) as the ultimately robust ABCC rule: MC is robust against all kinds of reasonable noise. AV follows in terms of robustness and seems to outperform other known ABCC rules (see Section 4). In contrast, the well-known approval Chamberlin-Courant (CC) rule is the least robust. On the other hand, all ABCC rules are robust if we restrict noise sufficiently (see
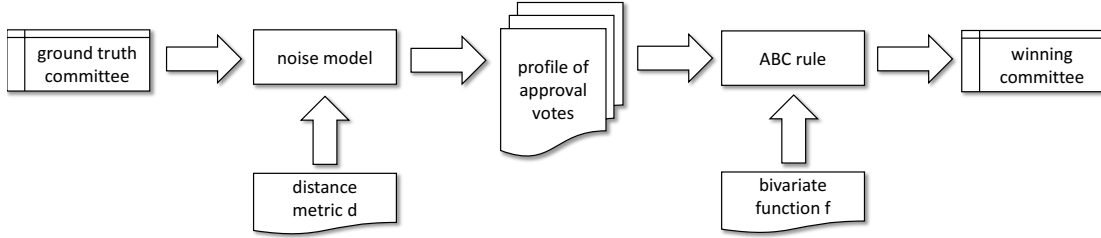
Figure 1: Our evaluation framework.

Section 5). We conclude with a discussion on open problems in Section 6.

Further related work. Approval-based multiwinner voting rules have been studied in terms of their computational complexity [Aziz *et al.*, 2015; Skowron *et al.*, 2016], axiomatic properties [Sánchez-Fernández *et al.*, 2017; Lackner and Skowron, 2018; Aziz *et al.*, 2017], as well as their applications [Brill *et al.*, 2017]. In particular, axiomatic work has focused on two different principles that govern multiwinner rules: diversity and individual excellence. Lackner and Skowron [2019] attempt a quantification of how close an approval-based multiwinner voting rule is to these two principles. We remark that the primary focus of the current paper is on individual excellence.

The robustness of approval voting has been previously evaluated against noise models, using either the MLE [Procaccia and Shah, 2015] or the sample complexity [Caragiannis and Micha, 2017] approach. These papers assume a ranking of the alternatives as ground truth, generate approval votes that consist of the top alternatives in rankings produced according to the noise model of Mallows [1957], and assess how well approval voting recovers the ground truth ranking. We believe that our framework is fairer to approval votes, as recovering an underlying ranking when voters have very limited power to rank is very demanding. The robustness of (non-approval) multiwinner voting against noise has been studied by Procaccia *et al.* [2012]. Different notions of robustness in multiwinner voting are considered by Gawron and Faliszewski [2019] and Bredereck *et al.* [2017].

2 Preliminaries

Throughout the paper, we denote by A the set of alternatives. We use $m = |A|$ and denote the committee size by k. The term committee refers to a set of exactly k alternatives.

Approval-based multiwinner voting. An approval vote is simply a subset of the alternatives (of any size). An approval-based multiwinner voting rule takes as input a profile of approval votes and returns one or more winning committees.

We particularly consider voting rules that belong to the class of approval-based counting choice rules (or, simply, ABCC rules), introduced by Lackner and Skowron [2018]. Such a rule is defined by a bivariate function f, with $f(x, y)$ indicating a non-negative score a committee gets from an approval vote containing y alternatives, x of which are common with the committee. f is non-decreasing in its first argument.

Formally, f is defined on the set $\mathcal{X}_{m,k}$, which consists of all pairs (x, y) of possible values of $|U \cap S|$ and $|S|$, given that U is k-sized and S can be any subset of the m alternatives of A. I.e., $\mathcal{X}_{m,k}$ is the set

$$\{(x, y) : y = 0, 1, ..., m, x = \max\{k + y - m, 0\}, ..., y\}.$$

The score of a committee is simply the total score it gets from all approval votes in a profile. Winning committees are those that have maximum score. We extensively use "the ABCC rule f" to refer to the ABCC rule that uses the bivariate function f. We denote the score that an ABCC rule f assigns to the committee U given a profile $\Pi = (S_i)_{i \in [n]}$ of n votes by $\mathrm{sc}_f(U, \Pi) = \sum_{i=1}^{n} f(|U \cap S_i|, |S|)$. With some abuse of notation, we use $\mathrm{sc}_f(U, S_i)$ to refer to the score U gets from vote S_i. Hence, $\mathrm{sc}_f(U, \Pi) = \sum_{i=1}^{n} \mathrm{sc}_f(U, S_i)$.

Well-known ABCC rules include:

- Multiwinner approval voting (AV), which uses the function $f_{\mathrm{AV}}(x, y) = x$.

- Approval Chamberlin-Courant (CC), which uses the function $f_{\mathrm{CC}}(x, y) = \min\{1, x\}$. The rule falls within a more general context considered by Chamberlin and Courant [1983].

- Proportional approval voting (PAV), which uses the function $f_{\mathrm{PAV}}(x, y) = \sum_{i=1}^{x} 1/i$.

These rules belong to the class of rules that originate from the work of Thiele [1895]. A Thiele rule uses a vector $\langle w_1, w_2, ..., w_k \rangle$ of non-negative weights to define $f(x, y) = \sum_{i=1}^{x} w_j$. Other known Thiele rules include the p-Geometric rule [Skowron *et al.*, 2016] and Sainte Laguë approval voting [Lackner and Skowron, 2019].

A well-known non-Thiele rule is the satisfaction approval voting (SAV) rule that uses $f_{\mathrm{SAV}}(x, y) = x/y$ for $y > 0$ and $f(x, y) = 0$ otherwise [Brams and Kilgour, 2014]. Let us also introduce the *modal committee* (MC) rule[1] which returns the committee (or committees) that has maximum number of appearances as approval votes in the profile. MC is also non-Thiele; it uses $f(k, k) = 1$ and $f(x, y) = 0$ otherwise.

Noise models. We employ noise models to generate approval votes, assuming that the ground truth is a committee. Denoting the ground truth by $U \subseteq A$, a noise model $\mathcal{M}$ produces random approval votes according to a particular distribution that defines the probability $\Pr_{\mathcal{M}}[S|U]$ to generate the set $S \subseteq A$ when the ground truth is U.

---

[1] We remark that MC is similar in spirit to the modal ranking rule considered by Caragiannis *et al.* [2014].

Let us give the following noise model $\mathcal{M}_p$ as an example. $\mathcal{M}_p$ uses a parameter $p \in (1/2, 1]$. Given a ground truth committee U, $\mathcal{M}_p$ generates a random set $S \subseteq A$ by selecting each alternative of U with probability p and each alternative in $A \setminus U$ with probability $1 - p$. Intuitively, the probability that a set will be generated depends on its "distance" from the ground truth: the higher this distance, the smaller this probability. To make this formal, we will need the set difference[2] distance metric $d_\Delta : 2^A \to \mathbb{R}_{\geq 0}$ defined as $d_\Delta(X, Y) = |X \setminus Y| + |Y \setminus X|$.

Claim 1. *For $S \subseteq A$, $\Pr_{\mathcal{M}_p}[S|U] = p^m \left(\frac{1-p}{p} \right)^{d_\Delta(U,S)}$.*

So, the probability $\Pr_{\mathcal{M}_p}[S|U]$ is decreasing in $d_\Delta(U, S)$. We will consider general noise models $\mathcal{M}$ with $\Pr_{\mathcal{M}}[S|U]$ depending on $d(U, S)$, where d is a distance metric defined over subsets of A.

Definition 1. *Let d be a distance metric over sets of alternatives. A noise model $\mathcal{M}$ is called d-monotonic if for any two sets $S_1, S_2 \subseteq A$, it holds $\Pr_{\mathcal{M}}[S_1|U] > \Pr_{\mathcal{M}}[S_2|U]$ if and only if $d(U, S_1) < d(U, S_2)$.*

Definition 1 implies that $\Pr_{\mathcal{M}}[S_1|U] = \Pr_{\mathcal{M}}[S_2|U]$ when $d(U, S_1) = d(U, S_2)$.

Besides the set difference metric used by $\mathcal{M}_p$, other well-known distance metrics[3] (see Deza and Deza [2016]) are:

- the normalized set difference or Jaccard metric d_J, defined as $d_J(X, Y) = \frac{d_\Delta(X,Y)}{|X \cup Y|}$,

- the maximum difference or Zelinka metric d_Z, defined as $d_Z(X, Y) = \max\{|X \setminus Y|, |Y \setminus X|\}$, and

- the normalized maximum difference or Bunke-Shearer metric d_{BS}, defined as $d_{BS}(X, Y) = \frac{d_Z(X,Y)}{\max\{|X|,|Y|\}}$.

Evaluating ABCC rules against noise models. We aim to evaluate the effectiveness of ABCC rules when applied to random profiles generated by large classes of noise models. To this end, we use *accuracy in the limit* as a measure.

Definition 2 (accuracy in the limit)**.** *An ABCC rule f is called* accuracy in the limit *for a noise model $\mathcal{M}$ if for every $\varepsilon > 0$ there exists n_ε such that, for every profile with at least n_ε approval votes produced by $\mathcal{M}$ with ground truth U, f returns U as the unique winning committee with probability at least $1 - \varepsilon$.*

Then, ABCC rules are evaluated in terms of robustness using the next definition.

Definition 3 (robustness)**.** *Let d be a distance metric over sets of alternatives. An ABCC rule f is monotone robust against d (or d-monotone robust) if it is accurate in the limit for all d-monotonic noise models.*

---

[2]Viewing sets as strings, the distance metric d_Δ is equivalent to the Hamming distance; see [Deza and Deza, 2016].

[3]Notice that $d(X, Y)$ for the four specific distance metrics defined here depends only on $|X \setminus Y|$, $|Y \setminus X|$, $|X|$, and $|Y|$. In a sense, these distance metrics are *alternative-independent*. Our results apply to the most general definition of distance, where $d(X, Y)$ can also depend on the contents of $X \setminus Y$, $Y \setminus X$, X, and Y.

3 MC is a Uniquely Robust ABCC Rule

We begin our technical exposition by identifying the unique ABCC rule that is monotone robust against *all* distance metrics. Our proofs, in the current and subsequent sections, make extensive use of the following lemma. The notation $S \sim \mathcal{M}(U)$ indicates that the random set S is drawn from the noise model $\mathcal{M}$ with ground truth U.

Lemma 2. *Let $\mathcal{M}$ be a noise model. An ABCC rule f is*

a. *accurate in the limit for $\mathcal{M}$ if $\mathbb{E}_{S \sim \mathcal{M}(U)}[\mathrm{sc}_f(U, S) - \mathrm{sc}_f(V, S)] > 0$ for every two different sets of alternatives $U, V \subseteq A$ with $|U| = |V| = k$.*

b. *not accurate in the limit for $\mathcal{M}$ if $\mathbb{E}_{S \sim \mathcal{M}(U)}[\mathrm{sc}_f(U, S) - \mathrm{sc}_f(V, S)] < 0$ for some pair of sets of alternatives $U, V \subseteq A$ with $|U| = |V| = k$.*

Due to lack of space, the proof of Lemma 2 is omitted. We are ready to present our first application of Lemma 2.

Theorem 3. *MC is the only ABCC rule that is monotone robust against any distance metric.*

Proof. Let $\mathcal{M}$ be a noise model that is d-monotonic for some distance metric d. Let $U, V \subseteq A$ be any two different k-sized sets of alternatives. By the definition of MC, we have

$$\mathbb{E}_{S \sim \mathcal{M}(U)}[\mathrm{sc}_{\mathrm{MC}}(U, S) - \mathrm{sc}_{\mathrm{MC}}(V, S)]$$
$$= \Pr_{\mathcal{M}}[U|U] - \Pr_{\mathcal{M}}[V|U] > 0.$$

By Lemma 2a, we obtain that MC is d-monotone robust.

We will now show that MC is the only ABCC rule f that has this property. Let f be an ABCC rule that is different than MC. This means that there exist integers x^* and y^* with $(x^* - 1, y^*), (x^*, y^*) \in \mathcal{X}_{m,k}$, $(x^*, y^*) \neq (k, k)$, and $f(x^*, y^*) > f(x^* - 1, y^*)$. We will construct a distance metric d and a d-monotonic noise model for which f is not accurate in the limit.

Rename the alternatives of A as $a_1, a_2, ..., a_m$ and let $U = \{a_1, a_2, ..., a_k\}$, $V = \{a_2, ..., a_{k+1}\}$, and $W = \{a_{k-x^*+2}, ..., a_{y^*+k-x+1}\}$. Notice that, by the definition of $\mathcal{X}_{m,k}$, $(x^* - 1, y^*) \in \mathcal{X}_{m,k}$ implies that $1 + \max\{y^* + k - m, 0\} \leq x^*$ and, equivalently, $y^* + k - x^* + 1 \leq m$; hence, the set W is well-defined. Clearly, $x^* \geq 1$; so sets V and W share at least one alternative.

We define a distance metric between subsets of A that has $d(X, Y) = 0$ if $X = Y$, $d(X, Y) \in \{1, 2\}$, otherwise, and in particular $d(U, V) = d(U, W) = 1$ and $d(U, S) = 2$ for every S different than U, V, or W.

We are ready to define the d-monotonic noise model $\mathcal{M}$. For simplicity, we use $p_0 = \Pr_{\mathcal{M}}[U|U]$, $p_1 = \Pr_{\mathcal{M}}[V|U] = \Pr_{\mathcal{M}}[W|U]$, and $p_2 = \Pr_{\mathcal{M}}[S|U]$ for every other set $S \subseteq A$ different than U, V, or W. For $\delta > 0$ (to be specified shortly), we set $p_0 = 1/3$, $p_1 = 1/3 - \delta$, and $p_2 = \frac{2\delta}{2^m - 3}$.

We now compute the quantity $\mathbb{E}_{S \sim \mathcal{M}(U)}[\mathrm{sc}_f(U, S) - \mathrm{sc}_f(V, S)]$; observe that $\mathrm{sc}_f(U, U) = \mathrm{sc}_f(V, V) = f(k, k)$, $\mathrm{sc}_f(U, V) = \mathrm{sc}_f(V, U) = f(k - 1, k)$, $\mathrm{sc}_f(U, W) = f(x^* - 1, y^*)$, and $f(V, W) = f(x^*, y^*)$. We obtain

$$\mathbb{E}_{S \sim \mathcal{M}(U)}[\mathrm{sc}_f(U, S) - \mathrm{sc}_f(V, S)]$$
$$= f(k, k)p_0 + f(k - 1, k)p_1 + f(x^* - 1, y^*)p_1$$

$$+ \sum_{S \neq U,V,W} f(|U \cap S|, |S|)p_2 - f(k-1, k)p_0$$

$$- f(k,k)p_1 - f(x^*, y^*)p1 - \sum_{S \neq U,V,W} f(|V \cap S|, |S|)p_2$$

$$\leq (p_0 - p_1)(f(k,k) - f(k-1,k))$$

$$- p_1(f(x^*, y^*) - f(x^* - 1, y^*))$$

$$- p_2 \sum_{S \neq U,V,W} f(|V \cap S|, |S|)$$

$$= \delta(f(k,k) - f(k-1,k))$$

$$- (1/3 - \delta)(f(x^*, y^*) - f(x^* - 1, y^*))$$

$$+ \frac{2\delta}{2^m - 3} \sum_{S \neq U,V,W} f(|V \cap S|, |S|). \tag{1}$$

Observe that the RHS of (1) is increasing in δ and approaches $-\frac{1}{3}(f(x^*, y^*) - f(x^* - 1, y^*)) < 0$ as δ approaches 0. Hence, for a sufficiently small positive δ, we have

$$\mathbb{E}_{S \sim \mathcal{M}(U)}[\mathrm{sc}_f(U, S) - \mathrm{sc}_f(V, S)] < 0.$$

By Lemma 2b, f is not accurate in the limit for $\mathcal{M}$. □

4 A Characterization for AV

In this section, we identify the class of distance metrics against which AV is monotone robust. We will need some additional notation that will be useful in several proofs.

For a distance metric d and a set of alternatives U, let $\mathrm{span}(d, U)$ be the number of different non-zero values the quantity $d(U, \cdot)$ can take. We denote these different distance values by $\delta_1(d, U)$, $\delta_2(d, U)$, ..., $\delta_{\mathrm{span}(d,U)}(d, U)$. We also use $\delta_0(d, U) = 0$. For $t = 0, 1, ..., \mathrm{span}(d, U)$ and alternatives $a, b \in A$, we denote by $N_{a|b}^t(d, U)$ the class of sets of alternatives S that contain alternative a but not alternative b and satisfy $d(U, S) \leq \delta_t(d, U)$.

Definition 4 (majority-concentricity). *A distance metric d is called majority-concentric[4] if for every k-sized set of alternatives U, it holds $N_{a|b}^t(d, U) \geq N_{b|a}^t(d, U)$ for every alternatives $a \in U$ and $b \notin U$ and $t = 0, 1, ..., \mathrm{span}(d, U)$.*

We are ready to prove our characterization for AV.

Theorem 4. *AV is d-monotone robust if and only if the distance metric d is majority-concentric.*

Proof. Let $\mathcal{M}$ be a d-monotonic noise model for a majority concentric distance metric d. Let U and V be two different sets with k alternatives each. By Lemma 2a, in order to show that AV is accurate in the limit for $\mathcal{M}$ (and, consequently, d-monotone robust), it suffices to show that $\mathbb{E}_{S \sim \mathcal{M}(U)}[\mathrm{sc}_{\mathrm{AV}}(U, S) - \mathrm{sc}_{\mathrm{AV}}(V, S)] > 0$.

We will need some additional notation. For $t = 0, 1, ..., \mathrm{span}(d, U)$, we denote by $\bar{N}^t(d, U)$) the class of sets of alternatives S that satisfy $d(U, S) = \delta_t(d, U)$. For alternatives $a, b \in A$, we denote $\bar{N}_a^t(d, U)$) the subclass of

---

$\bar{N}^t(d, U)$ consisting of sets of alternatives that include a and by $\bar{N}_{a|b}^t(d, U)$ the subclass of $\bar{N}_a^t(d, U)$ consisting of sets do not contain alternative b.

To simplify notation, we set $s = \mathrm{span}(d, U)$. Also, we drop (d, U) (e.g., we use $N_{a|b}^t$ instead of $N_{a|b}^t(d, U)$) from notation since it is clear from context. We have

$$\mathbb{E}_{S \sim \mathcal{M}(U)}[\mathrm{sc}_{\mathrm{AV}}(U, S)]$$

$$= \sum_{S \subseteq A} \mathrm{sc}_{\mathrm{AV}}(U, S) \cdot \mathrm{Pr}_{\mathcal{M}}[S|U] = \sum_{S \subseteq A} |U \cap S| \cdot \mathrm{Pr}_{\mathcal{M}}[S|U]$$

$$= \sum_{a \in U} \sum_{S \subseteq A : a \in S} \mathrm{Pr}_{\mathcal{M}}[S|U] = \sum_{a \in U} \sum_{t=0}^{s} \sum_{S \in \bar{N}_a^t} \mathrm{Pr}_{\mathcal{M}}[S|U]. \tag{2}$$

Now, observe that the probability $\mathrm{Pr}_{\mathcal{M}}[S|U]$ is the same for all sets $S \in \bar{N}^t$. In the following, we use $p_t = \mathrm{Pr}_{\mathcal{M}}[S|U]$ for all $S \in \bar{N}^t$, for $t = 0, 1, ..., s$. Hence, (2) becomes

$$\mathbb{E}_{S \sim \mathcal{M}(U)}[\mathrm{sc}_{\mathrm{AV}}(U, S)] = \sum_{a \in U} \sum_{t=0}^{s} |\bar{N}_a^t| \cdot p_t$$

Similarly, we have

$$\mathbb{E}_{S \sim \mathcal{M}(U)}[\mathrm{sc}_{\mathrm{AV}}(V, S)] = \sum_{a \in V} \sum_{t=0}^{s} |\bar{N}_a^t| \cdot p_t,$$

and, by linearity of expectation,

$$\mathbb{E}_{S \sim \mathcal{M}(U)}[\mathrm{sc}_{\mathrm{AV}}(U, S) - \mathrm{sc}_{\mathrm{AV}}(V, S)]$$

$$= \sum_{a \in U \setminus V} \sum_{t=0}^{s} |\bar{N}_a^t| \cdot p_t - \sum_{a \in V \setminus U} \sum_{t=0}^{s} |\bar{N}_a^t| \cdot p_t. \tag{3}$$

Let μ be a bijection that maps each alternative of $V \setminus U$ to a distinct alternative of $U \setminus V$. Then, (3) becomes

$$\mathbb{E}_{S \sim \mathcal{M}(U)}[\mathrm{sc}_{\mathrm{AV}}(U, S) - \mathrm{sc}_{\mathrm{AV}}(V, S)]$$

$$= \sum_{a \in U \setminus V} \sum_{t=0}^{s} |\bar{N}_a^t| \cdot p_t - \sum_{a \in U \setminus V} \sum_{t=0}^{s} |\bar{N}_{\mu(a)}^t| \cdot p_t$$

$$= \sum_{a \in U \setminus V} \sum_{t=0}^{s} \left(|\bar{N}_{a|\mu(a)}^t| - |\bar{N}_{\mu(a)|a}^t| \right) \cdot p_t$$

$$= \sum_{a \in U \setminus V} \left(|N_{a|\mu(a)}^0| - |N_{\mu(a)|a}^0| \right) \cdot p_0$$

$$+ \sum_{a \in U \setminus V} \sum_{t=1}^{s} \left(|N_{a|\mu(a)}^t| - |N_{a|\mu(a)}^{t-1}| - |N_{\mu(a)|a}^t| \right.$$

$$\left. + |N_{\mu(a)|a}^{t-1}| \right) \cdot p_t$$

$$= \sum_{a \in U \setminus V} \sum_{t=0}^{s-1} \left(|N_{a|\mu(a)}^t| - |N_{\mu(a)|a}^t| \right) \cdot (p_t - p_{t+1})$$

$$+ \left(|N_{a|\mu(a)}^s| - |N_{\mu(a)|a}^s| \right) \cdot p_s \tag{4}$$

$$\geq \sum_{a \in U \setminus V} \left(|N_{a|\mu(a)}^0| - |N_{\mu(a)|a}^0| \right) \cdot (p_0 - p_1) > 0.$$

---

[4]Majority-concentricity is similar in spirit with a property of distance metrics over rankings with the same name in [Caragiannis *et al.*, 2016].

The third equality follows since $\bar{N}^0_{a|\mu(a)} = N^0_{a|\mu(a)}$, $\bar{N}^0_{\mu(a)|a} = N^0_{\mu(a)|a}$, and $\bar{N}^t_{a|\mu(a)} = N^t_{a|\mu(a)} \setminus N^{t-1}_{a|\mu(a)}$ and $\bar{N}^t_{\mu(a)|a} = N^t_{\mu(a)|a} \setminus N^{t-1}_{\mu(a)|a}$ for $t = 1, ..., s$. The first inequality follows since d is majority concentric and since $p_t > p_{t+1}$ and, thus, all differences in (4) are non-negative. The last inequality follows after observing that since $|N^0_{a|\mu(a)}| = 1$ and $|N^0_{\mu(a)|a}| = 0$ for $a \in U \setminus V$ and since $p_0 > p_1$. This completes the "if" part of the proof.

Let us now consider a non-majority concentric distance metric d that satisfies $N^{t^*}_{a|b}(d, U) < N^{t^*}_{b|a}(d, U)$ for the k-sized set of alternatives U, some alternatives $a \in U$ and $b \notin U$, and some $t^* \in \{1, 2, ..., \text{span}(d, U)\}$. We show the "only if" part of the theorem by constructing a noise model $\mathcal{M}$ that satisfies $\mathbb{E}_{S \sim \mathcal{M}(U)}[\text{sc}_{\text{AV}}(U, S) - \text{sc}_{\text{AV}}(V, S)] < 0$ for $V = U \setminus \{a\} \cup \{b\}$.

Again, we use $p_t = \text{Pr}_{\mathcal{M}}[S|U]$ for every set of alternatives $S \in \bar{N}^t(d, U)$, $s = \text{span}(d, U)$, and drop (d, U) from notation. We define the model probabilities so that $\tau = p_0 > p_1 > ... > p_{t^*} = \tau - \epsilon$ and $2\epsilon = p_{t^*+1} > ... > p_s = \epsilon$. Notice that such a noise model exists for any arbitrarily small $\epsilon > 0$. Since there are 2^m sets of alternatives and τ is the probability that $\mathcal{M}$ returns the ground truth ranking, it must be $\tau > 1/2^m$. We now apply equality (4). Observe that, since $V = U \setminus \{a\} \cup \{b\}$, $\mu(a) = b$. We obtain

$$\mathbb{E}_{S \sim \mathcal{M}(U)}[\text{sc}_{\text{AV}}(U, S) - \text{sc}_{\text{AV}}(V, S)]$$

$$= \sum_{t=0}^{s-1} \left(|N^t_{a|b}| - |N^t_{b|a}| \right) \cdot (p_t - p_{t+1})$$
$$+ \left(|N^s_{a|b}| - |N^s_{b|a}| \right) \cdot p_s$$

$$= \sum_{t=0}^{t^*-1} \left(|N^t_{a|b}| - |N^t_{b|a}| \right) \cdot (p_t - p_{t+1})$$
$$+ \left(|N^{t^*}_{a|b}| - |N^{t^*}_{b|a}| \right) \cdot (p_{t^*} - p_{t^*+1})$$
$$+ \sum_{t=t^*+1}^{s-1} \left(|N^t_{a|b}| - |N^t_{b|a}| \right) \cdot (p_t - p_{t+1})$$
$$+ \left(|N^s_{a|b}| - |N^s_{b|a}| \right) \cdot p_s.$$

Now, observe that for $t \neq t^*$, it holds $|N^t_{a|b}| - |N^t_{b|a}| \leq 2^m$ (the total number of sets of alternatives) and $p_t - p_{t+1} \leq \epsilon$. Also, $|N^{t^*}_{a|b}| - |N^{t^*}_{b|a}| \leq -1$ and $p_{t^*} - p_{t^*+1} = \tau - 3\epsilon$. Setting specifically $\epsilon = \frac{1}{s8^m}$, we obtain that

$$\mathbb{E}_{S \sim \mathcal{M}(U)}[\text{sc}_{\text{AV}}(U, S) - \text{sc}_{\text{AV}}(V, S)]$$
$$\leq s2^m \epsilon - (\tau - 3\epsilon) \leq \frac{1}{4^m} - \frac{1}{2^m} + \frac{3}{s \cdot 8^m},$$

which is negative for $m \geq 2$ since $s \geq 1$. The proof of the "only if" part of the theorem now follows by Lemma 2b. $\square$

It is tempting to conjecture that AV and MC are the only ABCC rules that are monotone robust against all majority concentric distance metrics. However, this is not true as the next example, which uses a different ABCC rule, shows.

Example 1. *Let $A = \{a, b, c\}$ and $k = 2$. Consider the majority concentric distance metric d and the ABCC rule f with $f(1, 1) = 1$, $f(2, 2) = 2$, and $f(x, y) = 0$ otherwise. We will show that f is d-monotone robust against any majority concentric distance metric d. Without loss of generality, let us assume that $U = \{a, b\}$ and $V = \{a, c\}$. Observe that the quantity $\text{sc}_f(U, S) - \text{sc}_f(V, S)$ is equal to 0 when $S = \emptyset, \{a\}, \{b, c\}, \{a, b, c\}$, 1 when $S = \{b\}$, -1 when $S = \{c\}$, 2 when $S = \{a, b\}$, and -2 when $S = \{a, c\}$. Hence, for the d-monotonic noise model $\mathcal{M}$, we have $\mathbb{E}_{S \sim \mathcal{M}(U)}[\text{sc}_f(U, S) - \text{sc}_f(V, S)] = 2p_{ab} - 2p_{ac} + p_b - p_c$, where p_{ab}, p_{ac}, p_b, and p_c are abbreviations for the probabilities $\text{Pr}_{\mathcal{M}}[S|U]$ for $S = \{a, b\}$, $\{a, c\}$, $\{b\}$, and $\{c\}$, respectively.*

In order to have $N^t_{b|c} \geq N^t_{c|b}$ for $t = 0, 1, ..., \text{span}(d, U)$ as the definition of majority concentricity requires, it must be either $p_{ab} > p_b, p_c \geq p_{ac}$ or $p_{ab} > p_b, p_{ac} \geq p_c$. In the first case, we have $\mathbb{E}_{S \sim \mathcal{M}(U)}[\text{sc}_f(U, S) - \text{sc}_f(V, S)] = (p_{ab} - p_{ac}) + (p_{ab} - p_c) + (p_b - p_{ac}) > 0$. In the second case, we have $\mathbb{E}_{S \sim \mathcal{M}(U)}[\text{sc}_f(U, S) - \text{sc}_f(V, S)] = 2(p_{ab} - p_{ac}) + (p_b - p_c) > 0$. Accuracy in the limit of the ABCC rule f for the noise model $\mathcal{M}$ then follows by Lemma 2a.

5 Robustness of Other ABCC Rules

Our results for ABCC rules different than MC and AV involve two classes of distance metrics. We define the first one here.

Definition 5 (natural distance). *A distance metric d is called natural if for every three sets U, V, and S with $|U| = |V|$ such that $|U \cap S| > |V \cap S|$, it holds that $d(U, S) \leq d(V, S)$.*

The next observation follows easily by the definitions.

Claim 5. *Any natural distance metric is majority-concentric.*

The opposite is not true as the next example illustrates.

Example 2. *Let $A = \{a, b, c\}$ and consider the distance metric with $d(X, Y) = 0$ for every pair of sets with $X = Y$, $d(X, Y) = 1$ if $X \cap Y = \emptyset$ and $X \cup Y = A$, and $d(X, Y) = 2$, otherwise. It can be easily seen that the distance is majority-concentric; it suffices to observe that, within distance 1 from any set, each alternative appears in exactly one set. To see that is not natural, consider $U = \{a, b\}$, $V = \{a, c\}$ and $S = \{b\}$. We have $|U \cap S| > |V \cap S|$ but $d(U, S) = 2 > 1 = d(V, S)$.*

Lemma 7 below identifies the class of ABCC rules that are monotone robust against all natural distance metrics. The condition uses an appropriately defined bijection on sets of alternatives.

Definition 6. *Given two different sets U and V with $|U| = |V|$, a (U, V)-bijection $\mu : 2^A \rightarrow 2^A$ is defined as $\mu(S) = \{\mu'(a) : a \in S\}$, where $\mu' : A \rightarrow A$ is such that $\mu'(a) = a$ for every alternative $a \in U \cap V$ or $a \notin U \cup V$, $\mu'(a)$ is a distinct alternative in $V \setminus U$ for $a \in U \setminus V$, and $\mu'(a)$ is a distinct alternative in $U \setminus V$ for $a \in V \setminus U$.*

It is easy to see that a (U, V)-bijection μ has the following properties.

Claim 6. *Let $U, V \subseteq A$ with $|U| = |V|$ and let μ be a (U, V)-bijection. For every $S \subseteq A$, it holds $|S| = |\mu(S)|$, $|U \cap S| = |V \cap \mu(S)|$, and $|U \cap \mu(S)| = |V \cap S|$.*

Lemma 7. *An ABCC rule is d-monotone robust against a natural distance metric d if for every two different sets of alternatives $U, V \subseteq A$ with $|U| = |V| = k$ there exists a (U, V)-bijection μ on sets of alternatives and a set $S \subseteq A$ with $\mathrm{sc}_f(U, S) > \mathrm{sc}_f(V, S)$ and $d(U, S) < d(U, \mu(S))$.*

Proof. Let U and V be two different sets with k alternatives each. Let $\mathcal{S}_+$, $\mathcal{S}_-$, and $\mathcal{S}_0$ be the classes of sets of alternatives S with $|U \cap S| > |V \cap S|$, $|U \cap S| < |V \cap S|$, and $|U \cap S| = |V \cap S|$, respectively. Using this notation, we have

$$\mathbb{E}_{S \sim \mathcal{M}(U)}[\mathrm{sc}_f(U, S) - \mathrm{sc}_f(V, S)]$$
$$= \sum_{S \subseteq A} (\mathrm{sc}_f(U, S) - \mathrm{sc}_f(V, S)) \cdot \mathrm{Pr}_{\mathcal{M}}[S|U]$$
$$= \sum_{S \in \mathcal{S}_+} (\mathrm{sc}_f(U, S) - \mathrm{sc}_f(V, S)) \cdot \mathrm{Pr}_{\mathcal{M}}[S|U]$$
$$+ \sum_{S \in \mathcal{S}_0} (\mathrm{sc}_f(U, S) - \mathrm{sc}_f(V, S)) \cdot \mathrm{Pr}_{\mathcal{M}}[S|U]$$
$$+ \sum_{S \in \mathcal{S}_-} (\mathrm{sc}_f(U, S) - \mathrm{sc}_f(V, S)) \cdot \mathrm{Pr}_{\mathcal{M}}[S|U] \quad (5)$$

We will now transform the third sum in the RHS of (5) to one running over the sets of $\mathcal{S}_+$ like the first sum.

Let μ be a (U, V)-bijection on sets of alternatives; by Claim 6, μ maps every set of $\mathcal{S}_-$ to a set of $\mathcal{S}_+$ and vice-versa. Hence, instead of enumerating sets of $\mathcal{S}_-$, we could enumerate sets of $\mathcal{S}_+$ and apply the bijection μ on them. The third sum in the RHS of (5) then becomes

$$\sum_{S \in \mathcal{S}_-} (\mathrm{sc}_f(U, S) - \mathrm{sc}_f(V, S)) \cdot \mathrm{Pr}_{\mathcal{M}}[S|U]$$
$$= \sum_{S \in \mathcal{S}_+} (\mathrm{sc}_f(U, \mu(S)) - \mathrm{sc}_f(V, \mu(S))) \cdot \mathrm{Pr}_{\mathcal{M}}[\mu(S)|U]$$
$$= \sum_{S \in \mathcal{S}_+} (\mathrm{sc}_f(V, S) - \mathrm{sc}_f(U, S)) \cdot \mathrm{Pr}_{\mathcal{M}}[\mu(S)|U] \quad (6)$$

The second equality follows since, by Claim 6, $\mathrm{sc}_f(U, \mu(S)) = f(|U \cap \mu(S)|, |\mu(S)|) = \mathrm{sc}_f(|V \cap S|, |S|) = \mathrm{sc}_f(V, S)$ and, similarly, $\mathrm{sc}_f(V, \mu(S)) = \mathrm{sc}_f(U, S)$.

Now observe that $\mathrm{sc}_f(U, S) = \mathrm{sc}_f(V, S)$ when $S \in \mathcal{S}_0$. Hence, the second sum in the RHS of (5) is equal to 0. By combining (5) and (6), we get

$$\mathbb{E}_{S \sim \mathcal{M}(U)}[\mathrm{sc}_f(U, S) - \mathrm{sc}_f(V, S)]$$
$$= \sum_{S \in \mathcal{S}_+} (\mathrm{sc}_f(U, S) - \mathrm{sc}_f(V, S))$$
$$\cdot (\mathrm{Pr}_{\mathcal{M}}[S|U] - \mathrm{Pr}_{\mathcal{M}}[\mu(S)|U]). \quad (7)$$

Now observe that the RHS of (7) is always non-negative. This is due to the fact that $S \in \mathcal{S}_+$ which implies that $\mathrm{sc}_f(U, S) = f(|U \cap S|, |S|) \geq f(|V \cap S|, |S|) = \mathrm{sc}_f(V, S)$ since f is non-decreasing in its first argument and $d(U, S) \leq d(U, \mu(S))$ (and, consequently, $\mathrm{Pr}_{\mathcal{M}}[S|U] \geq \mathrm{Pr}_{\mathcal{M}}[\mu(S)|U]$) since d is natural and, by Claim 6, $|U \cap S| > |V \cap S| = |U \cap \mu(S)|$. The RHS of (7) is strictly positive if there exists a set $S \in \mathcal{S}_+$ that further satisfies $d(U, S) < d(U, \mu(S))$ (and, consequently, $\mathrm{Pr}_{\mathcal{M}}[S|U] > \mathrm{Pr}_{\mathcal{M}}[\mu(S)|U]$) and $\mathrm{sc}_f(U, S) > \mathrm{sc}_f(V, S)$. The lemma then follows by Lemma 2a. □

We now present two applications of Lemma 7.

Theorem 8. *An ABCC rule f is monotone robust against any natural distance metric if and only if $f(k, k) > f(k - 1, k)$.*

Notice that most popular ABCC rules from Section 2 satisfy the condition of Theorem 8. CC is an exception. The omitted proof of Theorem 8 implies that CC is not monotone robust for the natural distance metric d defined as $d(X, Y) = 0$ if $X = Y$ and $d(X, Y) = 1$, otherwise.

Our second application of Lemma 7 involves all non-trivial ABCC rules and an important subclass of natural distances.

Definition 7 (similarity distance). *A natural distance metric d is a similarity distance metric if for every three sets U, V, and S with $|U| = |V|$ such that $|U \cap S| > |V \cap S|$, it holds that $d(U, S) < d(V, S)$.*

Theorem 9. *Any non-trivial ABCC rule is monotone robust against any similarity distance metric.*

Proof. We apply Lemma 7 assuming a non-trivial ABCC rule f and a similarity distance metric d. Non-triviality of f implies that for every two different sets U and V with k alternatives each, there is a set S such that $\mathrm{sc}_f(U, S) > \mathrm{sc}_f(V, S)$. This yields $|U \cap S| > |V \cap S| = |U \cap \mu(S)|$, where μ is any (U, V)-bijection (see Claim 6), and implies that $d(U, S) < d(U, \mu(S))$ since d is a similarity distance. □

We can easily show that the four distance metrics set difference, Jaccard, Zelinka, and Bunke-Shearer that we defined in Section 2 are all similarity distance metrics. Using this observation and Theorem 9, we obtain the next statement.

Corollary 10. *Any non-trivial ABCC rule is monotone robust against the set difference, Jaccard, Zelinka, and Bunke-Shearer distance metrics.*

6 Epilogue

We believe that our approach complements nicely the axiomatic and quantitative analysis of approval-based multiwinner voting. The current paper leaves many interesting open problems. Besides identifying ABCC rules that are at least as robust as AV, applying our framework to non-ABCC rules deserves investigation. Beyond assessing the effects of noise in the limit, studying the sample complexity of approval-based multiwinner voting is important. This will require the design of concrete noise models like the $\mathcal{M}_p$ model that we presented in Section 2. In particular, models that simulate user behaviour in crowdsourcing platforms will be useful for evaluating approval-based voting in such environments. Even though the $\mathcal{M}_p$ model is very simple, we expect that implementation issues will emerge for more elaborate noise models. Similar issues in the implementation of the Mallows [1957] ranking model have triggered much non-trivial work; see, e.g., [Doignon *et al.*, 2004].

Acknowledgments

References

[Arrow, 1951] K. J. Arrow. *Social Choice and Individual Values*. Cowles Foundation, 1951.

[Aziz *et al.*, 2015] H. Aziz, S. Gaspers, J. Gudmundsson, S. Mackenzie, N. Mattei, and T. Walsh. Computational aspects of multi-winner approval voting. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 107–115, 2015.

[Aziz *et al.*, 2017] H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 48(2):461–485, 2017.

[Brams and Kilgour, 2014] S. J. Brams and D. M. Kilgour. Satisfaction approval voting. In R. Fara, D. Leech, and M. Salles, editors, *Voting Power and Procedures: Essays in Honour of Dan Felsenthal and Moshé Machover*, pages 323–346. Springer, 2014.

[Bredereck *et al.*, 2017] R. Bredereck, P. Faliszewski, A. Kaczmarczyk, R. Niedermeier, P. Skowron, and N. Talmon. Robustness among multiwinner voting rules. In *Proceedings of the 10th International Symposium on Algorithmic Game Theory, (SAGT)*, pages 80–92, 2017.

[Brill *et al.*, 2017] M. Brill, J.-F. Laslier, and P. Skowron. Multiwinner approval rules as apportionment methods. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, pages 414–420, 2017.

[Caragiannis and Micha, 2017] I. Caragiannis and E. Micha. Learning a ground truth ranking using noisy approval votes. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 149–155, 2017.

[Caragiannis *et al.*, 2014] I. Caragiannis, A. D. Procaccia, and N. Shah. Modal ranking: A uniquely robust voting rule. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence*, pages 616–622, 2014.

[Caragiannis *et al.*, 2016] I. Caragiannis, A. D. Procaccia, and N. Shah. When do noisy votes reveal the truth? *ACM Transactions on Economics and Computation*, 4(3):15:1–15:30, 2016.

[Chamberlin and Courant, 1983] J. R. Chamberlin and P. N. Courant. Representative deliberations and representative decisions: Proportional representation and the borda rule. *American Political Science Review*, 77(3):718–733, 1983.

[Conitzer and Sandholm, 2005] V. Conitzer and T. Sandholm. Common voting rules as maximum likelihood estimators. In *Proceedings of the 21st Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 145–152, 2005.

[Deza and Deza, 2016] M. M. Deza and E. Deza. *Encyclopedia of Distances*. Springer, 4th edition, 2016.

[Doignon *et al.*, 2004] J.-P. Doignon, A. Pakeč, and M. Regenwetter. The repeated insertion model for rankings: Missing link between two subset choice models. *Psychometrika*, 69(1):33–54, 2004.

[Faliszewski *et al.*, 2017] P. Faliszewski, P. Skowron, A. M. Slinko, and N. Talmon. Multiwinner voting: A new challenge for social choice theory. In U. Endriss, editor, *Trends in Computational Social Choice*, pages 27–47. AI Access, 2017.

[Gawron and Faliszewski, 2019] G. Gawron and P. Faliszewski. Robustness of approval-based multiwinner voting rules. In *Proceedings of the 6th International Conference on Algorithmic Decision Theory (ADT)*, pages 17–31, 2019.

[Kilgour, 2010] D. M. Kilgour. Approval balloting for multiwinner elections. In J.-F. Laslier and M. R. Sanver, editors, *Handbook on Approval Voting*, pages 105–124. Springer, 2010.

[Lackner and Skowron, 2018] M. Lackner and P. Skowron. Consistent approval-based multi-winner rules. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, pages 47–48, 2018.

[Lackner and Skowron, 2019] M. Lackner and P. Skowron. A quantitative analysis of multi-winner rules. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 407–413, 2019.

[Laslier and Sanver, 2010] J.-F. Laslier and M. R. Sanver, editors. *Handbook on Approval Voting*. Springer, 2010.

[Mallows, 1957] C. L. Mallows. Non-null ranking models. *Biometrika*, 44:114–130, 1957.

[Procaccia and Shah, 2015] A. D. Procaccia and N. Shah. Is approval voting optimal given approval votes? In *Proceedings of the 29th Annual Conference on Neural Information Processing Systems (NIPS)*, pages 1801–1809, 2015.

[Procaccia *et al.*, 2012] A. D. Procaccia, S. J. Reddi, and N. Shah. A maximum likelihood approach for selecting sets of alternatives. In *Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 695–704, 2012.

[Sánchez-Fernández *et al.*, 2017] L. Sánchez-Fernández, E. Elkind, M. Lackner, N. Fernández García, J. Arias-Fisteus, P. Basanta-Val, and P. Skowron. Proportional justified representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, pages 670–676, 2017.

[Skowron *et al.*, 2016] P. Skowron, P. Faliszewski, and J. Lang. Finding a collective set of items: From proportional multirepresentation to group recommendation. *Artificial Intelligence*, 241:191—216, 2016.

[Thiele, 1895] T. N. Thiele. On multiple choice (in Danish). *Bulletin of the Royal Danish Academy of Sciences and Letters*, pages 415–441, 1895.

[Young, 1988] H. P. Young. Condorcet's theory of voting. *American Political Science Review*, 82(4):1231–1244, 1988.