ECAI 2024
U. Endriss et al. (Eds.)

doi:10.3233/FAIA240870

Multiwinner Temporal Voting with Aversion to Change

Valentin Zech[a,*], Niclas Boehmer[b], Edith Elkind[a,c] and Nicholas Teh[a]

[a]University of Oxford, UK
[b]Hasso Plattner Institute, University of Potsdam, Germany
[c]Alan Turing Institute, UK

Abstract. We study two-stage committee elections where voters have dynamic preferences over candidates; at each stage, a committee is chosen under a given voting rule. We are interested in identifying a winning committee for the second stage that overlaps as much as possible with the first-stage committee. We show a full complexity dichotomy for the class of Thiele rules: this problem is tractable for Approval Voting (AV) and hard for all other Thiele rules (including, in particular, Proportional Approval Voting and the Chamberlin–Courant rule). We extend this dichotomy to the greedy variants of Thiele rules. We also explore this problem from a parameterized complexity perspective for several natural parameters. We complement the theory with experimental analysis: e.g., we investigate the average number of changes in the committee as a function of changes in voters' preferences and the role of ties.

1 Introduction

A local town council advisory committee is responsible for making decisions on various community issues such as education, infrastructure, and public services. Elections are held biennially to fill the positions on this committee. Numerous residents, each with their own platform and priorities for the town's development, step forward as candidates for the election to this advisory committee. Voters from different neighborhoods and demographics then go to the polls to elect members of this committee, to make decisions on their behalf.

Between election cycles, due to varying campaign performances and evolving community concerns, some voters change their preferences over the candidates. While the voting rule to be used to select the advisory committee is fixed by the bylaws, it often results in multiple tied committees, i.e., it does not fully determine the election outcome. There are many ways to break these ties; in particular, one may want to maintain contiguity by prioritizing committees that have a substantial overlap with the previous committee, so as to build on the existing expertise and maintain stability, while remaining representative of the population's preferences.

This problem can be viewed through the lens of *multiwinner temporal voting*, a natural temporal extension of the well-studied multiwinner voting model (see the taxonomy of Boehmer and Niedermeier [2] and a survey by Elkind et al. [19]). In multiwinner temporal elections, a set of voters have dynamic preferences over a set of candidates, and we want to elect a committee at each timestep. In this work, we seek to study, given a voting rule, how winning committees adapt to changes in the voters' preferences in situations where it is undesirable to replace many committee members at once.

* Corresponding Author. Email: zech@vzech.de.

Our Contributions We consider two-stage elections that use a fixed voting rule to select a winning committee at each stage; the voters have approval preferences that may evolve between the stages. When the first-stage committee is elected, the second-stage preferences are not yet known. Hence, a winning committee is chosen arbitrarily from the committees tied for winning. In the second stage, the goal is to identify a committee that wins under the new preferences, yet has as much overlap as possible with the first-stage committee.

We consider this problem for the well-known class of Thiele rules [32] (this class includes, in particular, Approval Voting (AV) [4], Proportional Approval Voting (PAV) [41], and the Chamberlin–Courant rule (CCAV) [11]) and their greedy variants. In Section 3, we present a full complexity classification of our decision problem for Thiele rules. In particular, we obtain a dichotomy: the decision version of our problem is in P for Approval Voting (AV) (Section 3.1), and coNP-hard for all Thiele rules other than AV (Section 3.2 and Section 3.3). We then extend this dichotomy to all *greedy* Thiele rules (Section 4); surprisingly, our problem remains hard for all rules other than AV, even though under greedy Thiele rules, computing a single winning committee is computationally tractable. To complement our hardness results, we provide parameterized complexity results for a selection of natural parameters and their combinations (Section 5). Finally, we use experiments to obtain quantitative results (Section 6): we measure the amount of change in the committees as a function of change in voters' preferences, and investigate the role of ties. In particular, our experiments show that simply breaking ties lexicographically is far from optimal with respect to the contiguity objective, thereby justifying our theoretical analysis.

All missing proofs and additional experimental results can be found in the full version of our paper [44].

Related Work Our analysis extends the line of work initiated by Bredereck et al. [8], but our model and contributions differ from theirs in several key aspects: Bredereck et al. [8] study a model where (i) committees are not selected from the set of winning committees output by a voting rule, but need to satisfy a lower bound on their score (via a *single non-transferable vote* committee scoring rule), (ii) the subsequent committees need not be of the same size, and (iii) voters preferences in all stages are known at the start. In contrast, we consider a large class of popular voting rules and require each selected committee to win under a given rule for a fixed committee size; moreover, our model is more realistic in that it does not assume that the second-stage preferences are known initially.

Other papers in this vein include the work of Bredereck et al. [6], that considered maximizing the changes made to a committee (the "revolutionary" setting) to find diverse committees, and the paper by

Deltl et al. [15], which looks into treating agents fairly.

Our work is also related to the study of *robustness* in temporal voting [7]. Bredereck et al. [7] conducted an axiomatic study of several voting rules, focusing on the worst-case change that may have to be made to a winning committee after a single change was made to some agent's preference. Conversely, for a given election and a voting rule, they ask for the minimum number of changes to the agents' preferences so that the election outcome changes (in any way) under the given rule. This framework was subsequently adapted and studied in approval-based elections [24], for greedy approval-based rules [21], and for nearly-structured preferences [34].

Our work is different in that we go beyond this worst-case measure and conduct a more fine-grained analysis, where we ask *by how much* a winning committee needs to change in a given altered election and study the related computational problems. Moreover, Bredereck et al. [7] and Faliszewski et al. [21] assume a fixed tie-breaking order of candidates, whereas we study rules under parallel-universe tie-breaking. Further, while most of the prior work focuses on a few easy-to-compute rules, in our work, we focus on the important class of Thiele rules and their greedy variants.

There are also models of temporal voting where a single winner is chosen at each timestep. These works generally consider temporal extensions of popular multiwinner voting rules and study various axiomatic properties [30, 31], welfare measures [18, 35], or extensions of the *justified representation* axioms and its variants [9, 16, 12, 20]. Other works look into the special case of fair scheduling, where preferences and the outcome are permutations of candidates [17].

We note that similar problems have also been studied (albeit under different names) in the context of stable matching [5], coloring [27], clustering [13, 33], and reoptimization [10, 39].

2 Preliminaries

Given a logical expression φ, we use the Iverson bracket notation $[\varphi]$ to denote the evaluation of that expression: $[\varphi] = 1$ if φ is true and $[\varphi] = 0$ otherwise. We assume familiarity with the basics of classic complexity theory [37] and parameterized complexity [22, 36].

Elections Let $C = \{c_1, \dots, c_m\}$ be a set of m candidates, and let $N = \{1, \dots, n\}$ be a set of n voters. Each voter $i \in N$ has *approval preferences* over candidates in C, captured by a *ballot* $v_i \subseteq C$; we require that $v_i \neq \emptyset$. Let $V = (v_1, \dots, v_n)$. We refer to the pair (C, V) as an *election*.

Rules A *multiwinner voting rule* $\mathcal{R}$ maps an election $E = (C, V)$ and an integer $k \in [|C|]$ to a non-empty family of sets $\mathcal{R}(E, k)$, where each set in $\mathcal{R}(E, k)$ is a size-k subset of C. The sets in $\mathcal{R}(E, k)$ are called the *winning committees* for E under $\mathcal{R}$.

In this work, we focus on a well-studied class of voting rules known as Thiele rules [29, 32], and their greedy variants. We say that a function $\lambda : \mathbb{N}^+ \to [0, 1]$ is an *Ordered Weighted Averaging (OWA) function* if $\lambda(1) = 1$ and λ is non-increasing, i.e. for all $i, j \in \mathbb{N}^+$ it holds that $i > j$ implies $\lambda(i) \leq \lambda(j)$ [21]. In what follows, we assume that all OWA functions λ we consider take values in $\mathbb{Q} \cap [0, 1]$ and are polynomial-time computable. Each OWA function λ defines a *Thiele rule* $\mathcal{R}_\lambda$ as follows. Given an election $E = (C, V)$ and a committee size $k \in [|C|]$, the λ*-score* of a candidate set $S \subseteq C$ in E is defined by: $\lambda\text{-score}^E(S) = \sum_{i \in N} \left(\sum_{j=1}^{|S \cap v_i|} \lambda(j) \right)$. The rule $\mathcal{R}_\lambda$ then outputs all size-k subsets of C with the maximum λ-score: $\mathcal{R}_\lambda(E, k) = \operatorname{argmax}\{\lambda\text{-score}(S) \mid S \in \binom{C}{k}\}$. We omit λ and/or E whenever it is clear from the context. This framework captures many well-known voting rules. Specifically,

- $\lambda(i) = 1$ for all $i \in \mathbb{N}^+$ corresponds to the Approval Voting (AV) rule [4];
- $\lambda(i) = 1/i$ for all $i \in \mathbb{N}^+$ corresponds to the Proportional Approval Voting (PAV) rule [41]; and
- $\lambda(i) = [i = 1]$ for all $i \in \mathbb{N}^+$ corresponds to the Chamberlin–Courant Approval Voting (CCAV) rule [11].

The AV rule is appropriate when the aim is to select a group of individually excellent candidates, whereas CCAV aims to represent as many voters as possible; the PAV rule provides strong proportional representation guarantees [32]. We say that an OWA function λ (and the associated Thiele rule $\mathcal{R}_\lambda$) is *unit-decreasing* if $\lambda(1) = 1 > \lambda(2)$. AV is not unit-decreasing, whereas PAV and CCAV are. Intuitively, unit-decreasing rules capture that the voters' marginal utility from having an additional representative in the winning committee is lower than their utility from being represented at all.

Since committees under both PAV and CCAV are NP-hard to compute [1, 38, 42], there exist greedy approximation variants of these rules, which are based on the notion of *marginal contributions*. Given an election $E = (C, V)$, a subset $S \subseteq C$ of candidates and an OWA function λ, we define the *marginal contribution* (or *points*) of a candidate $c \in C \setminus S$ with respect to S as $\lambda\text{-score}(S \cup \{c\}) - \lambda\text{-score}(S)$. Then, the greedy variant of a Thiele rule $\mathcal{R}_\lambda$, which we denote by Greedy-$\mathcal{R}_\lambda$, outputs all committees that can be obtained by the following iterative procedure: start with an empty committee S and perform k iterations; at each iteration, add a candidate in $C \setminus S$ with maximum marginal contribution to S under λ (at each iteration there may be multiple candidates with maximum marginal contribution to the current committee; a committee is in the output of the rule if it can be obtained by some way of breaking these ties at each iteration). Clearly, a committee in the output of Greedy-$\mathcal{R}_\lambda$ can be computed in polynomial time (recall that we assume that λ itself is polynomial-time computable). We note that Greedy-AV is equivalent to AV; however, other Thiele rules differ from their greedy variants.

Distances Given two committees $S, S' \subseteq C$ of equal size k, we define the *distance* between S and S' as $\operatorname{dist}(S, S') = k - |S \cap S'|$.

To this end, we define elementary ADD and REMOVE operations on elections. An ADD operation adds a previously unapproved candidate to the ballot of a single voter. A REMOVE operation removes a single candidate from a single voter's ballot. Then, the distance between two elections E and E', denoted $\operatorname{Dist}(E, E')$, is defined as the length of the shortest sequence of ADD and REMOVE operations that transforms E into E' (and $+\infty$ if E cannot be transformed into E' using these two operations).

Decision Problem We are now ready to present the family of decision problems we are interested in.

RESILIENT COMMITTEE ELECTIONS ($\mathcal{R}$-RCE):
Input: Elections $E = (C, V)$ and $E' = (C, V')$ over the same set of candidates C, a committee size $k \in \mathbb{N}$, a winning committee $S \in \mathcal{R}(E, k)$, and a distance bound $\ell \in \mathbb{N}$.
Question: Does there exist a committee $S' \in \mathcal{R}(E', k)$ such that $\operatorname{dist}(S, S') \leq \ell$?

3 A Dichotomy for Thiele Rules

We present a full complexity classification of RCE for Thiele rules. RCE is tractable for AV and hard for all other Thiele rules. To show this, we proceed in three steps: first (Section 3.1) we present a polynomial-time algorithm for AV, then (Section 3.2) we give a hardness proof for all unit-decreasing Thiele rules, and finally (Section 3.3) we extend it to all Thiele rules other than AV. Our hardness result for unit-decreasing Thiele rules also establishes that this problem is coW[1]-hard with respect to the committee size k.

3.1 Tractability for Approval Voting

We first observe that RCE is easy for Approval Voting.

Proposition 3.1. *AV-*RCE *admits a polynomial-time algorithm.*

Proof. Consider two elections $E = (C, V)$ and $E' = (C, V')$ over a candidate set $C = \{c_1, \dots, c_m\}$. Let $k \in \mathbb{N}$ be the committee size and let $S \in \mathcal{R}(E, k)$ be a winning committee for E. For each $c \in C$, let $s(c)$ be the approval score of c in E'. Without loss of generality, assume that $s(c_1) \geq \dots \geq s(c_m)$. Then, we partition C into three disjoint sets: $C_{\text{above}} = \{c_j \mid s(c_j) > s(c_k)\}$, $C_{\text{equal}} = \{c_j \mid s(c_j) = s(c_k)\}$, and $C_{\text{below}} = \{c_j \mid s(c_j) < s(c_k)\}$. A winning committee under AV in E' must include all of the candidates in C_{above} and $k - |C_{\text{above}}|$ candidates from C_{equal}; by construction, $0 < k - |C_{\text{above}}| \leq |C_{\text{equal}}|$. Thus, we construct a committee S^* by first including all candidates from C_{above} as well as $\min\{k - |C_{\text{above}}|, |C_{\text{equal}} \cap S|\}$ candidates from $C_{\text{equal}} \cap S$. We then fill the committee with arbitrary candidates from C_{equal}. Obviously, the resulting committee S^* wins in election E' under AV. Furthermore, our approach ensures that S^* contains as many candidates from S as possible. Finally, we check if $\operatorname{dist}(S, S^*) \leq \ell$. As all steps are computable in polynomial time, this concludes the proof. □

3.2 Hardness for Unit-Decreasing Thiele Rules

We will now present our hardness result for unit-decreasing Thiele rules. Our proof also shows parameterized hardness with respect to the committee size, and applies even to the case where E and E' differ only in a single approval.

Theorem 3.2. *For every unit-decreasing Thiele rule $\mathcal{R}$ and every fixed value of $\ell \in [k-1]$ the problem $\mathcal{R}$-*RCE *is* coNP*-hard and* coW[1]*-hard when parameterized by the committee size k, even if every voter approves at most two candidates, and even if $\operatorname{Dist}(E, E') = 1$.*

Proof. We reduce from Independent Set (IS). An instance of IS is a pair (G, κ), where $G = (\mathcal{V}, \mathcal{E})$ is an arbitrary graph and $\kappa \in \mathbb{N}$ is a non-negative integer. It is a yes-instance if there is an independent set of size κ in G, and a no-instance otherwise. IS is NP-hard and W[1]-hard when parameterized by the solution size κ. Given an instance (G, κ) of IS where $G = (\mathcal{V}, \mathcal{E})$ and $|\mathcal{V}| = \nu$, we construct an instance of $\mathcal{R}$-RCE as follows.

We set committee size k to κ. In election E, the set of candidates C is defined as $C_{\mathcal{V}} \cup D$, where $C_{\mathcal{V}} = \{c_w \mid w \in \mathcal{V}\}$ is the set of *vertex candidates*, and D is a set of k *dummy candidates*. Let $\alpha = \lambda(2)$, where λ is the underlying OWA function of $\mathcal{R}$, and let $t = \left\lceil \frac{2}{1-\alpha} \right\rceil$.

We introduce the following four voter groups in election E.

1. For every edge $\{u, w\} \in \mathcal{E}$, there are t *edge voters* who approve the vertex candidates c_u and c_w.
2. For every vertex $w \in \mathcal{V}$, there are $(\nu - \deg(w)) \cdot t$ voters who approve the vertex candidate c_w.
3. For every pair of candidates $d \in D$ and $c_w \in C_{\mathcal{V}}$, there are t voters who approve both d and c_w.
4. For every dummy candidate $d \in D$, there are $k \cdot t$ voters who approve d.

Note that every candidate is approved by exactly $(\nu + k) \cdot t$ voters.

Since no two dummy candidates are approved by the same voter, the size-k committee D has the maximum possible score of $k \cdot (\nu + k) \cdot t$ in election E. Therefore, $D \in \mathcal{R}(E, k)$.

Now, consider an election $E' = (C, V')$, where V' is obtained by picking an arbitrary candidate $d^* \in D$ and an arbitrary voter i who approves d^* and some vertex candidate, and removing d^* from i's ballot. Then the score of the committee D in E' is $k \cdot (\nu + k) \cdot t - 1$. We will show that D wins in election E' if and only if there is no size-κ independent set in graph G and that otherwise, every winning committee $S \in \mathcal{R}(E', k)$ is entirely disjoint from D, i.e., $D \cap S = \emptyset$. Note that this generalizes the statement to all values of $\ell \in [k-1]$.

($\Rightarrow$) Assume that there is an independent set I of size κ in graph G, and let $S_I = \{c_w \mid w \in I\}$ be the size-k committee that corresponds to the vertices in I. Since I is an independent set, no two candidates in S_I are approved by the same voter, so the score of S_I in E' is $k \cdot (\nu + k) \cdot t$. Moreover, $S_I \cap D = \emptyset$. Assume for contradiction that there is a committee $S \in \mathcal{R}(E', k)$ such that $S \cap D \neq \emptyset$, i.e., there is a candidate $d \in S \cap D$. Since $\operatorname{score}(D, E') = k \cdot (\nu + k) \cdot t - 1$, it must hold that $S \neq D$, i.e., there is a candidate $c \in S \cap C_{\mathcal{V}}$. However, then there are t voters who approve both d and c. Therefore, the score of S is at most $k \cdot (\nu + k) \cdot t - (1 - \alpha) \cdot t$. Now, since $\mathcal{R}$ is a unit-decreasing Thiele rule, we have $\alpha < 1$ and, hence, $\operatorname{score}(S_I, E') > \operatorname{score}(S, E')$. Thus, $S \notin \mathcal{R}(E', k)$, a contradiction.

($\Leftarrow$) Assume that there is no independent set I of size κ in G. Assume for contradiction that there is a size-k committee $S \subseteq C$ such that $\operatorname{score}(S, E') > \operatorname{score}(D, E')$, i.e., $D \notin \mathcal{R}(E', k)$. Then $S \cap C_{\mathcal{V}} \neq \emptyset$; let c be some candidate in $S \cap C_{\mathcal{V}}$. Now, if $S \cap D \neq \emptyset$, there exists a candidate $d \in S$, and t voters who approve both c and d. Similarly, if $S \cap D = \emptyset$, we have $S \subseteq C_{\mathcal{V}}$. Since S does not correspond to an independent set, there are at least t edge voters who approve two candidates in S. Therefore, in either case, the score of S is at most $k \cdot (\nu + k) \cdot t - (1 - \alpha) \cdot t$. But then $t = \left\lceil \frac{2}{1-\alpha} \right\rceil$ implies that the quantity $\operatorname{score}(D, E') - \operatorname{score}(S, E') = (k \cdot (\nu + k) \cdot t - 1) - (k \cdot (\nu + k) \cdot t - (1 - \alpha) \cdot t) = (1 - \alpha) \cdot t - 1 = (1 - \alpha) \cdot \left\lceil \frac{2}{1-\alpha} \right\rceil - 1 \geq 1$, i.e., D has a strictly higher score than S in E', a contradiction.

Note that the committee size k in our constructed election is equal to the solution size κ of the given IS instance. Therefore, our reduction is parameter-preserving. □

Note that we do not claim that $\mathcal{R}$-RCE is coNP-complete; because in the naive guess-and-check approach, one would first guess a committee S' with a low enough distance from the original committee S, which indicates to the class NP. Then to verify that the chosen committee S' wins in the altered election, one would guess a second committee S'' and check if S'' has a higher score than S' in the altered election, which indicates to the class coNP.

3.3 Beyond Unit-Decreasing Thiele Rules

Consider a Thiele rule $\mathcal{R}_\lambda$ that is not unit-decreasing. This means that $\lambda(1) = \lambda(2) = 1$. Then either $\lambda(i) = \lambda(j)$ for all $i, j \in \mathbb{N}$ (i.e., $\mathcal{R}_\lambda$ is the Approval Voting rule) or there exists an $s > 1$ such that $\lambda(j) = 1$ for all $j \leq s$ and $\lambda(s+1) < 1$.

In the latter case, we can modify the proof of Theorem 3.2 by (1) adding a set of $s-1$ candidates F that are approved by all voters, and (2) increasing the committee size by $s - 1$. Then in both E and E', every committee with the maximum score would contain F. Moreover, $F \cup D$ is optimal for $\mathcal{R}_\lambda$ in E, and it remains optimal in E' if and only if the underlying graph does not admit an independent set of size κ. This establishes that $\mathcal{R}_\lambda$-RCE is coNP-hard.

We are now ready to state our dichotomy result.

Theorem 3.3. *Consider a Thiele rule $\mathcal{R}$ associated with the OWA λ. If $\lambda(i) = 1$ for all $i \in \mathbb{N}^+$, then the problem $\mathcal{R}$-RCE is polynomial-time solvable. Otherwise, it is* coNP*-hard and* coW[1]*-hard when parameterized by k. These hardness results hold for all fixed $\ell \in [k-s]$ where $s \in \mathbb{N}^+$ is the smallest number such that $\lambda(s+1) < 1$, and even if $\mathrm{Dist}(E, E') = 1$.*

4 A Dichotomy for Greedy Thiele Rules

In this section, we focus on greedy Thiele rules, and establish a dichotomy result that is similar to Theorem 3.3: if Greedy-$\mathcal{R}_\lambda$ is a greedy Thiele rule, Greedy-$\mathcal{R}_\lambda$-RCE is NP-hard unless $\lambda(i) = 1$ for all $i \in \mathbb{N}^+$ (i.e., unless Greedy-$\mathcal{R}_\lambda$ is AV). This is despite greedy Thiele rules having better computational properties than Thiele rules: e.g., it is easy to find a winning committee under a greedy Thiele rule.

However, our argument becomes much more involved. Again, we start by establishing a hardness result for unit-decreasing rules. Our hardness reduction for this class of rules proceeds in two steps. We first define a new problem, which we call CANDIDATE INCLUSION (CI) and show it to be NP-hard for greedy unit-decreasing Thiele rules. We then give a reduction from CI to RCE.

For a fixed voting rule $\mathcal{R}$, an instance of $\mathcal{R}$-CI comprises of an election $E = (C, V)$, a committee size $k \in \mathbb{N}$ and a set of candidates $P \subseteq C$; it is a yes-instance if there exists a winning committee $S \in \mathcal{R}(E, k)$ such that $P \subseteq S$, and a no-instance otherwise. This problem can be seen as a generalization of the WINNER CHECKING (WC) problem studied by Aziz et al. [1], i.e., the task of checking whether a given committee is among the winners in a given election.

We provide a high-level idea of our hardness proof of CI, and defer the full proof (which is quite technical) to the full version [44].

Proposition 4.1. *For every greedy unit-decreasing Thiele rule Greedy-$\mathcal{R}$ and size of $|P| \in [1, k]$, Greedy-$\mathcal{R}$-*CI *is* NP*-hard.*

Proof idea. Fix a greedy unit-decreasing Thiele rule Greedy-$\mathcal{R}$. We reduce from RESTRICTED EXACT COVER BY THREE SETS (RX3C) [26], which is a variant of EXACT COVER BY THREE SETS [23]. An instance of RX3C comprises of a finite set of elements $\mathcal{U} = \{u_1, \dots, u_{3h}\}$ and a family $\mathcal{M} = \{M_1, \dots, M_{3h}\}$ of size-3 subsets of $\mathcal{U}$ such that every element of $\mathcal{U}$ belongs to exactly three sets in $\mathcal{M}$; it is a yes-instance if there is a selection of exactly h sets from $\mathcal{M}$ whose union is $\mathcal{U}$, and a no-instance otherwise.

Given an instance $(\mathcal{U}, \mathcal{M})$ of RX3C, we construct an instance of Greedy-$\mathcal{R}$-CI with an election E, a subset P of candidates, and a committee size k. Our set of candidates contains a *set candidate* for every set in $\mathcal{M}$, as well as three candidates p, d, and x. We set the committee size k to $3h + 2$. We construct voters so that x is chosen in the first iteration, followed by a selection of h set candidates. Then, in the $(h + 1)$-th iteration, we reach the critical point where candidate p *can* be selected if and only if the previously selected set candidates correspond to an exact cover of $\mathcal{U}$, and otherwise, candidate d is selected as a default. In the final $2h$ iterations, all remaining set candidates are selected. The set P consists of candidate p and an arbitrary number of set candidates. Then, there is a winning committee $S \in$ Greedy-$\mathcal{R}(E, k)$ with $P \subseteq S$ if and only if $\mathcal{U}$ can be covered with h sets from $\mathcal{M}$. □

By reducing CI to RCE, we establish the following:

Theorem 4.2. *For every greedy unit-decreasing Thiele rule Greedy-$\mathcal{R}_\lambda$ and for every distance between committees $\ell \in [k-1]$, Greedy-$\mathcal{R}$-*RCE *is* NP*-hard, even if $\mathrm{Dist}(E, E') = 1$.*

Proof idea. We give a reduction from Greedy-$\mathcal{R}$-CI to Greedy-$\mathcal{R}$-RCE. Fix a greedy unit-decreasing Thiele rule Greedy-$\mathcal{R}$. Given an instance $(\tilde{E} = (\tilde{C}, \tilde{V}), \tilde{P}, \tilde{k})$ of Greedy-$\mathcal{R}$-CI, we construct an instance of Greedy-$\mathcal{R}$-RCE as follows. We create two *wrapper elections* E and E' at distance of 1 from each other. In E and E', we include all candidates in $\tilde{C}$ and all voters in $\tilde{V}$, an additional set B of $\tilde{k} - |\tilde{P}|$ candidates, as well as two *control candidates* x and y. We set the committee size k to $\tilde{k} + 2$. We construct voters in election E so that candidate x can be chosen in the first iteration. Once x is selected, all candidates in $\tilde{C} \setminus \tilde{P}$ lose sufficiently many points that they will not be selected in any of the subsequent iterations. Thus, all candidates in $B \cup \tilde{P}$ need to be chosen, before y is chosen in the final iteration. Therefore, the size-k committee $S = B \cup \tilde{P} \cup \{x, y\}$ wins in election E, i.e., $S \in$ Greedy-$\mathcal{R}(E, k)$.

In contrast, we construct the voters in E' so that y must be selected in the first iteration. Once y is selected, all candidates in B lose sufficiently many points that they will not be selected in any of the subsequent iterations. Then, in the following $\tilde{k}$ iterations, candidates from $\tilde{C}$ must be chosen, before x is chosen in the final iteration. Thus, the intersection between S and a winning committee S' in E' can only contain x, y, and candidates in $\tilde{P}$.

After y has been chosen and before x is chosen, Greedy-$\mathcal{R}$ operates on E' in the same way as it would on $\tilde{E}$. This implies that Greedy-$\mathcal{R}(E', k) = \{\tilde{S} \cup \{x, y\} \mid \tilde{S} \in$ Greedy-$\mathcal{R}(\tilde{E}, \tilde{k})\}$, i.e., a committee S' is winning in E' if and only if it consists of candidates x and y and all candidates from a winning committee in $\tilde{E}$. We set the allowed difference ℓ between committees in our RCE instance to $k - |\tilde{P}| - 2$, i.e., at least $|\tilde{P}| + 2$ candidates need to appear in both winning committees. Since x and y will always be chosen, this implies that a selection of at least $|\tilde{P}|$ candidates from $\tilde{S}$ need to be present in both S and S', which are exactly the candidates from $\tilde{P}$. This ensures that the given instance of Greedy-$\mathcal{R}$-CI is a *yes*-instance if and only if our constructed Greedy-$\mathcal{R}$-RCE instance is a *yes*-instance.

The above establishes hardness for all values of $\ell \in [k-3]$. Recall that Greedy-$\mathcal{R}$-CI is NP-hard for every size of $|P| \geq 1$, and we remark that the distance between committees in E and E' can be increased by 2 by setting k to $\tilde{k} + 1$. Then, under Greedy-$\mathcal{R}$, y will not be chosen in the last iteration on E, and x will not be chosen in the last iteration on E'. Thus, one can verify that we obtain hardness for the complete range of $\ell \in [k-1]$. □

Recall that Greedy-AV is equivalent to AV and hence Greedy-AV-RCE is polynomial-time solvable. On the other hand if $\lambda(s) = 1$, $\lambda(s+1) < 1$, we can use the same construction as in the proof of Theorem 3.3, i.e., modify the proof of Theorem 4.2 by increasing the committee size by s and adding s candidates approved by all voters. We obtain the following corollary.

Corollary 4.3. *Consider a greedy Thiele rule Greedy-$\mathcal{R}$ associated with the OWA λ. If $\lambda(i) = 1$ for all $i \in \mathbb{N}^+$, then Greedy-$\mathcal{R}$-*RCE *is polynomial-time solvable. Otherwise, it is* NP*-hard. The result holds for all $\ell \in [k-s]$, where $s \in \mathbb{N}^+$ is the smallest number such that $\lambda(s+1) < 1$, and even if $\mathrm{Dist}(E, E') = 1$.*

Previously, we have motivated the RCE problem with settings in which it is costly to replace any member of an already implemented winning committee. In these settings, we are mostly interested in the computational complexity for small values of the parameter ℓ, i.e., we allow for only very few candidates to be replaced. However, in other scenarios, one might mostly be concerned that there is at least *some* intersection between subsequent winning committees. For instance, imagine a scenario where the board of directors of an organization is elected periodically. When a new board is elected without a candidate who was also part of the previous board, there might not be an adequate handover. As this would likely lead to a great reduction in productivity since the new board members will have to be acquainted with their roles completely independently, one might suggest that at least a few board members should be part of two subsequent boards of directors. We can think of this scenario in the light of a *transition of power*. With regard to the RCE problem, these types of scenarios motivate the study of the computational complexity for small values of $k - \ell > 0$, i.e., situations in which we want at least $k - \ell$ candidates to stay in the committee. However, we see that such a problem is still computationally hard, with the following result.

Theorem 4.4. *For every greedy unit-decreasing Thiele rule Greedy-$\mathcal{R}_\lambda$ and for every fixed value of $k - \ell \in \mathbb{N}^+$, parameterized by the solution size k, Greedy-$\mathcal{R}$-*RCE *is* W[1]*-hard, even if every voter approves at most two candidates, and even if* $\mathrm{Dist}(E, E') = 1$.

5 Parameterized Complexity Results

Next, we consider RCE for Thiele rules and their greedy variants from a parameterized complexity perspective. We present tractability results (FPT and XP) for some parameters, which are not already ruled out by the results in the previous section.

5.1 Thiele Rules

Fix a Thiele rule $\mathcal{R}$. If $\mathcal{R}$ is unit-decreasing, then, according to Theorem 3.2, unless FPT = coW[1], no algorithm can solve an instance $\mathcal{I}$ of $\mathcal{R}$-RCE in $\mathcal{O}\left(f(k) \cdot |I|^{\mathcal{O}(1)}\right)$ time, where k is the committee size and f is some computable function. Thus, we cannot hope for an FPT algorithm with respect to k that works for all Thiele rules. However, $\mathcal{R}$-RCE admits a simple algorithm that is XP with respect to k and FPT with respect to $|C| = m$.

Proposition 5.1. *For every Thiele rule $\mathcal{R}$ the problem $\mathcal{R}$-*RCE *is* FPT *in m and* XP *in k.*

Proof. Given an $\mathcal{R}$-RCE instance (E, E', S, k, ℓ), we can go over all size-k subsets of C, evaluate their scores in E', and check if one of the committees with the maximum score is at distance at most ℓ from S. There are $\binom{m}{k} \leq m^k \leq m^m$ committees to consider; for each, its $\mathcal{R}$-score in E' can be computed in polynomial time. □

Further, our problem is fixed-parameter tractable with respect to the combined parameter $n + k$, where n is the number of voters. This proof, as well as some of the subsequent proofs, is based on the idea that we can partition the candidates in E' into at most 2^n non-empty *candidate classes*, so that all candidates in each class are approved by the same voters.

Proposition 5.2. *For every Thiele rule $\mathcal{R}$ the problem $\mathcal{R}$-*RCE *is* FPT *in $n + k$.*

Proof. Given a Thiele rule $\mathcal{R}$ and an RCE instance (E, E', S, k, ℓ), we construct an election $\tilde{E} = (\tilde{C}, \tilde{V})$ with $|\tilde{C}| \leq k \cdot 2^n$ such that there is a committee $\tilde{S} \in \mathcal{R}(\tilde{E}, k)$ with $\mathrm{dist}(S, \tilde{S}) \leq \ell$ if and only if there is an $S' \in \mathcal{R}(E', k)$ with $\mathrm{dist}(S, S') \leq \ell$. The candidate set $\tilde{C}$ contains all k candidates in S, and, for every candidate class K, an arbitrary selection of at most $k - |K \cap S|$ of candidates from $K \setminus S$. The profile $\tilde{V}$ is then obtained by restricting V to $\tilde{C}$. Since $S \subseteq \tilde{C}$, and all candidates within each class are interchangeable, we can assume without loss of generality that a committee in $\mathcal{R}(E', k)$ that minimizes the distance to S is a subset of $\tilde{C}$. We can therefore use the same approach as in the proof of Proposition 5.1, i.e., go through all size-k subsets of $\tilde{C}$, compute their score in $\tilde{E}$ and distance to S. The bound on the running time follows from the analysis in Proposition 5.1 and the fact that $|\tilde{C}| \leq 2^n \cdot k$. □

Whether this result can be strengthened to a fixed-parameter tractable algorithm with respect to n alone remains an open question. However, we can place our problem in FPT with respect to n for a specific well-studied rule, namely, CCAV.

Proposition 5.3. *CCAV-*RCE *is* FPT *in n.*

5.2 Greedy Thiele Rules

Greedy Thiele rules are less computationally demanding than Thiele rules. As such, all easiness results from Section 5.1 extend to greedy Thiele rules; moreover, some variants of our problem that are hard for Thiele rules (under a suitable complexity assumption) admit FPT algorithms for greedy Thiele rules. For instance, by Theorem 4.4, if $\mathcal{R}$ is a Thiele rule, $\mathcal{R}$-RCE is coW[1]-hard with respect to k even for fixed ℓ. In contrast, our next proof shows that Greedy-$\mathcal{R}$-RCE is FPT in k for any fixed value of ℓ.

Proposition 5.4. *For every greedy Thiele rule Greedy-$\mathcal{R}$ the problem Greedy-$\mathcal{R}$-*RCE *is* FPT *in k for every fixed value of $\ell \in \mathbb{N}$, as well as* FPT *in m and* XP *in k.*

Proof. Given an RCE instance (E, E', S, k, ℓ), we guess a subset of candidates $S^- \subseteq S$, $|S^-| \leq \ell$ to be replaced, and a subset of candidates $S^+ \subseteq C \setminus S$, $|S^+| = |S^-|$, to replace them. Note that for $S' = (S \setminus S^-) \cup S^+$ we have $\mathrm{dist}(S, S') \leq \ell$, and there are at most $\binom{k}{\ell} \cdot \binom{m-k}{\ell}$ pairs (S^-, S^+) to consider (polynomially many for constant ℓ, and at most $2^k \cdot m^k$, as $\ell \leq k$). We then guess a permutation π of S' and check if the rule Greedy-$\mathcal{R}$ can select the candidates in S', in the order specified by π. There are $k!$ permutations to consider, so the bounds on the running time follow. □

Just as for Thiele rules, we combine the approach of Proposition 5.4 with the idea of partitioning candidates into classes to design an algorithm that is FPT in $n + k$.

Proposition 5.5. *For every greedy Thiele rule Greedy-$\mathcal{R}$ the problem Greedy-$\mathcal{R}$-*RCE *is* FPT *in $n + k$.*

For Greedy-CC, we can strengthen this result from FPT in $n + k$ to FPT in n; it remains an open problem if a similar tractability result holds for other greedy Thiele rules.

Proposition 5.6. *Greedy-CC-*RCE *is* FPT *in n.*

6 Experiments

To complement our theoretical analysis, we conduct experiments to gain insights into the practical facets of our problem. Thereby, we also contribute to the small, growing body of experimental work on approval-based elections [3]. Given that computing a winning committee for most Thiele rules is already computationally intractable, we focus on two popular greedy Thiele rules: Greedy-CC and Greedy-PAV.

Our experiments focus on the following three questions. **Q1:** How resilient are winning committees under Greedy-CC and Greedy-PAV, i.e., how much do they need to change when votes do? **Q2:** How good are solutions of RCE obtained by employing lexicographical tie-breaking when computing the original and updated winning committee? **Q3:** Is there a correlation between (i) the round in which the greedy rule included a candidate in the committee and (ii) how often the candidate gets replaced after changes in the votes occur?

In Experiments 1 and 3, we assume *lexicographical tie-breaking* in the computation of Greedy-CC and Greedy-PAV, i.e., we break ties based on some fixed order of candidates whenever multiple candidates have the same marginal contribution. This implies that both rules become resolute, i.e., they return a unique winning committee. In particular, this allows us to solve RCE in polynomial time. The code for our experiments is available online [43].

Experimental Design We consider two different models for generating approval-based preferences, both of which are well-studied in the literature [3, 14, 25, 40].

- In the *1D-Euclidean Model (1D)*, each voter v (resp. candidate c) is assigned (uniformly at random) a point p_v (resp. p_c) in the interval $[0, 1]$. The model is parameterized by a *radius* $\tau \in [0, 1]$. A voter v approves of a candidate c if and only if $|p_v - p_c| \leq r$.
- In the *2D-Euclidean Model (2D)*, each voter and candidate is assigned (uniformly at random) a point in the unit square $[0, 1] \times [0, 1]$. A voter v approves of a candidate c if and only if the Euclidean distance between their points is at most the radius τ.

We focus on elections with $n = 1000$, $m = 100$, and committee size $k = 10$, which is standard in the literature [3]. We set the radius for the 1D (resp. 2D) model to 0.051 (resp. 0.195), so that, on average, every voter approves around 10 candidates. In all three experiments, we sample 100 elections for each of the two models.

We also conduct experiments for a greater range of radii, a sampling method known as *Resampling*, as well as for 1D and 2D models in combination with Resampling. Due to space constraints, we defer the model definitions and detailed trends to our full version [44], and only briefly highlight the differences to the above sampling models.

To capture change in the votes, we consider three different operations: *ADD*, *REMOVE*, and *MIX*. Given a number r of changes to be performed, for *ADD*, we uniformly at random add r new approvals to the elections (i.e., we sample an r-subset of all voter-candidate pairs where the voter does not approve the candidate). For *REMOVE*, we uniformly at random delete r existing approvals from the election, whereas for *MIX*, we add and remove $\lfloor r/2 \rfloor$ approvals each.

We consider different levels of change as determined by a change percentage $p \in [0\%, 10\%]$. For an election E, a change percentage of p corresponds to making $r = \lfloor \text{app}(E) \cdot p \rfloor$ changes, where $\text{app}(E)$ is the total number of approvals in E. We consider 15 change percentages, quadratically scaled, to ensure that smaller amounts of changes are captured in greater detail.

Experiment 1: Resilience of Greedy Thiele Rules We analyze the resilience of winning committees, i.e., how much they change when voters' preferences change, and how this depends on the voting rule and operation type. For this, for each considered level of change p and election E, we perform 100 iterations. For each iteration, we sample an election E' by applying $r = \lfloor \text{app}(E) \cdot p \rfloor$ changes and compute the distance between the winning committee in E and E'.

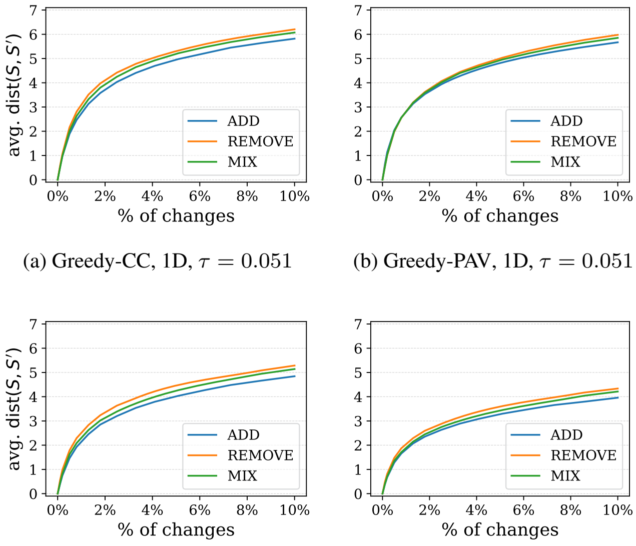

(a) Greedy-CC, 1D, $\tau = 0.051$ (b) Greedy-PAV, 1D, $\tau = 0.051$

(c) Greedy-CC, 2D, $\tau = 0.195$ (d) Greedy-PAV, 2D, $\tau = 0.195$

Figure 1: Results of Experiment 1. x-axis is the percentage change between the original election E and the adapted election E'; y-axis is the average distance between the two winning committees.

The average distance between winning committees in the original and the modified election can be found in Figure 1.

Examining Figure 1, we find that in all considered settings, winning committees are highly non-resilient. In particular, changing only 1% of the approvals *at random* leads to (on average) the replacement of two of the ten committee members. If we increase the change percentage further to 10%, around half of the committee gets replaced. These observations hint at a general non-robustness of Greedy-CC and Greedy-PAV, and a high fragility of produced outcomes.

In fact, winning committees under Greedy-CC and Greedy-PAV tend to produce—on average—committees of similar resilience when elections are sampled with the 1D model. However, while elections sampled with the 2D model are generally more resilient for both voting rules, Greedy-PAV has a slight edge over Greedy-CC.

Turning to the different operation types, one might intuitively expect that removing approvals leads to greater changes in the winning committee, as randomly removed approvals, generally speaking, hurt winning candidates with a higher probability. However, while this is indeed the case, the observed difference is not very prominent.

For 1D and 2D with Resampling, the trends are very similar, but both rules produce slightly more resilient committees (on average, around 0.5 to 1 fewer candidates need to be replaced given a change rate of 10%). For Resampling, the produced committees are a lot more resilient (the highest measured average of the number of candidates that need to be replaced was just over 2 for Greedy-CC), and the outcome is highly dependent on the choice of sampling parameters (see our full version [44] for a detailed discussion).

Experiment 2: The Role of Ties In Experiment 1, we circumvented the intractability of RCE for Greedy-CC and Greedy-PAV by applying lexicographic tie-breaking. In fact, this can be seen as a natural heuristic to solve RCE (compute S' using the same tie-breaking rule that was used to pick S). Ties play a surprisingly important role in winner determination of greedy Thiele rules [28]. Hence, our second experiment (Figure 2) investigates the effectiveness of this strategy and the importance of tie-breaking for the RCE problem.

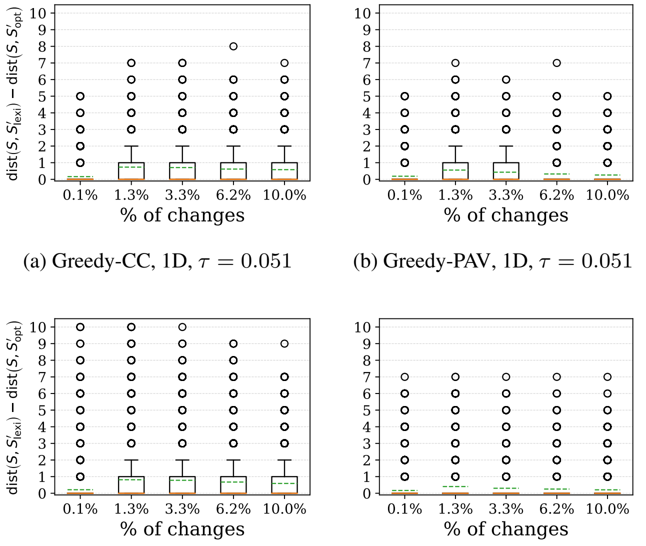

(a) Greedy-CC, 1D, $\tau = 0.051$ (b) Greedy-PAV, 1D, $\tau = 0.051$

(c) Greedy-CC, 2D, $\tau = 0.195$ (d) Greedy-PAV, 2D, $\tau = 0.195$

Figure 2: Results of Experiment 2. Focus is only on MIX operation. Orange lines represent the median and dashed green lines the mean. x-axis is the percentage change between the original election E and the adapted election E'; y-axis is the $\mathrm{dist}(S, S'_{\mathrm{lexi}}) - \mathrm{dist}(S, S'_{\mathrm{opt}})$, where S, and S'_{lexi} are the respective winning committees in E and E' under lexicographic tie-breaking, and S'_{opt} is chosen out of 100 tied winning committees in E' to be closest to S.

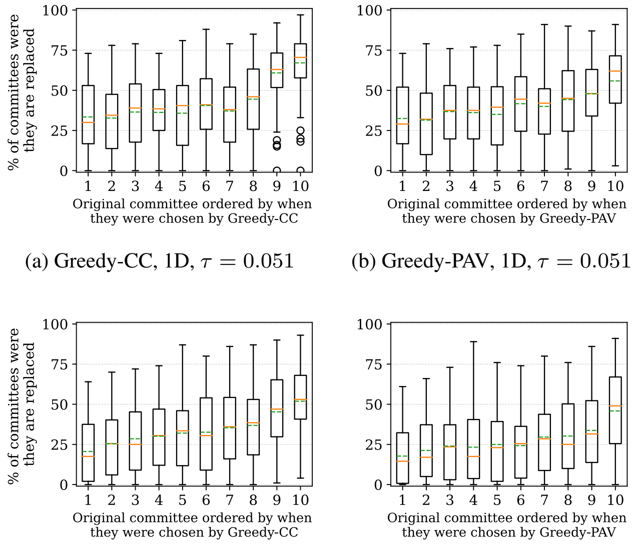

(a) Greedy-CC, 1D, $\tau = 0.051$ (b) Greedy-PAV, 1D, $\tau = 0.051$

(c) Greedy-CC, 2D, $\tau = 0.195$ (d) Greedy-PAV, 2D, $\tau = 0.195$

Figure 3: Results of Experiment 3. Focus is only on MIX operation and changes of 2.5% in relation to the original number of approvals were applied. Orange lines represent the median and dashed green lines the mean. x-axis corresponds to candidates from the original winning committee, ordered by when they were chosen in the given greedy Thiele rule; y-axis is the percentage of winning committees in the adapted elections where each candidate is replaced.

For each election E, we make S the winning committee in E under lexicographical tie-breaking. Subsequently, for all considered change percentages, we sample 100 elections E'. For each of these elections E', we compute up to 100 committees winning in the election and pick S'_{opt} to be the one closest to S. Further, let S'_{lexi} be the committee winning under lexicographic tie-breaking in E'. Using this, we compute the quantity $\mathrm{dist}(S, S'_{\mathrm{lexi}}) - \mathrm{dist}(S, S'_{\mathrm{opt}})$, i.e., the number of additional candidates that need to be replaced in S'_{lexi} compared to S'_{opt}. This can be interpreted as a lower bound on the "price" paid for using our heuristic instead of solving RCE optimally.

Across all considered sampling methods and voting rules, we find instances that show drastic differences in the contiguity offered by S'_{lexi} and S'_{opt}. Specifically, for a percentage change of 1.3%, the difference was non-zero in $\sim 1/3$ of cases, and at least 3 in $\sim 7.8\%$ of cases. For both voting rules, Greedy-CC and Greedy-PAV, the outliers tend to be more extreme on the 2D, compared to the 1D model.

The surprising trend for the difference to decrease again for higher changes to the underlying elections is due to these elections being much "noisier" than pure 1D and 2D elections. Thus, they tend to have far fewer ties. However, while this behaviour in terms of ties is not present in the 1D and 2D with Resampling models, they produce an almost identical picture, with similarly drastic outliers. Despite the generally higher resilience in the Resampling model, here we also witness outliers up to the value of 6. The results show that, for greedy Thiele rules, lexicographic tie-breaking does not constitute a reliable approximation for finding a committee that is close to the original one, highlighting the prominence of ties in these rules and motivating the search for optimal solutions via the RCE problem.

Experiment 3: Who Gets Replaced? We try to determine if some members of the initial winning committee get replaced more often when changes occur. In light of the round-based nature of greedy Thiele rules, one might expect the candidates chosen in later rounds to be generally weaker and hence more likely to be replaced. To address this hypothesis, we fix a change percentage of 2.5% for the *MIX* operation, and for each election E, sample 100 elections by applying $r = \lfloor \mathrm{app}(E) \cdot 2.5\% \rfloor$ changes to E. For each member c of the winning committee for E, we determine how many winning committees of the 100 sampled elections contain c. Figure 3 shows the results of this experiment as boxplots, grouped by the round in which candidates were added to the initially winning committee.

While the suspected correlation is present, the dependence is surprisingly weak. The most notable trend is that the last selected candidate is consistently the weakest, especially so for Greedy-CC under 1D elections. While, compared to the 1D model, the correlation is slightly more prominent in the 1D with Resampling model, Resampling has seemingly no effect on the correlation in the 2D model.

7 Conclusion

We presented a complexity dichotomy (along with parameterized complexity results) for RCE under Thiele rules and their greedy variants. We also conducted three experiments that unveiled interesting practical insights for our problem; in particular, they show that ties are highly prevalent and therefore the decision on how to break them to achieve a particular goal (as captured by RCE) is an important one.

Natural directions for future work include considering weighted cost in the replacement of candidates, or investigating RCE with different classes of rules, as well as removing or adding candidates or votes. It would also be interesting to explore the resilience of voting rules designed for temporal settings. A canonical candidate is, e.g., the voting mechanism for the AAAI Executive Council, where one-third of the positions are up for election each time.

Acknowledgements

NB was supported by the DFG project ComSoc-MPMS (NI 369/22). EE was supported by the AI Programme of The Alan Turing Institute and an EPSRC Grant EP/X038548/1.

References

[1] 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 Multi-Agent Systems (AAMAS)*, pages 107–115, 2015.

[2] N. Boehmer and R. Niedermeier. Broadening the Research Agenda for Computational Social Choice: Multiple Preference Profiles and Multiple Solutions. In *Proceedings of the 20th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1–5, 2021.

[3] N. Boehmer, P. Faliszewski, Łukasz Janeczko, A. Kaczmarczyk, G. Lisowski, G. Pierczyński, S. Rey, D. Stolicki, S. Szufa, and T. Wąs. Guide to Numerical Experiments on Elections in Computational Social Choice. In *Proceedings of the 33rd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 7962–7970, 2024.

[4] S. J. Brams and P. C. Fishburn. Approval Voting. *The American Political Science Review*, 72(3):831–847, 1978.

[5] R. Bredereck, J. Chen, D. Knop, J. Luo, and R. Niedermeier. Adapting Stable Matchings to Evolving Preferences. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1830–1837, 2020.

[6] R. Bredereck, A. Kaczmarczyk, and R. Niedermeier. Electing Successive Committees: Complexity and Algorithms. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1846–1853, 2020.

[7] R. Bredereck, P. Faliszewski, A. Kaczmarczyk, R. Niedermeier, P. Skowron, and N. Talmon. Robustness Among Multiwinner Voting Rules. *Artificial Intelligence*, 290:103403, 2021.

[8] R. Bredereck, T. Fluschnik, and A. Kaczmarczyk. When Votes Change and Committees Should (Not). In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 144–150, 2022.

[9] L. Bulteau, N. Hazon, R. Page, A. Rosenfeld, and N. Talmon. Justified Representation for Perpetual Voting. *IEEE Access*, 9:96598–96612, 2021.

[10] H.-J. Böckenhauer, E. Burjons, M. Raszyk, and P. Rossmanith. Reoptimization of Parameterized Problems. *Acta Informatica*, 59:427–450, 2022.

[11] J. R. Chamberlin and P. N. Courant. Representative Deliberations and Representative Decisions: Proportional Representation and the Borda Rule. *The American Political Science Review*, 77(3):718–733, 1983. ISSN 0003-0554. Publisher: [American Political Science Association, Cambridge University Press].

[12] N. Chandak, S. Goel, and D. Peters. Proportional Aggregation of Preferences for Sequential Decision Making. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, pages 9573–9581, 2024.

[13] M. Charikar, C. Chekuri, T. Feder, and R. Motwani. Incremental Clustering and Dynamic Information Retrieval. *SIAM Journal on Computing*, 33(6):1417–1440, 2004.

[14] T. Delemazure, J. Lang, J. Laslier, and M. R. Sanver. Approval with Runoff. In *Proceedings of the 31th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 230–236. ijcai.org, 2022.

[15] E. M. Deltl, T. Fluschnik, and R. Bredereck. Algorithmics of Egalitarian versus Equitable Sequences of Committees. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2651–2658, 2023.

[16] V. Do, M. Hervouin, J. Lang, and P. Skowron. Online Approval Committee Elections. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 251–257, 2022.

[17] E. Elkind, S. Kraiczy, and N. Teh. Fairness in Temporal Slot Assignment. In *Proceedings of the 15th International Symposium on Algorithmic Game Theory (SAGT)*, pages 490–507, 2022.

[18] E. Elkind, T. Y. Neoh, and N. Teh. Temporal Elections: Welfare, Strategyproofness, and Proportionality. In *Proceedings of the 27th European Conference on Artificial Intelligence (ECAI)*, 2024.

[19] E. Elkind, S. Obraztsova, and N. Teh. Temporal Fairness in Multiwinner Voting. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, pages 22633–22640, 2024.

[20] E. Elkind, S. Obraztsova, and N. Teh. Verifying Proportionality in Temporal Voting. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 2246–2248, 2024.

[21] P. Faliszewski, G. Gawron, and B. Kusek. Robustness of Greedy Approval Rules. In *Proceedings of the 19th European Workshop on Multi-Agent Systems (EUMAS)*, pages 116–133, 2022.

[22] J. Flum and M. Grohe. *Parameterized Complexity Theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2006.

[23] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979. ISBN 978-0-7167-1044-8.

[24] G. Gawron and P. Faliszewski. Robustness of Approval-Based Multiwinner Voting Rules. In *Proceedings of the 6th International Conference on Algorithmic Decision Theory (ADT)*, volume 11834, pages 17–31. Springer, 2019.

[25] M. T. Godziszewski, P. Batko, P. Skowron, and P. Faliszewski. An Analysis of Approval-Based Committee Rules for 2D-Euclidean Elections. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5448–5455. AAAI Press, 2021.

[26] T. F. Gonzalez. Clustering to Minimize the Maximum Intercluster Distance. *Theoretical Computer Science*, 38:293–306, 1985. doi: 10.1016/0304-3975(85)90224-5.

[27] S. Hartung and R. Niedermeier. Incremental List Coloring of Graphs, Parameterized by Conservation. *Theoretical Computer Science*, 494: 86–98, 2013.

[28] L. Janeczko and P. Faliszewski. Ties in Multiwinner Approval Voting. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2765–2773, 2023.

[29] D. M. Kilgour. Approval Balloting for Multi-winner Elections. In J.-F. Laslier and M. R. Sanver, editors, *Handbook on Approval Voting*, Studies in Choice and Welfare, pages 105–124. Springer, Berlin, Heidelberg, 2010. doi: 10.1007/978-3-642-02839-7_6.

[30] M. Lackner. Perpetual Voting: Fairness in Long-Term Decision Making. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, pages 2103–2110, 2020.

[31] M. Lackner and J. Maly. Proportional Decisions in Perpetual Voting. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5722–5729, 2023.

[32] M. Lackner and P. Skowron. *Multi-Winner Voting with Approval Preferences - Artificial Intelligence, Multiagent Systems, and Cognitive Robotics*. Springer Briefs in Intelligent Systems. Springer, 2023. ISBN 978-3-031-09015-8.

[33] J. Luo, H. Molter, A. Nichterlein, and R. Niedermeier. Parameterized Dynamic Cluster Editing. *Algorithmica*, 83:1–44, 2021.

[34] N. Misra and C. Sonar. Robustness Radius for Chamberlin-Courant on Restricted Domains. In *Proceedings of the 45th International Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM)*, volume 11376, pages 341–353. Springer, 2019.

[35] T. Y. Neoh and N. Teh. Welfare Maximization in Perpetual Voting (Student Abstract). In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, pages 23597–23599, 2024.

[36] R. Niedermeier. *Invitation to Fixed-Parameter Algorithms*. Oxford University Press, 2006.

[37] C. H. Papadimitriou. *Computational Complexity*. Academic Internet Publ., 2007. ISBN 978-1-4288-1409-7.

[38] A. D. Procaccia, J. S. Rosenschein, and A. Zohar. On the Complexity of Achieving Proportional Representation. *Social Choice and Welfare*, 30(3):353–362, 2008.

[39] B. Schieber, H. Shachnai, G. Tamir, and T. Tamir. A Theory and Algorithms for Combinatorial Reoptimization. *Algorithmica*, 80:576–607, 2018.

[40] S. Szufa, P. Faliszewski, L. Janeczko, M. Lackner, A. Slinko, K. Sornat, and N. Talmon. How to Sample Approval Elections? In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 496–502, 2022.

[41] T. N. Thiele. *Om flerfoldsvalg. In Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandlinger*. 1895.

[42] Y. Yang and J. Wang. Parameterized Complexity of Multi-winner Determination: More Effort Towards Fixed-Parameter Tractability. In *Proceedings of the 16th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 2142–2144, 2018.

[43] V. Zech and N. Teh. ECAI: Multiwinner Temporal Voting with Aversion to Change. https://github.com/ValentinZech/ECAI_Multiwinner_Temporal_Voting_with_Aversion_to_Change, 2024.

[44] V. Zech, N. Boehmer, E. Elkind, and N. Teh. Multiwinner Temporal Voting with Aversion to Change. *arXiv preprint arXiv:2408.11017*, 2024.