

Selecting Matchings via Multiwinner Voting: How Structure Defeats a Large Candidate Space

Niclas Boehmer, Markus Brill, and Ulrike Schmidt-Kraepelin

Abstract

Given a set of agents with approval preferences over each other, we study the task of finding k matchings fairly representing everyone’s preferences. We model the problem as an approval-based multiwinner election where the set of candidates consists of matchings of the agents, and agents’ preferences over each other are lifted to preferences over matchings. Due to the exponential number of candidates in such elections, standard algorithms for classical sequential voting rules (such as those proposed by Thiele and Phragmén) are rendered inefficient. We show that the computational tractability of these rules can be regained by exploiting the structure of the approval preferences. Moreover, we establish algorithmic results and axiomatic guarantees that go beyond those obtainable in the general multiwinner setting: Assuming that approvals are symmetric, we show that *Proportional Approval Voting (PAV)*, a well-established but computationally intractable voting rule, becomes polynomial-time computable, and its sequential variant (*seq-PAV*), which does not provide any proportionality guarantees in general, fulfills a rather strong guarantee known as *extended justified representation*. Some of our algorithmic results extend to other types of compactly representable elections with an exponential candidate space.

1 Introduction

Matching problems involving preferences occur in a wide variety of applications, and the literature has identified a host of criteria for choosing a “fair” matching [33]. In contrast to most of this work, we are interested in situations where *multiple* matchings between agents need to be chosen based on the preferences of agents over each other. Such situations occur naturally in applications where agents need to be matched multiple times, either successively or simultaneously. For instance, teachers often divide students into pairs for partner work, and multiple matchings might be required for different learning activities and different subjects. Several matchings also need to be found in pair programming, if, for example, one pairing is selected per project milestone. Other natural applications occur in workplaces where shifts are executed in pairs, which is often the case for security reasons (e.g., police officers or pilots usually work in shifts as pairs).

We model scenarios of this type as the problem of finding k matchings between agents based on the agents’ dichotomous (i.e., approval/disapproval) preferences over each other. More concretely, we associate with each agent an approval set, i.e., a subset of other agents that are approved by the agent. In the student/teacher scenario, approval sets of students could, for example, consist of all students they like, or of all students that are deemed compatible by the teacher. Preferences over agents are then lifted to preferences over matchings in a straightforward way: An agent approves a matching if and only if she is matched to an agent she approves. If the task were to find only a single matching, it would be natural to select a matching maximizing the number of approvers. However, as we are interested in finding multiple matchings, it is often possible to balance interests of agents. More concretely, this allows us to strive for the ideal of proportional representation: A group that makes up a p -fraction of the agents ($p \in [0, 1]$) should not be “less happy” than if this group

could decide on $\lfloor p \cdot k \rfloor$ of the matchings. This objective leads to considerations that are quite different from the classical goals of the matching literature such as stability or popularity. For a discussion and concrete example we refer to the related work section and Section 3.1.

The described type of fairness we strive for is captured by proportionality axioms from the approval-based multiwinner literature. By interpreting matchings as candidates and agents as voters in an election, our setting can be viewed as a special case of approval-based multiwinner elections [3]. As a consequence, voting rules and axiomatic results from this more general framework are applicable to our setting, to which we refer to as *matching elections*. In matching elections, we explicitly allow that a single candidate (i.e., matching) can be selected multiple times. This is in contrast to general approval-based multiwinner elections, where candidates can be selected at most once. As a rationale for our decision, observe that such a constraint would be rather artificial in our setting: Two matchings which only differ in a few pairs would already be considered as two distinct candidates in a matching election. Allowing matchings to be selected multiple times positions matching elections within the class of *party-approval elections* [13], a recently introduced subclass of approval-based multiwinner elections for which stronger axiomatic guarantees are obtainable.

Matching elections exhibit two characteristics that give rise to several interesting theoretical questions: First, the number of candidates in a matching election is exponential in the number of agents (and thus in the size of the description of an instance). As a consequence, a number of standard algorithms for applying voting rules or checking axiomatic guarantees no longer run efficiently, as they iterate over the candidate space. Second, preferences of agents have a very specific structure. For instance, it is possible to combine certain parts of two matchings, thereby obtaining a “compromise” candidate that is approved by some approvers of the first and some approvers of the second matching. Exploiting this structure has the potential to not only recover the computational tractability of voting rules, but also to prove proportional representation guarantees that go beyond those obtainable in more general multiwinner settings.

We also consider two natural special cases of matching elections: *symmetric* matching elections, where agents’ approvals are mutual, and *bipartite* matching elections, where agents are partitioned into two groups and agents only approve members of the opposite group. The previously described applications yield symmetric matching elections if, for example, approvals encode compatibility constraints. Similarly, bipartite matching elections arise whenever matched agents are required to have different attributes regarding professional experience, educational background, gender, etc.

1.1 Related Work

Recent years have witnessed a considerable amount of interest in approval-based multiwinner elections (see [30] for a survey); a particular focus has been on axiomatic properties capturing the notion of proportional representation [3, 11, 38, 12, 19, 14, 27, 34, 1, 35].

A fundamental challenge in computational social choice is to model settings where agents are presented an exponential number of possibilities. One method to deal with this is to assume that there exists some compact representation of the agents’ preferences that can be systematically lifted to preferences over all possibilities. This approach has been used, for example, in the study of hedonic games [8, 5, 6], fair division [10, 2], and single-winner voting in combinatorial domains [15, 32]. To the best of our knowledge, multiwinner elections with exponentially many candidates have not yet been considered.

In the following, we mention several established optimality criteria for selecting matchings based on ordinal preferences of agents [33] and discuss how they relate to the ideal of proportional representation (see Footnote 4 in Section 3.1 for a concrete example). Most prominently, stable matchings [20] as well as their fractional relaxation [37] are motivated

by the underlying “threat” that pairs of agents can block a matching. In contrast, proportionality prescribes that a pair of agents has the power to decide on $\lfloor 2 * (k/n) \rfloor$ matchings in the committee. An advantage of the latter is that we can represent the preferences of all agents, even those who would not be matched in any stable matching. Another related criterion is popularity [23, 28, 16]: A (fractional) matching is popular if it is preferred to any other matching by a majority of the agents (in expectation). While this is well-motivated for selecting a single matching, it leads to a “dictatorship of the majority” in the multiwinner case (as 51% of the agents could decide on the entire committee).

Bogomolnaia and Moulin [9] consider a setting that is similar to ours, except that *probability distributions* over matchings are chosen (rather than multiple matchings). They focus on the *egalitarian solution* [9], which chooses probability distributions maximizing the utility of the worst-off agent (breaking ties according to the leximin order). It was recently shown that such a probability distribution can be computed in polynomial time [22]. Bogomolnaia and Moulin [9] only consider bipartite and symmetric¹ instances and show that, under these restrictions, the egalitarian solution satisfies strong fairness and incentive properties. However, for non-symmetric instances, the fairness ideal behind the egalitarian solution is not completely satisfactory, as it ignores how hard it is to satisfy agents. Our axioms, in contrast, implicitly reward groups that can be matched easily to agents they approve.

1.2 Our Contributions

We establish matching elections as a novel subdomain of approval-based multiwinner elections with an exponential candidate space and initiate their computational and axiomatic study. By doing so, we are able to focus on a dimension of fairness which, to the best of our knowledge, has not been studied within the matching literature before. We consider several established (classes of) approval-based multiwinner rules (Thiele rules, Phragmén’s sequential rule, and Rule X) and proportionality axioms (PJR, EJR, and core stability). Exploiting the structure of matching elections, we prove a number of positive results. In particular, we show that all considered sequential rules can be computed in polynomial time despite the exponential candidate space. In fact, we show the slightly more general result that those rules are tractable in all elections where a candidate maximizing a weighted approval score can be found efficiently. We furthermore show that non-sequential Thiele rules such as PAV can be computed efficiently in symmetric and in bipartite matching elections, whereas they are computationally intractable in general matching elections (with a general matching election we mean an election that is not necessarily bipartite nor symmetric).

The additional structure of symmetric matching elections has axiomatic ramifications as well: We show that a large class of sequential Thiele rules satisfies EJR in this setting. This is particularly surprising as these rules are known to violate even significantly weaker axioms in general multiwinner elections. On the other hand, Phragmén’s sequential rule and Rule X do not satisfy stronger proportionality axioms compared to the general setting.

The proofs (or their completions) for results marked by (★) and some further results can be found in the full version of this paper [7].

2 Preliminaries

In this section, we define party-approval elections and recap some approval-based multiwinner voting rules and proportionality axioms. Let $\mathbb{N} = \{1, 2, \dots\}$ and $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$. For $n \in \mathbb{N}$, let $[n]$ denote the set $\{1, \dots, n\}$.

¹More precisely, Bogomolnaia and Moulin [9] do allow asymmetric preferences but assume that agents can only be matched if they approve each other, effectively rendering the setting symmetric.

2.1 Party-Approval Elections

A *party-approval election* [13] is a tuple (N, C, A, k) , where N is a set of agents, C a set of candidates, $A = (A_a)_{a \in N}$ a preference profile with $A_a \subseteq C$ denoting the approval set of agent $a \in A$, and $k \in \mathbb{N}$ the committee size.² A *committee* $W : C \rightarrow \mathbb{N}_0$ is a multiset of candidates, with the interpretation that $W(c)$ is the number of copies of candidate c contained in committee W . The *size* of a committee W is given by $\sum_{c \in C} W(c)$. For an agent $a \in N$ and a committee W , we let the happiness score $h_a(W)$ of a denote the number of (copies of) candidates from W approved by a , i.e., $h_a(W) = \sum_{c \in A_a} W(c)$. Moreover, $N_c = \{a \in N \mid c \in A_a\}$ denotes the set of *approvers* (also called supporters) of c , and $|N_c|$ is called the *approval score* of c . A *voting rule* maps a party-approval election (N, C, A, k) to a set of committees of size k . All committees output by a voting rule are considered tied for winning. Party-approval elections differ from the more general approval-based multiwinner elections [3] in that candidates can appear in a committee multiple times.

It is usually assumed that instances of an election are described by listing all candidates and approval sets explicitly. Since we will deal with elections with an exponential candidate space, we relax this assumption and only require that a representation of an election is given from which the full election can be reconstructed. We will show that several computational problems we consider in the following can be reduced to solving the following problem:

WEIGHTED APPROVAL WINNER

Input: A representation of a party-approval election (N, C, A, k) and a weight function $\omega : N \mapsto \mathbb{R}_{\geq 0}$.

Output: A candidate maximizing the total weight of its approvers, i.e., an element of $\operatorname{argmax}_{c \in C} \sum_{a \in N_c} \omega(a)$.

We let r_{waw} denote the running time of solving this problem.

2.2 Voting Rules from Multiwinner Voting

We describe several methods for computing committees. The output of a voting rule consists of all committees that can result from this method for some way of breaking ties.

Thiele Rules [39, 26] The class of w -Thiele rules is parameterized by a *weight sequence* w , i.e., an infinite sequence of non-negative numbers $w = (w_1, w_2, \dots)$ such that $w_1 = 1$ and $w_i \geq w_{i+1}$ for all i . Given a weight sequence w , the score of a committee W is defined as $sc_w(W) = \sum_{a \in N} \sum_{i=1}^{h_a(W)} w_i$. The rule w -Thiele selects committees maximizing this score. Setting $w_i = 1/i$ for all $i \in \mathbb{N}$ yields the arguably most popular w -Thiele rule known as *Proportional Approval Voting (PAV)*.

Sequential w -Thiele Rules (seq- w -Thiele rules) [39, 26] These variants of w -Thiele rules start with the empty committee and add candidates iteratively. Given a multiset W of already selected candidates, the *marginal contribution* of a candidate c is defined as $sc_w(W \cup \{c\}) - sc_w(W)$. In each step, seq- w -Thiele adds a candidate with a maximum marginal contribution. Setting $w_i = 1/i$ for all $i \in \mathbb{N}$, we obtain the rule *seq-PAV*.

Phragmén's Sequential Rule (seq-Phragmén) [36, 26] In seq-Phragmén, all agents start without money and continuously earn money (i.e., budget) at an equal and constant speed. As soon as there is a candidate c such that the group N_c jointly owns one dollar, such a candidate is added to the committee W and the budget of the group N_c is reduced to zero. All remaining agents keep their budget. This is repeated until the committee has size k .

Rule X [34] Initially, every agent a has a budget b_a of k/n dollars. Each candidate costs one dollar and a candidate c is said to be q -affordable if $\sum_{a \in N_c} \min\{b_a, q\} \geq 1$. In each

²To avoid trivial instances, we always assume that there exists at least one agent $a \in N$ with $A_a \neq \emptyset$.

round, we add a candidate which is q -affordable for minimum q and reduce the budget of the agents from N_c accordingly. The rule stops when there exists no q -affordable candidate for any $q > 0$. Note that Rule X might create a committee of size smaller than k ; in this case, the committee can be completed by choosing the remaining candidates arbitrarily [34].

Since seq- w -Thiele rules, seq-Phragmén, and Rule X add candidates to the committee one by one, we refer to these rules as *sequential rules*.

2.3 Axioms from Multiwinner Voting

Consider a party-approval election (N, C, A, k) . For $\ell \in [k]$, a set of agents $S \subseteq N$ is ℓ -cohesive if $|S| \geq \ell \frac{n}{k}$ and $\bigcap_{a \in S} A_a \neq \emptyset$. We consider three axioms capturing proportional representation [3, 38]:

Proportional Justified Representation A committee W provides *proportional justified representation (PJR)* if there does not exist an $\ell \in [k]$ and an ℓ -cohesive group S such that W contains strictly less than ℓ (copies of) candidates that are approved by at least one agent in S , i.e., $\sum_{c \in \bigcup_{a \in S} A_a} W(c) < \ell$.

Extended Justified Representation A committee W provides *extended justified representation (EJR)* if there does not exist an $\ell \in [k]$ and an ℓ -cohesive group S such that $h_a(W) < \ell$ for all $a \in S$.

Core Stability Given a committee W , we say that a group of agents $S \subseteq N$ *blocks* W if $|S| \geq \ell \frac{n}{k}$ for some $\ell \in [k]$ and there exists a committee W' of size ℓ such that $h_a(W') > h_a(W)$ for all $a \in S$. A committee W is *core stable* if it is not blocked by any group of agents.

Core stability implies EJR [3], and EJR implies PJR [38]. As it is standard in the literature [30], we say that a voting rule satisfies PJR/EJR/core stability if *all* committees in its output satisfy the respective condition.

3 Matching Elections

We now formally introduce matching elections and establish them as a special case of party-approval elections by giving a formal embedding. We familiarize ourselves with the newly introduced setting by proving some first observations on the special structure of the candidate space as well as showing that the weighted approval winner problem can be solved efficiently.

A *matching election* is a tuple (N, A, k) , where N is a set of agents, $A = (A_a)_{a \in N}$ a preference profile with $A_a \subseteq N \setminus \{a\}$ denoting the set of agents that are approved by agent a , and $k \in \mathbb{N}$ the number of matchings to be chosen. We let n denote the number of agents $|N|$. For notational convenience, we also call (N, A) a matching election.

A *matching* M is a subset of (unordered) pairs of agents, i.e., $M \subseteq \{\{a, b\} \mid a, b \in N, a \neq b\}$, such that no agent is included in more than one pair. If $\{a, b\} \in M$, we say that a is b 's *partner* or a is *matched to* b in M . A matching M is *perfect* if every agent has a partner. An agent a *approves* a matching M if a is matched to some agent b in M and a approves b , i.e., $b \in A_a$. We let N_M denote the set of agents approving matching M . We call a matching M *Pareto optimal* if there does not exist another matching M' such that $N_M \subsetneq N_{M'}$. We call a matching *minimal* if there does not exist another matching M' such that $M' \subset M$ and $N_M = N_{M'}$. An outcome of a matching election is a multiset (or committee) \mathcal{M} of k Pareto optimal and minimal matchings.³

³Minimality is only a formal restriction introduced for the sake of consistency, as any minimal matching can be extended to a (nearly) perfect matching by adding pairs of unmatched agents. On the other hand, Pareto optimality enforces that no clearly suboptimal matchings are part of the committee. We can convert any matching M into a Pareto optimal matching M' with $N_M \subseteq N_{M'}$ by solving one instance of WEIGHTED APPROVAL WINNER. For details, we refer to the proof of Lemma 1.




Figure 1: The figure on the left depicts the approval graph of the matching election (N, A) with $N = \{a_1, \dots, a_6\}$ and approval sets $A_{a_1} = \{a_2\}$, $A_{a_2} = \{a_3\}$, $A_{a_3} = \{a_4\}$, $A_{a_4} = \{a_3\}$, $A_{a_5} = \{a_3\}$, and $A_{a_6} = \{a_4\}$. The figure on the right depicts the three candidates c_1 , c_2 , and c_3 in the corresponding party-approval election.

Approval Graph The *approval graph* of a matching election (N, A) is a mixed graph defined as follows. The nodes of the approval graph are the agents in N and the edges depict the approval preferences: For two agents $a, b \in N$, there is an undirected edge $\{a, b\}$ if a approves b and b approves a ; and there is a directed edge (a, b) if a approves b but b does not approve a . For an example, see the illustration on the left side of Figure 1. Observe that a matching is minimal if and only if it contains only pairs which are connected by an (undirected or directed) edge in the approval graph. Every minimal and Pareto optimal matching is in particular a maximal matching in the approval graph when all edges are interpreted as undirected. Observe that the reverse direction is not true, i.e., not every maximal matching in the approval graph is Pareto optimal.

Bipartite and Symmetric Matching Elections We consider two natural domain restrictions for matching elections. A matching election (N, A) is called *bipartite* if there exists a partition of the agents $N = N_1 \dot{\cup} N_2$ such that each agent approves only agents from the other set, i.e., if $a \in N_i$ for $i \in \{1, 2\}$, then $A_a \subseteq N \setminus N_i$. Furthermore, we call a matching election (N, A) *symmetric* if agents' approvals are mutual, i.e., for two agents $a, b \in N$, $b \in A_a$ implies $a \in A_b$.

3.1 Embedding into Party-Approval Elections

A matching election (N, A, k) can be transformed into a party-approval election (N', C', A', k') with $N' = N$ and $k' = k$, and C' being the set of all Pareto optimal and minimal matchings in (N, A) and A' being the preference profile where each agent approves all candidates corresponding to matchings she approves. As we thereby establish matching elections as a subclass of party-approval elections, voting rules and axioms for party-approval elections directly translate to matching elections.

To illustrate the described transformation, we convert the matching election with six agents, whose approval graph is depicted on the left side of Figure 1, into a party-approval election. The candidates of the corresponding party-approval election are the three Pareto optimal and minimal matchings $c_1 = \{\{a_1, a_2\}, \{a_3, a_4\}\}$, $c_2 = \{\{a_1, a_2\}, \{a_3, a_5\}, \{a_4, a_6\}\}$, and $c_3 = \{\{a_2, a_3\}, \{a_4, a_6\}\}$, which are marked on the right side of Figure 1. The approval sets of the agents in the corresponding party-approval election are $A_{a_1} = \{c_1, c_2\}$, $A_{a_2} = \{c_3\}$, $A_{a_3} = A_{a_4} = \{c_1\}$, $A_{a_5} = \{c_2\}$, and $A_{a_6} = \{c_2, c_3\}$.

To get a feeling for proportionality in this election, let us set $k = 3$. Observe that the groups $\{a_3, a_4\}$ and $\{a_5, a_6\}$ make up one third of the electorate while at the same time, each of the groups can agree on a matching they commonly approve. In other words, both groups are 1-cohesive. Since a_3 and a_4 only approve c_1 , this is a strong argument in favor of choosing c_1 at least once. Given that c_1 is chosen at least once, adding c_2 seems preferable over adding c_3 , since c_2 is approved by three agents, two of which are completely unhappy so far, whereas c_3 is approved by only two so far completely unhappy agents. Lastly, there

is the choice between selecting c_3 , which would lead to every agent being satisfied at least once, and selecting one of the more popular matchings c_1 or c_2 again. In fact, all three resulting committees are core stable. PAV and seq-PAV both select $\{c_1, c_2, c_3\}$ in this example, whereas seq-Phragmén returns $\{c_1, c_2, c_3\}$ and $\{c_1, c_1, c_2\}$ as tied winners. Rule X terminates after adding c_1 and c_2 to the committee, which can be interpreted as a three-way tie between $\{c_1, c_1, c_2\}$, $\{c_1, c_2, c_2\}$, and $\{c_1, c_2, c_3\}$.⁴

While the focus of this paper is on matching elections, we note that some of our results apply to general party-approval elections. In particular, we establish our algorithmic results in Section 4.1 by reducing the computational problem at hand to solving instances of WEIGHTED APPROVAL WINNER (which is polynomial-time solvable for matching elections as shown in Section 3.2).

3.2 First Observations on the Candidate Space

In this subsection, we make some general first observations about features of our candidate space and the agents' approval sets. We start with an observation about the richness of the candidate space. Given a candidate (i.e., a matching) M and an agent a disapproving M , it is possible to obtain a new candidate M' that is approved by a and by all agents approving M except at most three: Assuming that a approves at least one agent, say b , to construct M' , we remove the pair from M containing b , say $\{b, c\}$ (if it exists), as well as the pair containing a , say $\{a, d\}$ (if it exists). Finally, we insert the pair $\{a, b\}$. Observe that, for the approval of a , we lost at most three approvals from M , namely the ones of b, c , and d .

Observation 1. *Given a matching election (N, A) with $n \geq 4$, let M be a matching and $a \in N \setminus N_M$ an agent with $A_a \neq \emptyset$. There exists a matching M' which is approved by a and all but at most three agents from N_M .*

Using this exchange argument, we can show that the number of approvals of each Pareto optimal matching M is at least $\frac{1}{3}$ of the number of approvals of any other matching M' .

Observation 2 (★). *Let (N, A) be a matching election and M be a Pareto optimal matching. For any other matching M' , it holds that $|N_M| \geq \frac{1}{3}|N_{M'}|$.*

Thus, we know that all candidates in a matching election are approved by the same number of agents up to a factor of three. For symmetric matching elections, it is even possible to tighten this bound: Here, all candidates are approved by the same number of agents and it is possible to perform one-to-one exchanges. This is also the key observation that helps proving that many seq- w -Thiele rules satisfy EJR. To see why the observation holds, recall that, in a symmetric matching election, the set of agents approving a minimal matching is exactly the set of matched agents. For the sake of contradiction, assume that there exist two minimal Pareto optimal matchings M and M' where M matches more agents than M' . Then, the symmetric difference of M and M' contains at least one path of odd length starting and ending with an edge from M . By augmenting M' along this path, it is possible to match an additional agent, which contradicts that M' is Pareto optimal.

Observation 3. *In symmetric matching elections, all candidates have the same approval score and correspond to maximum matchings in the approval graph.*

⁴A modified version of this example shows that stability, popularity and the egalitarian ideal (see Section 1.1) are incompatible with the ideal of proportional representation: Restrict the matching election from Figure 1 to the agents $\{a_1, a_2, a_3, a_4\}$. This election has two candidates $c = \{\{a_1, a_2\}, \{a_3, a_4\}\}$ and $c' = \{\{a_2, a_3\}\}$. If the committee size is $k = 4$, the proportional representation ideal implies that c is selected three times and c' is selected once. In contrast to this, both the only fractional stable and the only mixed popular solution would select c with probability 1. The egalitarian solution would select each of c and c' with probability $1/2$, therefore forcing an overrepresentation of a_2 .

The first part of Observation 3 already implies that symmetric matching elections have a strong structure. The second part has even further implications on the distribution of approvals of agents. These follow from the Gallai-Edmonds Structure Theorem [21, 17], which describes the structure of maximum matchings in undirected graphs. For our setting, the theorem implies that we can partition the agents into three sets W, X , and Y such that all agents from X and W approve every Pareto optimal matching. Moreover, in every Pareto optimal matching, all agents from X are matched to agents from Y and agents from W are matched among themselves. Using this theorem, we can convert every symmetric matching election into an essentially equivalent bipartite matching election. Here, the agents Y form one part of the bipartition and agents from X (plus some dummy agents) form the other part.

Weighted Approval Winner Problem For matching elections, we can solve WEIGHTED APPROVAL WINNER by solving two MAXIMUM WEIGHTED MATCHING instances: Given a matching election (N, A) and a weight function ω on the agents, we define a weight function w on the edges of the approval graph G of (N, A) such that for every matching M in G it holds that $\sum_{e \in M} w(e) = \sum_{a \in N_M} \omega(a)$. Clearly, if M is a maximum weight matching with respect to w , then it also maximizes $\sum_{a \in N_M} \omega(a)$. However, M might not be a candidate in the matching election, as it is not guaranteed to be Pareto optimal. We introduce a second weight function ω' on the agents giving all agents in N_M a weight of $n + 1$, and all agents in $N \setminus N_M$ a weight of 1. Again, we derive a weight function on the edges of G , w' , guaranteeing $\sum_{e \in M} w'(e) = \sum_{a \in N_M} \omega'(a)$. We show: If M' is a maximum weight matching with respect to w' , then M' is a solution to the WEIGHTED APPROVAL WINNER problem for the matching election (N, A) and weight function ω .

Lemma 1 (★). *Given a matching election (N, A) and a weight function ω , WEIGHTED APPROVAL WINNER is solvable in $\mathcal{O}(n^3)$ -time.*

Note that there exist other elections with an exponential candidate space for which WEIGHTED APPROVAL WINNER is polynomial-time solvable. For instance, for all party-approval elections (N, C, A, k) where the independent set system $(N, \{S \mid S \subseteq N_c \text{ for some } c \in C\})$ forms a matroid, WEIGHTED APPROVAL WINNER reduces to finding a maximum weight independent set. This problem is polynomial-time solvable if the independence of a set $S \subseteq N$ can be checked efficiently [29].

4 Computational Complexity of Winner Determination

In this section, we analyze the computational complexity of computing winning committees for different voting rules. While some of our results are tailored to matching elections, our algorithmic results in Section 4.1 are applicable to a wider class of elections with an exponential number of candidates. We start by considering sequential rules before we turn to w -Thiele rules. For w -Thiele rules, we first consider the general then the bipartite and lastly the symmetric setting.

4.1 Sequential Rules

For all considered sequential voting rules, we show that finding the next candidate to be added to the committee reduces to solving WEIGHTED APPROVAL WINNER. Recall that r_{waw} denotes the running time of solving the latter problem.

For sequential w -Thiele rules, this reduction is straightforward: Given a multiset W of already selected candidates, we set the weight of an agent a to its marginal contribution




Figure 2: Illustration of the situation in the proof of Theorem 1. The example depicts the budget curves for three different candidates c_1 , c_2 , and c_3 . The functions $f_{c_1}(t)$, $f_{c_2}(t)$, and $f_{c_3}(t)$ are depicted by a solid, dotted, and dashed line, respectively. The optimal value curve $f(t)$ is marked in blue.

to the score in case that a candidate in A_a is added to W , i.e., $\omega(a) = w_{h_a(W)+1}$. The candidate returned by WEIGHTED APPROVAL WINNER is then added to the committee.

Observation 4. *Given a party-approval election (N, C, A, k) and a weight sequence w , a committee that is winning under seq- w -Thiele can be computed in $\mathcal{O}(k \cdot r_{waw})$ -time.*

We show in the full version that a similar reduction also works for a local search variant of PAV [4]. As this variant satisfies core stability in party-approval elections [13], a core-stable outcome in a matching election can be computed efficiently.

Observation 5 (★). *Given a party-approval election (N, C, A, k) , a committee satisfying core stability can be computed in $\mathcal{O}(nk^4 \ln(k) \cdot r_{waw})$ -time.*

Our algorithm for Phragmén’s sequential rule employs WEIGHTED APPROVAL WINNER in a more involved way.

Theorem 1 (★). *Given a party-approval election (N, C, A, k) , a committee that is winning under seq-Phragmén can be computed in $\mathcal{O}(kn \cdot r_{waw})$ -time.*

Proof sketch. In each iteration, the problem of finding a candidate to be added to the committee can be described as follows. Every agent has a budget of $\beta_a \geq 0$ and constantly earns additional money. Thus, at time $t \in [0, 1]$, agent a owns $\beta_a + t$ dollars. The total budget of the approvers of a candidate c can be expressed as an affine linear function $f_c(t) = |N_c| \cdot t + \sum_{a \in N_c} \beta_a$. Moreover, $f(t) = \max_{c \in C} f_c(t)$ is the *optimal value curve*, taking the value of the maximum budget of any supporter group for a candidate at time t . Define t^* as the minimum value $t \in [0, 1]$ such that $f(t^*) = 1$. A candidate c^* with $f_{c^*}(t^*) = f(t^*) = 1$ is a feasible choice under seq-Phragmén in this iteration. See Figure 2 for an illustration. We argue that t^* can be computed by using a classical parametric optimization method, solving WEIGHTED APPROVAL WINNER as a subroutine.

The crux of finding t^* is that $f(t)$ is the maximum of exponentially many functions. However, $f(t)$ is a convex piecewise linear function with at most n breaking points, as the slope of $f(t)$ can take at most $n+1$ different values: for each candidate c , $|N_c| \in \{0, \dots, n\}$. If we knew all of the breaking points, we could iterate over all linear subintervals of $f(t)$ in order to find t^* . As described in the full version, the Eisner-Severance method [18] can be employed to find all breaking points, using $\mathcal{O}(n)$ calls to WEIGHTED APPROVAL WINNER. \square

By slightly modifying the above approach, we obtain a similar algorithm for Rule X. Here, for some fixed budgets of the agents, we need to find the minimum $q \in \mathbb{R}$ such that

the supporters of some candidate jointly have one dollar, assuming that each of them pays at most q . We again define the optimal value curve as the maximum budget of all supporter groups dependent on q . Unfortunately, in this case, the optimal value curve may neither be concave nor convex. However, by observing that we can partition the domain into n intervals such that the optimal value curve is a convex function in each interval, we can solve the problem using again the Eisner-Severance method as in the previous proof.

Theorem 2 (★). *Given a party-approval election (N, C, A, k) , a committee that is winning under Rule X can be computed in $\mathcal{O}(kn \cdot r_{\text{vaw}})$ -time.*

4.2 Non-Sequential Thiele Rules

In this section, we show that finding a winning committee in a general matching election is NP-hard for most w -Thiele rules. By contrast, as shown subsequently, this task becomes polynomial-time solvable for bipartite or symmetric matching elections.

In the party-approval setting, computing a winning committee of non-constant size under PAV is NP-hard [13]. However, if k is constant, the task can be solved in polynomial-time by iterating over all size- k committees. This is in contrast to our setting, where we prove NP-hardness of computing a winning committee under a large class of w -Thiele rules including PAV, even for $k = 2$. We reduce from the problem of deciding whether a 3-regular graph admits two edge-disjoint perfect matchings [25].

Theorem 3 (★). *Let w be a weight sequence with $w_1 > w_2 > 0$. Given a matching election (N, A, k) and some number $\alpha \in \mathbb{R}$, deciding whether there exists a committee \mathcal{M} of size k with $sc_w(\mathcal{M}) \geq \alpha$ is NP-complete for $k = 2$ and even if each agent approves at most three agents.*

In contrast to this, all w -Thiele rules are tractable in bipartite matching elections. The general idea of the algorithm is to construct all k matchings simultaneously with the help of a meta-election. In the meta-election, each agent is replaced by k copies. We then solve the WEIGHTED APPROVAL WINNER problem for the meta-election with appropriate agent weights to obtain a single matching which matches all k copies of each agent. From this, using Hall's theorem [24], we construct k matchings in the original instance.

Theorem 4 (★). *Let w be a weight sequence. In a bipartite matching election (N, A, k) , a winning committee under w -Thiele can be computed in $\mathcal{O}((kn)^3)$ -time.*

Unfortunately, the algorithm from the proof of Theorem 4 does not work for symmetric matching elections, as not every (non-bipartite) k -regular graph can be partitioned into k perfect matchings. Nevertheless, it is still possible to extend the algorithm by reducing each symmetric matching election to an essentially equivalent bipartite matching election.

Recall from Observation 3 that Pareto optimal matchings in symmetric matching elections have a strong structure, as they are, in particular, maximum matchings in the (undirected) approval graph. Using this, we can apply the Gallai-Edmonds Structure Theorem [21, 17] to obtain a partition of the agents into three sets W, X , and Y such that all agents from X and W approve every Pareto optimal matching. Moreover, in every Pareto optimal matching, all agents from X are matched to agents from Y and agents from W are matched among themselves. Using this, it is possible to transform every symmetric matching election into a bipartite one by putting agents from Y on the one side and agents from X and some dummy agents on the other side. It is then possible to construct from each winning committee under w -Thiele in the constructed bipartite election, a winning committee under w -Thiele in the original symmetric election. Using this, we can extend the algorithm from Theorem 4 to symmetric instances:

Corollary 1 (★). *Let w be a weight sequence. In a symmetric matching election (N, A, k) , a winning committee under w -Thiele can be computed in $\mathcal{O}((kn)^3)$ -time.*

5 Axiomatic Results

As matching elections are also party-approval elections, axiomatic guarantees from the latter setting still apply, i.e., PAV satisfies core stability, Rule X satisfies EJR, and seq-Phragmén satisfies PJR. Below, we study whether stronger axiomatic guarantees are obtainable for our subdomain. We focus on symmetric matching elections, as they exhibit a particularly strong structure. We start with a surprising positive result: A large class of sequential w -Thiele rules (including seq-PAV, which fails all considered axioms in general) satisfy EJR.

Theorem 5. *Let w be a weight sequence with $w_i > w_{i+1}$ for all $i \in \mathbb{N}$. Seq- w -Thiele satisfies EJR in all symmetric matching elections.*

Proof. Let (N, A, k) be a symmetric matching election. In Section 3.2 we have observed that the set N of agents can be partitioned into three sets W , X , and Y , such that in any Pareto optimal matching, all agents in $W \cup X$ are matched, agents in X are matched to agents in Y , and agents in W are matched among themselves. Thus, a group of agents violating EJR can only contain agents from Y .

Let $\mathcal{M} = \{M_1, \dots, M_k\}$ be some output of seq- w -Thiele (we assume that seq- w -Thiele selected matching i in iteration i). Let $\mathcal{M}_{<i} = \{M_1, \dots, M_{i-1}\}$. Assume for contradiction that there exists an EJR violation, i.e., for some $\ell \in [k]$, there is a set $S \subseteq N$ with $|S| \geq \ell n/k$, a Pareto optimal matching \widetilde{M} with $S \subseteq N_{\widetilde{M}}$ and $h_a(\mathcal{M}) < \ell$ for all $a \in S$.

We claim that the existence of S implies that in every iteration i , at least $|S|$ agents in Y which are matched in this iteration approve at most $\ell - 1$ matchings from $\mathcal{M}_{<i}$:

Claim. *For every $i \in [k]$, there exists a group $S_i \subseteq Y \cap N_{M_i}$ with $|S_i| = |S|$ and $h_a(\mathcal{M}_{<i}) \leq \ell - 1$ for all $a \in S_i$.*

Proof of Claim. Fix $i \in [k]$. If all agents in S are matched in M_i , setting $S_i = S$, the claim holds. Consider some $a \in S$ which is not matched in M_i . Since M_i and \widetilde{M} are maximum matchings in the approval graph of the instance, their symmetric difference consists of alternating cycles and even-length paths. In particular, there exists an even-length path starting in a and ending in some $b \in Y$ which is matched in M_i but not in \widetilde{M} . If $h_b(\mathcal{M}_{<i}) > h_a(\mathcal{M}_{<i})$, we could strictly increase the marginal contribution of M_i by augmenting along this path, as this would lead to a approving M_i at the cost of b disapproving it. Hence, $h_b(\mathcal{M}_{<i}) \leq h_a(\mathcal{M}_{<i})$. Since all even-length paths in the symmetric difference of M_i and \widetilde{M} are disjoint, we can construct S_i as follows: For every $a \in S$ choose a itself if $a \in N_{M_i}$ and else the agent at the other end of the corresponding even-length alternating path. \diamond

Let \mathcal{S} be the multiset of groups of agents S_i from the claim, i.e., $\mathcal{S} := \{S_1, \dots, S_k\}$. We define $g_a(\mathcal{S}) := |\{i \in [k] \mid a \in S_i\}|$ as the number of sets in \mathcal{S} that include agent a . By construction, we know that $g_a(\mathcal{S}) \leq \ell$ for all $a \in Y$: No group S_i contains an agent that is already included in ℓ of the groups S_1, \dots, S_{i-1} , as this would imply that a approves at least ℓ of the matchings in $\mathcal{M}_{<i}$. Since $S_i \subseteq N_{M_i}$ for all $i \in [k]$, we have $g_a(\mathcal{S}) \leq h_a(\mathcal{M}) \leq \ell - 1$ for all $a \in S$. Moreover, $\sum_{a \in Y} g_a(\mathcal{S}) = k|S|$, since every group S_i contains exactly $|S|$ agents from Y . We get the following contradiction (where the last step holds as $|Y| \leq n$):

$$\begin{aligned} \sum_{a \in Y} g_a(\mathcal{S}) &= \sum_{a \in S} g_a(\mathcal{S}) + \sum_{a \in Y \setminus S} g_a(\mathcal{S}) \\ &\leq (\ell - 1)|S| + \ell(|Y| - |S|) = \ell|Y| - |S| \leq \frac{k|S||Y|}{n} - |S| < k|S|. \quad \square \end{aligned}$$

However, sequential w -Thiele rules do not satisfy core stability in symmetric instances:

Proposition 1 (★). *Let w be a weight sequence. Committees returned by seq- w -Thiele are not guaranteed to be core stable, even if the given matching election is symmetric.*

Furthermore, Rule X and seq-Phragmén do not satisfy stronger guarantees in (symmetric) matching elections, compared to general party-approval elections.

Proposition 2 (★). *In symmetric matching election, committees returned by seq-Phragmén are not guaranteed to provide EJR and committees returned by Rule X are not guaranteed to be core stable.*

In the counterexamples for seq- w -Thiele, seq-Phragmén, and Rule X, there also exist other winning committees under these rules that satisfy the respective notion. Presumably, this is due to the richness of the candidate space, combined with a high number of ties in the execution of all three rules. It remains an open question whether the rules always return at least one winning committee satisfying the respective property.

6 Conclusion

We initiated the study of a multiagent problem at the intersection of social choice and matching theory: Given preferences of agents over each other, we model the problem of finding a representative multiset of matchings as a multiwinner election. Notwithstanding the difficulty presented by an exponential candidate space, we exploit the structure of the election domain to recover the computational tractability of some sequential rules, and also establish computational and axiomatic results that do not hold in the general setting.

There are several intriguing directions for future work on matching elections. First, one could consider axioms that are tailored to the specific structure of the setting. For example, a natural relaxation of core stability could only allow groups of agents to be matched among themselves in a deviation. Second, it would be natural to allow agents to rank-order potential matching partners and apply ordinal multiwinner voting procedures. Third, it would be interesting to identify other relevant multiwinner voting domains involving compactly representable preferences over an exponential candidate space.

Finally, in some applications, one is interested in finding multiple matchings of the same set of agents to be implemented one after the other. It is therefore natural to try to find a *sequence* of matchings, rather than simply a multiset (as done in this paper). While an arbitrary ordering of a proportional committee still provides proportionality if assessed as a whole, in such temporal settings, it might also be desirable to satisfy proportionality constraints for every sliding window of the sequence. One potential way to achieve this is to introduce depreciation weights to sequential rules, capturing the amount and recency of representation that agents have observed so far. Similar ideas have been recently explored within the context of approval-based multiwinner elections [31].

Acknowledgments

This work was supported by the Deutsche Forschungsgemeinschaft under grants NI 369/19 and BR 4744/2-1. We thank Ágnes Cseh for helpful discussions.

References

- [1] H. Aziz and B. E. Lee. Proportionally representative participatory budgeting with ordinal preferences. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*. AAAI Press, 2021. Forthcoming.
- [2] H. Aziz, S. Gaspers, S. Mackenzie, and T. Walsh. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence*, 227(2857):71–92, 2015.
- [3] 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.
- [4] H. Aziz, E. Elkind, S. Huang, M. Lackner, L. Sánchez-Fernández, and P. Skowron. On the complexity of extended and proportional justified representation. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, pages 902–909. AAAI Press, 2018.
- [5] H. Aziz, F. Brandl, F. Brandt, P. Harrenstein, M. Olsen, and D. Peters. Fractional hedonic games. *ACM Transactions on Economics and Computation*, 7(2):6:1–6:29, 2019.
- [6] N. Boehmer and E. Elkind. Stable roommate problem with diversity preferences. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 96–102. ijcai.org, 2020.
- [7] N. Boehmer, M. Brill, and U. Schmidt-Kraepelin. Selecting matchings via multi-winner voting: How structure defeats a large candidate space. Technical report, arXiv:2102.07441 [cs.GT], 2021.
- [8] A. Bogomolnaia and M. O. Jackson. The stability of hedonic coalition structures. *Games and Economic Behavior*, 38(2):201–230, 2002.
- [9] A. Bogomolnaia and H. Moulin. Random matching under dichotomous preferences. *Econometrica*, 72(1):257–279, 2004.
- [10] S. Bouveret, U. Endriss, and J. Lang. Fair division under ordinal preferences: Computing envy-free allocations of indivisible goods. In *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI)*, pages 387–392. IOS Press, 2010.
- [11] M. Brill, R. Freeman, S. Janson, and M. Lackner. Phragmén’s voting methods and justified representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 406–413. AAAI Press, 2017.
- [12] M. Brill, J.-F. Laslier, and P. Skowron. Multiwinner approval rules as apportionment methods. *Journal of Theoretical Politics*, 30(3):358–382, 2018.
- [13] M. Brill, P. Gözl, D. Peters, U. Schmidt-Kraepelin, and K. Wilker. Approval-based apportionment. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1854–1861. AAAI Press, 2020.
- [14] Y. Cheng, Z. Jiang, K. Munagala, and K. Wang. Group fairness in committee selection. In *Proceedings of the 20th ACM Conference on Economics and Computation (ACM-EC)*, pages 263–279. ACM, 2019.
- [15] Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. Preference handling in combinatorial domains: From AI to social choice. *AI Magazine*, 29(4):37–46, 2008.

- [16] Á. Cseh. Popular matchings. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 6, pages 105–122. AI Access, 2017.
- [17] J. Edmonds. Paths, trees and flowers. *Canadian Journal of Mathematics*, 17:449–467, 1965.
- [18] M. J. Eisner and D. G. Severance. Mathematical techniques for efficient record segmentation in large shared databases. *Journal of the Association for Computing Machinery*, 23(4):619–635, 1976.
- [19] B. Fain, K. Munagala, and N. Shah. Fair allocation of indivisible public goods. In *Proceedings of the 19th ACM Conference on Economics and Computation (ACM-EC)*, pages 575–592. ACM, 2018.
- [20] D. Gale and L. S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.
- [21] T. Gallai. Maximale systeme unabhängiger kanten. *Publications of the Mathematical Institute of the Hungarian Academy of Sciences, Series A*, 9:401–413, 1964.
- [22] D. García-Soriano and F. Bonchi. Fair-by-design matching. *Data Mining and Knowledge Discovery*, 34:1291–1335, 2020.
- [23] P. Gärdenfors. Match making: Assignments based on bilateral preferences. *Behavioral Science*, 20(3):166–173, 1975.
- [24] P. Hall. On representatives of subsets. *Journal of the London Mathematical Society*, 10:26–30, 1935.
- [25] I. Holyer. The NP-completeness of edge-colouring. *SIAM Journal on Computing*, 10(4):718–720, 1981.
- [26] S. Janson. Phragmén’s and Thiele’s election methods. Technical report, arXiv:1611.08826 [math.HO], 2016.
- [27] Z. Jiang, K. Munagala, and K. Wang. Approximately stable committee selection. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 463–472. ACM, 2020.
- [28] T. Kavitha, J. Mestre, and M. Nasre. Popular mixed matchings. *Theoretical Computer Science*, 412(24):2679–2690, 2011.
- [29] B. Korte and J. Vygen. *Combinatorial Optimization: Theory and Algorithms*. Springer, 5th edition, 2012.
- [30] M. Lackner and P. Skowron. Approval-based committee voting: Axioms, algorithms, and applications. Technical report, arXiv:2007.01795 [cs.GT], 2020.
- [31] Martin Lackner. Perpetual voting: Fairness in long-term decision making. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, pages 2103–2110. AAAI Press, 2020.
- [32] J. Lang and L. Xia. Voting in combinatorial domains. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 9. Cambridge University Press, 2016.

- [33] D. F. Manlove. *Algorithmics of Matching Under Preferences*. World Scientific Publishing Company, 2013.
- [34] D. Peters and P. Skowron. Proportionality and the limits of welfarism. In *Proceedings of the 21st ACM Conference on Economics and Computation (ACM-EC)*, pages 793–794. ACM, 2020.
- [35] D. Peters, G. Pierczynski, N. Shah, and P. Skowron. Market-based explanations of collective decisions. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*. AAAI Press, 2021. Forthcoming.
- [36] E. Phragmén. Sur une méthode nouvelle pour réaliser, dans les élections, la représentation proportionnelle des partis. *Öfversigt af Kongliga Vetenskaps-Akademiens Förhandlingar*, 51(3):133–137, 1894.
- [37] A. E. Roth, U. G. Rothblum, and J. H. Vande Vate. Stable matchings, optimal assignments, and linear programming. *Mathematics of Operations Research*, 18(4):803–828, 1993.
- [38] L. Sánchez-Fernández, E. Elkind, M. Lackner, N. Fernández, J. A. Fisteus, P. Basanta Val, and P. Skowron. Proportional justified representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 670–676. AAAI Press, 2017.
- [39] T. N. Thiele. Om flerfoldsvalg. *Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandlingar*, pages 415–441, 1895.

Niclas Boehmer
Research Group Algorithmics and Computational Complexity, TU Berlin
Berlin, Germany
Email: niclas.boehmer@tu-berlin.de

Markus Brill
Research Group Efficient Algorithms, TU Berlin
Berlin, Germany
Email: brill@tu-berlin.de

Ulrike Schmidt-Kraepelin
Research Group Efficient Algorithms, TU Berlin
Berlin, Germany
Email: u.schmidt-kraepelin@tu-berlin.de