



Proportional representation in matching markets: selecting multiple matchings under dichotomous preferences

Niclas Boehmer¹ · Markus Brill^{2,3} · Ulrike Schmidt-Kraepelin^{3,4}

Received: 23 April 2022 / Accepted: 13 January 2023 / Published online: 5 April 2023
© The Author(s) 2023

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. To formalize fairness, we apply the concept of proportional representation as studied in approval-based multiwinner elections. To this end, we model the problem as a 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 classical approval-based 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 that its sequential variant, which does not provide any proportionality guarantees in general, fulfills a rather strong guarantee known as extended justified

A preliminary version of this paper has appeared in the proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (Boehmer et al. 2022).

✉ Ulrike Schmidt-Kraepelin
u.schmidt-kraepelin@tu-berlin.de

Niclas Boehmer
niclas.boehmer@tu-berlin.de

Markus Brill
markus.brill@warwick.ac.uk

- ¹ Algorithmics and Computational Complexity, Technische Universität Berlin, 10587 Berlin, Germany
- ² Department of Computer Science, University of Warwick, Coventry CV4 7AL, UK
- ³ Efficient Algorithms, Technische Universität Berlin, 10587 Berlin, Germany
- ⁴ Centro de Modelamiento Matemático, Universidad de Chile, Santiago, Chile

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 single “fair” matching (Manlove 2013). In contrast to most of this work, we are interested in situations where *multiple* matchings between agents can be chosen, which allows to consider new dimensions of “fairness.” 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 can also be chosen in pair programming, for example, one pairing 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 with them 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 in the matching. If the task were to find only a single matching, it would be natural to select a matching maximizing the number of approvers (which, naturally, some agents might not approve). However, when selecting multiple matchings, it is possible to choose different matchings to balance interests of agents and to strive for *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, where k denotes the total number of matchings to be selected. This objective leads to considerations that are quite different from the classical goals of the matching literature such as stability or popularity. We review classical solution concepts in Sect. 1.1, formalize proportional representation in Sect. 2.3, and discuss a concrete example to illustrate the differences between them in Sect. 3.

The type of proportional representation-based fairness we strive for is captured by proportionality axioms from the approval-based multiwinner voting literature (Lackner and Skowron 2022). By interpreting matchings as candidates and agents as voters, our setting can be viewed as a special case of approval-based multiwinner elections. 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*. We explicitly allow that a single candidate (i.e., matching) can be selected multiple times and we refer to multisets of matchings as *committees*. This is in

contrast to general approval-based multiwinner elections, where candidates can be selected at most once.¹

This positions matching elections within the class of *party-approval elections* (Brill et al. 2022a), a recently introduced subclass of approval-based multiwinner elections for which stronger axiomatic guarantees are obtainable.

Matching elections exhibit two characteristics that make them an intriguing subdomain of party-approval elections and 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 the general party-approval setting.

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

The matching literature has established a variety of optimality criteria for selecting a single matching based on ordinal preferences of agents. In the following, we mention several of these criteria and discuss how they relate to the ideal of proportional representation (see Remark 2 in Sect. 3.1 for a concrete example). Most prominently, stable matchings (Gale and Shapley 1962) as well as their fractional relaxation (Roth et al. 1993) 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, where n denotes the

¹ 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. Moreover, allowing each matching to be selected only once can lead to inefficiencies that are hard to justify. For instance, in case there is a single matching that all agents approve, it would be rather unnatural to forbid choosing this matching multiple times (recall the initially mentioned applications of our problems, e.g., to the planning of shifts). Lastly, the constraint that each matching can be selected only once is prohibitive from a computational point of view: Since counting the number of different matchings in a graph is #P-hard (Valiant 1979), it would be computationally intractable to decide whether committees of a given size exist.

number of agents. 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 (Gärdenfors 1975; Kavitha et al. 2011; Cseh 2017). 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 (2004) 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* (Bogomolnaia and Moulin 2004), 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 (García-Soriano and Bonchi 2020). Bogomolnaia and Moulin (2004) only consider bipartite and symmetric² 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.

Proportional representation is traditionally studied in the context of multiwinner elections (Chamberlin and Courant 1983; Monroe 1995; Faliszewski et al. 2017). Recent years have witnessed a considerable amount of interest in multiwinner elections based on dichotomous preferences (Lackner and Skowron 2022). Within this setting, a particular focus has been on defining axiomatic properties capturing proportional representation (Aziz et al. 2017; Sánchez-Fernández et al. 2017; Brill et al. 2018; Peters et al. 2021) and on algorithms for guaranteeing (approximately) representative outcomes (Brill et al. 2017; Cheng et al. 2019; Jiang et al. 2020; Peters and Skowron 2020). Matching elections constitute a subdomain of *party-approval* elections Brill et al. (2022a).

The elections we consider in this paper have an exponential number of candidates, and thus require a way to represent agents’ preferences succinctly. Similar approaches involving compactly represented preferences of agents have been used, for example, in the study of hedonic games (Bogomolnaia and Jackson 2002; Aziz et al. 2019; Boehmer and Elkind 2020), fair division (Bouveret et al. 2010; Aziz et al. 2015b), and single-winner voting in combinatorial domains (Chevalerey et al. 2008; Lang and Xia 2016). To the best of our knowledge, multiwinner elections with exponentially many candidates have not yet been considered.

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

² Bogomolnaia and Moulin (2004) allow asymmetric preferences but assume that agents can only be matched if they approve each other, hence rendering the setting symmetric.

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 the method of equal shares (also known as Rule X); see Sect. 2.2 for definitions) and proportionality axioms (PJR, EJR, and core stability; see Sect. 2.3 for definitions). 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 matching elections and in bipartite matching elections, whereas they are computationally intractable in general³ matching elections. We present these results in Sect. 4, which we start with Table 1 summarizing our computational results.

The strong structure of symmetric matching elections has axiomatic ramifications as well: We show that a large class of sequential Thiele rules satisfy EJR in this setting. This is particularly surprising as these rules are known to violate even significantly weaker axioms in general approval-based multiwinner and party-approval elections. On the other hand, Phragmén’s sequential rule and the method of equal shares do not satisfy stronger proportionality axioms compared to the general setting. We present these results in Sect. 5, which we start with Table 2 summarizing our axiomatic results.

Lastly, in Sect. 6, we show that in matching elections it can be checked efficiently whether a committee satisfies EJR, whereas checking core stability or PJR is intractable. The problem of checking PJR is our only example for a computational problem that is polynomial-time solvable in the party-approval setting but NP-complete in the setting of matching elections.

2 Preliminaries

We define party-approval elections in Sect. 2.1 and recap some approval-based multiwinner voting rules in Sect. 2.2 and proportionality axioms in Sect. 2.3. Let $\mathbb{N} = \{1, 2, \dots\}$, $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$, and for all $n \in \mathbb{N}$ let $[n] = \{1, \dots, n\}$.

2.1 Party-approval elections

A *party-approval election* (Brill et al. 2022a) is a tuple (N, C, A, k) , where $N = [n]$ 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 $a \in N$, and $k \in \mathbb{N}$ the committee size.⁴ For each $a \in N$, we say that a *approves* all candidates from A_a and *disapproves* all other candidates. A *committee* $W : C \rightarrow \mathbb{N}_0$ is a multiset of candidates,

³ A “general” matching election is a matching election that is neither bipartite nor symmetric.

⁴ To avoid trivialities, we always assume that there exists at least one agent $a \in N$ with $A_a \neq \emptyset$.

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 nonempty 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 (Aziz et al. 2017) in that candidates can appear in a committee multiple times.

It is usually assumed that instances of a party-approval 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 (as we will argue later, for matching elections its representation consists of listing for each agent the agents it approves of). 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 $\omega : N \rightarrow \mathbb{Q}_{\geq 0}$.

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 four methods for computing committees. For each method the output of the corresponding voting rule consists of all committees that can result for some way of breaking ties.

Thiele Rules (*w*-Thiele) (Thiele 1895; Janson 2016)

The class of 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 Thiele rule known as *Proportional Approval Voting (PAV)*.

Sequential Thiele Rules (*seq-w-Thiele*) (Thiele 1895; Janson 2016)

These variants of 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 $i \in \mathbb{N}$ yields *seq-PAV*.

Phragmén's Sequential Rule (*seq-Phragmén*) (Phragmén 1894; Janson 2016)

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 of supporters of 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 .

The Method of Equal Shares (Equal Shares) (Peters and Skowron 2020)

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 round, we add a candidate c to the committee which is q -affordable for minimum q and reduce the budget of each agent $a \in N_c$ by $\min\{b_a, q\}$. The rule stops when there exists no q -affordable candidate for any $q > 0$. Note that Equal Shares might create a committee of size smaller than k ; in this case, the committee can be completed by choosing the remaining candidates arbitrarily (Peters and Skowron 2020).

Since sequential Thiele rules, *seq-Phragmén*, and Equal Shares add candidates to the committee one by one, we refer to these rules as *sequential rules*. See Table 1 for an overview of the computational complexity of the introduced rules in party-approval elections.

2.3 Axioms from multiwinner voting

Consider a party-approval election (N, C, A, k) . For $\ell \in [k]$, a set $S \subseteq N$ of agents is ℓ -cohesive if $|S| \geq \ell \frac{n}{k}$ and $\bigcap_{a \in S} A_a \neq \emptyset$. The intuition is that a group of $\ell \frac{n}{k}$ agents makes up an $\frac{\ell}{k}$ -fraction of the agents and should hence be entitled to decide on ℓ of the k members of a committee. If the agents of the group approve one joint candidate, they could chose ℓ copies of this candidate, yielding a happiness score of ℓ for each member of the group. Hence, ideally, in a committee satisfying proportional representation each agent that is included in some ℓ -cohesive group (for any $\ell \in [k]$) has a happiness score of at least ℓ . Since this ideal is not always achievable (Aziz et al. 2017; Brill et al. 2022b), the axioms *proportional justified representation* (Sánchez-Fernández et al. 2017) and *extended justified representation* (Aziz et al. 2017) capture relaxations of this ideal. *Core stability* strengthens these two axioms by additionally requiring proportional representation for non-cohesive groups:

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 fewer 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 $|S| \geq \ell^n$ 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 (Aziz et al. 2017), and EJR implies PJR (Sánchez-Fernández et al. 2017). As it is standard in the literature on approval-based multiwinner elections (Lackner and Skowron 2022), we say that a voting rule satisfies PJR/EJR/core stability if *all* committees in its output always satisfy the respective condition.

3 Matching elections

In this section, we formally introduce matching elections and related notation. In Sect. 3.1, we establish matching elections 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 (Sect. 3.2) as well as showing that the Weighted Approval Winner problem can be solved efficiently (Sect. 3.3).

A *matching election* is a tuple (N, A, k) , where $N = [n]$ 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. For notational convenience, we also call (N, A) a matching election.

A *matching* M is a set 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., $b \in A_a$, and *disapproves* a matching if a is unmatched or matched to an agent it does not approve of. (Note that agents are indifferent between being unmatched and being matched to an agent they do not approve of.) 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




Fig. 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

M' such that $M' \subsetneq M$ and $N_{M'} = N_M$. An outcome of a matching election is a multi-set (or committee) \mathcal{M} of k Pareto optimal and minimal matchings.⁵

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 their 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 Fig. 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 $a \in A_b$. Note that the approval graph of a symmetric matching election contains only undirected edges.

Remark 1 As described above, we only allow Pareto optimal and minimal matchings to be part of a committee. In principle, one could also require that the selected matchings satisfy some other criteria, e.g., the well-studied criterion of *stability*. A

⁵ Minimality 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. Note that 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). Thus, every matching can be easily replaced by a “better” minimal and Pareto optimal matching.

matching is called (*weakly*) *stable* if no pair of agents exists both strictly preferring each other to their currently assigned partner (Manlove 2013). For approval preferences this translates to the absence of a pair of agents that approve each other but both do not approve their partner in the matching. While for symmetric approvals, every Pareto optimal matching is stable, this is not the case for asymmetric approvals: For example, the matching $\{\{a_1, a_2\}, \{a_3, a_5\}, \{a_4, a_6\}\}$ in Fig. 1 is Pareto optimal but not stable. In fact, restricting the candidate space to stable matchings for asymmetric matching elections has drastic computational consequences: In Appendix 1, we show that the Weighted Approval Winner problem is no longer polynomial-time solvable for this variant of the model. Thus, none of our algorithmic results translate to this setting.

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 $k' = k$ by defining C' as the set of all Pareto optimal and minimal matchings in (N, A) and A' as 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 Fig. 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 Fig. 1. The approval sets of the agents in the 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 each, and at the same time, the members of each group can agree on a matching they commonly approve. In other words, both groups are 1-cohesive and are thus entitled to be represented at least once. 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. Equal Shares 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 $\{c_1, c_2, c_3\}$.

Remark 2 A modified version of this example shows that stability, popularity and the egalitarian ideal (see Sect. 1.1) are incompatible with the ideal of

proportional representation: Restrict the matching election from Fig. 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 fractional 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 .

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 Sect. 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 Sect. 3.3).

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:

Observation 1 Given a matching election (N, A) , 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 .

Proof 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 . \square

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 $|N_M| \geq \frac{1}{3}|N_{M'}|$.

Proof Let M be a Pareto optimal matching and M' some other matching. To prove the observation, we create a third matching \tilde{M} by the following procedure: Initially, set $\tilde{M} = M'$. As long as there exists an agent $a \in N_M$ not approving \tilde{M} , insert into \tilde{M} the pair from M containing a , say $\{a, b\}$, and delete the pairs containing a and b

from \tilde{M} (if they exist). This procedure terminates in $|N_M|$ steps, since every agent in N_M is considered at most once. After termination, $N_M \subseteq N_{\tilde{M}}$ and because M is Pareto optimal, also $N_M = N_{\tilde{M}}$. Since in each iteration the number of approvals went down by at most two, we get that $|N_M| \geq |N_{M'}| - 2|N_M|$. Thus, the observation follows. \square

From this, 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 sequential Thiele rules satisfy EJR in symmetric matching elections.

Observation 3 In symmetric matching elections, all candidates (*i*) have the same approval score and (*ii*) are maximum matchings in the approval graph.

Proof 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. \square

While part (*i*) of Observation 3 implies that symmetric matching elections have a strong structure, part (*ii*) has further implications on the distribution of approvals of agents. These follow from the Gallai–Edmonds Structure Theorem (Gallai 1964; Edmonds 1965), which describes the structure of maximum matchings in undirected graphs.

Gallai–Edmonds decomposition

Let $G = (V, E)$ be an undirected graph and W, X , and Y be a partition of the set of nodes V , such that Y is the (potentially empty) set of nodes which are not matched in all maximum matchings, X are their neighbors from $V \setminus Y$, and $W = V \setminus (Y \cup X)$. Concerning the notation, for some subset $S \subseteq V$ of nodes, we denote by $G[S]$ the subgraph induced by S , i.e., the graph $(S, E[S])$, where $E[S]$ is the set of all edges from E having both end nodes in S . The decomposition theorem (Gallai 1964; Edmonds 1965) says that

1. the graph $G[W]$ contains a perfect matching;
2. the connected components of $G[Y]$ are all factor-critical, i.e., removing any node from a connected component of $G[Y]$ results in a graph containing a perfect matching; and
3. in every maximum matching, all nodes from X are matched to distinct connected components of $G[Y]$.

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 (see the proof of Lemma 2 for a formal description).

3.3 Weighted approval winner problem

In the proof of Lemma 1 we show that for matching elections, we can solve Weighted Approval Winner by solving two maximum weighted matching instances.

Lemma 1 *Given a matching election (N, A) and a weight function $\omega : N \rightarrow \mathbb{Q}_{\geq 0}$, Weighted Approval Winner is solvable in $\mathcal{O}(n^3)$ -time.*

Proof Given a matching election (N, A) and a weight function ω on the agents, let $G = (N, \hat{E}, E)$ be the corresponding approval graph. Recall that G is a mixed graph, where N is the set of nodes, \hat{E} is the set of directed edges and E is the set of undirected edges.

We denote by $\bar{G} = (N, \bar{E})$ the undirected graph induced by G . More precisely, $\bar{E} = \{\{a, b\} \mid \{a, b\} \in E \text{ or } (a, b) \in \hat{E}\}$. We show how to solve the Weighted Approval Winner problem by computing two maximum weight matchings in \bar{G} , with respect to two different weight functions. We start by defining the first weight function on the edges $w : \bar{E} \rightarrow \mathbb{R}_{\geq 0}$. For every directed edge in G , $(a, b) \in \hat{E}$, let $w(\{a, b\}) = \omega(a)$ and for every undirected edge in G , $\{a, b\} \in E$, let $w(\{a, b\}) = \omega(a) + \omega(b)$.

By construction of the weight function w , for every matching M in \bar{G} , it holds that the weight of M with respect to w is equal to the weighted sum of all agents under ω that approve M , that is, $\sum_{e \in M} w(e) = \sum_{a \in N_M} \omega(a)$. Let M be a maximum weight matching in \bar{G} with respect to w . By the above observation, M also maximizes the weighted approval sum under ω among all matchings of the agents. Recall that, in order for M to be a candidate in the matching election (N, A) , it needs to be minimal and Pareto optimal. While M clearly satisfies minimality (every edge included in M is approved by at least one agent), Pareto optimality is not guaranteed since there might exist agents $a \in N$ with $\omega(a) = 0$.

In the following, we construct a matching M' in \bar{G} based on the matching M , that is minimal, satisfies $N_M \subseteq N_{M'}$ (and thus maximizes the weighted sum of approvals), and is Pareto optimal. To this end, we first define a second weight function on the agents, i.e., $\omega' : N \rightarrow \mathbb{R}_{\geq 0}$. More precisely, $\omega'(a) = n + 1$ if $a \in N_M$ and $\omega'(a) = 1$ if $a \in N \setminus N_M$. Again, we derive a weight function $w' : \bar{E} \rightarrow \mathbb{R}_{\geq 0}$ on the edges of \bar{G} as follows. For every directed edge in G , $(a, b) \in \hat{E}$, let $w'(\{a, b\}) = \omega'(a)$ and for every undirected edge in G , $\{a, b\} \in E$, let $w'(\{a, b\}) = \omega'(a) + \omega'(b)$.

Again, by construction of w' it holds for every matching M' in \bar{G} that $\sum_{e \in M'} w'(e) = \sum_{a \in N_{M'}} \omega'(a)$.

Let M' be a maximum weight matching in \bar{G} with respect to w' . It follows directly from the construction of the weight functions w' and ω' that

$N_M \subseteq N_{M'}$. Hence,

$$\sum_{e \in M} w(e) = \sum_{a \in N_M} \omega(a) \leq \sum_{a \in N_{M'}} \omega(a) = \sum_{e \in M'} w'(e),$$

and by the maximality of M with respect to w the two sides are equal. Hence, M' also maximizes the weighted approval sum with respect to ω . Moreover, M' is minimal, since every edge in M' is approved by at least one agent. Lastly, M' is also Pareto optimal since $\omega'(a) > 0$ for all agents $a \in N$. To see this, assume for contradiction that there exists a matching M'' with $N_{M'} \subsetneq N_{M''}$. However, since ω' is strictly positive for all agents, this would imply

$$\sum_{e \in M''} w'(e) = \sum_{a \in N_{M''}} \omega'_a > \sum_{a \in N_{M'}} \omega'_a = \sum_{e \in M'} w'(e),$$

a contradiction to the maximality of M' with respect to w' . We conclude that M' is a weighted approval winner for the matching election (N, A) and the weight function ω .

Thus, the Weighted Approval Winner problem for any matching election can be solved by computing two maximum weight matchings. This can be done in $\mathcal{O}(n^3)$ -time (Korte and Vygen 2012). □

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 (Korte and Vygen 2012).

4 Computational complexity of winner determination

In this section, we analyze the computational complexity of computing winning committees for different voting rules. We give an overview of our results from this section in Table 1. While some of our results are tailored to matching elections, our algorithmic results in Sect. 4.1 are applicable to a wider class of elections with an exponential number of candidates. We start by considering sequential rules in Sect. 4.1 before we turn to Thiele rules in Sect. 4.2. For Thiele rules, we first consider the general then the bipartite and lastly the symmetric setting.

Note that for all considered *sequential* rules, it was already shown in previous works (Aziz et al. 2015a; Brill et al. 2017; Peters and Skowron 2020) that the

Table 1 Summary of results on the complexity of computing a winning committee for several multiwinner voting rules

Rules	Party-approval elec.	Matching elec.	Sym. matching elec.	Bip. matching elec.
w-Thiele	NP-hard (Brill et al. 2022a)	NP-hard (Thm. 3)	P (Cor. 1)	P (Thm. 4)
seq-w-Thiele	P (Aziz et al. 2015a)	P (Obs. 4, Lemma 1)	P	P
seq-Phragmén	P (Brill et al. 2017)	P (Thm. 1, Lemma 1)	P	P
Equal Shares	P (Peters and Skowron 2020)	P (Thm. 2, Lemma 1)	P	P

The first column contains results for party-approval elections for which the input contains a complete list of the candidates, as originally defined by (Brill et al. 2022a). We remark that these previously known results within the setting of party-approval elections do not have any implications for matching elections. Our hardness result (Theorem 3) holds for w-Thiele rules satisfying $w_1 > w_2 > 0$

problem of finding a winning committee is solvable in time polynomial in the number of voters and the number of candidates. However, as in matching elections the number of candidates can be exponential in the input size, these algorithms are rendered inefficient. To regain the tractability of sequential rules, we propose new algorithms that use Weighted Approval Winner as a subroutine.

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 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 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 $\mathcal{O}(k \cdot r_{\text{waw}})$ -time.

A similar reduction also works for a local search variant of PAV (Aziz et al. 2018). Since this variant satisfies core stability in party-approval elections (Brill et al. 2022a), a core-stable outcome in a matching election can thus be computed efficiently.

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

Proof Brill et al. (2022a) showed that for a party-approval election (N, C, A, k) a core-stable committee can be computed by running a parameterized local search variant of PAV. (The method was originally introduced by Aziz et al. (2018) for general approval-based multiwinner elections.) In the following, we present this method tailored to the party-approval setting and show that computing a winning committee can be reduced to solving the Weighted Approval Winner problem $\mathcal{O}(nk^3 \ln(k))$ times.

Let w be the weight sequence corresponding to PAV, i.e., $w_i = 1/i$ for all $i \in \mathbb{N}$. The method *LS-PAV* starts by selecting an arbitrary size- k committee W . Then, it checks whether there exists an *improving swap* defined as follows. A *swap* replaces one (copy of a) candidate c which occurs at least once in W by (a copy of) some other candidate $c' \neq c$. Let W' be the committee obtained from W by removing (one copy of) c and adding one copy of c' . The swap replacing c by c' is called *improving* iff

$$sc_w(W') \geq sc_w(W) + \epsilon,$$

where $\epsilon := \frac{1}{(1+2(k-1))(k-1)k}$. LS-PAV searches for an improving swap (c, c') and, if an improving swap exists, updates the committee by exchanging (one copy of) c for (one copy of) c' . This procedure is repeated until there do not exist any improving swaps.

We claim that we can check whether there exists an improving swap (and if so, find one) in $\mathcal{O}(k \cdot r_{waw})$ -time: For a given committee W , iterate over all c that are selected at least once in W . Define \tilde{W} as the committee obtained from W by deleting (one copy of) c . We create a WEIGHTED APPROVAL WINNER instance by setting the weights of the agents to $\omega(a) = w_{n_a(\tilde{W})+1}$. Let c' be a weighted approval winner of this instance. Then, there exists an improving swap replacing c iff (c, c') is an improving swap. Moreover, Brill et al. (2022a) showed that the algorithm always terminates after performing at most $\mathcal{O}(nk^3 \ln(k))$ improving swaps and that the outcome is guaranteed to satisfy core stability. □

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 In each iteration, the problem of finding a candidate to be added to the committee can be described as follows. Each agent $a \in N$ has accumulated a budget of $\beta_a \geq 0$ in previous rounds and earns additional money in this iteration at constant speed. That is, at time $t \in \mathbb{R}_{\geq 0}$ of this iteration agent a owns $b_a(t) = \beta_a + t$ dollars. The total budget of the approvers N_c of a candidate $c \in C$ at time t can thus be expressed as an affine linear function $f_c(t) = \sum_{a \in N_c} \beta_a + |N_c| \cdot t$, which we call the candidate’s budget curve. Moreover, we define $f(t) = \max_{c \in C} f_c(t)$ as 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 $t \in \mathbb{R}_{\geq 0}$ such that $f(t) = 1$. Such




Fig. 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

a value always exists and lies in the real interval $[0, 1]$ since $f(0) \leq 1$ (by the definition of seq-Phragmén), $f(1) \geq 1$, and $f(t)$ is continuous on $[0, 1]$. A candidate c^* with $f_{c^*}(t^*) = f(t^*) = 1$ is a feasible choice to be added to the committee under seq-Phragmén in this iteration. See Fig. 2 for an illustration. In the following we argue that t^* and c^* can be computed by using a classical method from parametric optimization and solving Weighted Approval Winner as a subroutine.

Observe that the function $f(t)$ is increasing, piecewise linear, and convex, where the latter holds because taking the pointwise maximum of convex functions results in a convex function. In order to make parts of this proof also applicable to the proof of Theorem 2, we are only going to use that $f(t)$ is non-decreasing (and not that it is increasing) in the following. Note that for a given point $t \geq 0$, we can evaluate $f(t)$ by employing the Weighted Approval Winner problem using $b_a(t)$ as the weight of each agent $a \in N$ and computing the weight of the returned candidate. This also yields a candidate c with $f_c(t) = f(t)$.

The crux of finding t^* is that $f(t)$ is the maximum of exponentially many functions. However, we observe that the piecewise linear function $f(t)$ has at most n breaking points because the slope of $f(t)$ can take at most $n + 1$ different values: for each candidate $c, |N_c| \in \{0, \dots, n\}$ and thus the slope of each candidate’s budget curve lies in $\{0, \dots, n\}$. Hence, if we knew all breaking points of $f(t)$, we could find t^* by evaluating the resulting $\mathcal{O}(n)$ linear subintervals of $f(t)$. The Eisner-Severance method (Eisner and Severance 1976) can be employed to find the breaking points of $f(t)$, using $\mathcal{O}(n)$ calls to WEIGHTED APPROVAL WINNER.⁶

Since $f(t)$ is non-decreasing, we do not always have to find all breaking points in order to find t^* . Even though this does not improve the worst-case running time, we describe in the following an algorithm to find t^* , which mixes the idea of the Eisner-Severance method with a binary search approach.

⁶ More formally, the Eisner-Severance method takes as input a piecewise linear convex function g defined on an interval \mathcal{I} . Additionally, we need to include a method to *evaluate* g on some point $\tau \in \mathcal{I}$ which returns the value $g(\tau)$ and an affine linear function h with $g(\tau) = h(\tau)$ and $g(\tau') \geq h(\tau')$ for all $\tau' \in \mathcal{I}$. Given this input, the Eisner Severance method finds all breaking points of g on \mathcal{I} using $\mathcal{O}(z)$ evaluations, where z is the number of breaking points of g .

We start by searching for two candidates \underline{c} and \bar{c} such that $f_{\underline{c}}(0) = f(0)$ and $f_{\bar{c}}(1) = f(1)$ (by solving the Weighted Approval Winner problem twice). If $f_{\underline{c}}(0) = 1$, we are done. Moreover, if $f_{\underline{c}}(t) = f_{\bar{c}}(t)$ for all $t \in [0, 1]$, then there is no breaking point of $f(t)$ within the interval $[0, 1]$ and we can find t^* by solving $f_{\underline{c}}(t^*) = 1$. Otherwise, we calculate the intersection point of $f_{\underline{c}}(t)$ and $f_{\bar{c}}(t)$, say \hat{t} . By definition of f , we have $f_{\underline{c}}(\hat{t}) \leq f(\hat{t})$ and we distinguish the following two cases:

If $f_{\underline{c}}(\hat{t}) = f(\hat{t})$, we have found a breaking point of $f(t)$ and there is no other breaking point within the intervals $[0, \hat{t}]$ or $[\hat{t}, 1]$. Then, if $f(\hat{t}) \geq 1$, we find t^* by solving $f_{\underline{c}}(t^*) = 1$, and if $f(\hat{t}) < 1$, we find t^* by solving $f_{\bar{c}}(t^*) = 1$.

If $f_{\underline{c}}(\hat{t}) < f(\hat{t})$, we find a candidate \hat{c} such that $f_{\hat{c}}(\hat{t}) = f(\hat{t})$ (by solving Weighted Approval Winner). Then, if $f(\hat{t}) \geq 1$, we repeat the process for the pair $\{\underline{c}, \hat{c}\}$ and the interval $[0, \hat{t}]$. If $f(\hat{t}) < 1$, we repeat the process for the pair $\{\hat{c}, \bar{c}\}$ and the interval $[\hat{t}, 1]$. We can restrict ourselves to searching within one of the two intervals because $f(t)$ is non-decreasing. This recursive procedure yields a worst-case running time of $\mathcal{O}(n \cdot r_{\text{waw}})$, as we might iterate over all breaking points.

We have to execute the above procedure for each candidate to be added to the committee, and thus k times in total. This leads to an overall running time of $\mathcal{O}(kn \cdot r_{\text{waw}})$. □

By slightly modifying the above approach, we obtain a similar algorithm for Equal Shares. 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 Equal Shares can be computed in $\mathcal{O}(kn \cdot r_{\text{waw}})$ -time.*

Proof At any of the iterations within the execution of Equal Shares, the problem of finding a next candidate c^* to be added to the committee (or deciding to stop) can be described as follows: Each agent $a \in N$ has some leftover budget $b_a \leq k/n$ at the beginning of the iteration. Then, the budget of the supporters of a candidate $c \in C$ under the restriction that every agent pays at most $q \in \mathbb{R}$ can be expressed as $f_c(q) = \sum_{a \in N_c} \min\{b_a, q\}$. Similarly as in the proof of Theorem 1, we define the *optimal value curve* as $f(q) = \max_{c \in C} f_c(q)$. If $f(k/n) < 1$, there exists no q -affordable candidate for any q and Equal Shares terminates. Otherwise, we aim to find the minimum q^* in the real interval $[0, k/n]$ such that $f(q^*) = 1$. Such a value exists because $f(0) = 0$, $f(k/n) \geq 1$ (by the above assumption), and $f(q)$ is continuous on $[0, k/n]$. Then, a candidate c^* satisfying $f_{c^*}(q^*) = 1$ is a feasible next choice for Equal Shares. Given q^* , such a c^* can be found by one call to Weighted Approval Winner.




Fig. 3 Illustration of the situation in the proof of Theorem 2. The example depicts budget curves for candidates c_1, c_2 , and c_3 . The functions $f_{c_1}(q), f_{c_2}(q)$, and $f_{c_3}(q)$ are depicted by solid, dotted, and dashed lines, respectively. The optimal value curve $f(q)$ is marked in blue. Breaking points of type (i) and (ii) are marked by squares and circles, respectively. Intervals in which the optimal value curve is convex are marked by gray rectangles

Observe that $f(q)$ is non-decreasing, since $f_c(q)$ is non-decreasing for all $c \in C$. However, in contrast to the proof of Theorem 1, $f(q)$ is in general neither convex nor concave. As a consequence, we cannot directly apply the Eisner-Severance method (see Fig. 3 for an illustration). More concretely, consider some $q' \in [0, k/n]$ at which the piecewise linear function $f(q)$ has a breaking point. There are two distinct causes that can lead to a breaking point in $f(q)$: Cause (i) is a breaking point within the function f_c for some candidate c . This happens when at least one supporter of c has budget exactly q' . Cause (ii) is a cutting point of two functions f_{c_1} and f_{c_2} at point q' , where c_1 has less supporter with a budget higher than q' than c_2 has. While cause (i) leads to a decrease of the slope of $f(q)$, cause (ii) leads to an increase of the slope of $f(q)$. Importantly, if there exists no agent $a \in N$ with $b_a = q'$, we can be sure that cause (i) does not apply, and hence, the slope of $f(q)$ can only increase at q' .

We reindex the agents according to their budget, i.e., $b_{a_1} \leq b_{a_2} \leq \dots \leq b_{a_n}$. Due to the above reasoning, within each interval $[b_{a_i}, b_{a_{i+1}}]$ the function $f(q)$ is convex and its slope in this interval can take at most $n + 1$ distinct values. In order to find q^* , we now evaluate $f(q)$ at the borders of all of the intervals $[b_{a_i}, b_{a_{i+1}}]$ for all $i \in [n - 1]$ and select the left-most interval with $f(b_{a_{i^*}}) \leq 1 \leq f(b_{a_{i^*+1}})$. Then, we apply the Eisner-Severance method or its modified version as described in the proof of Theorem 1 in order to find the smallest $q^* \in [b_{a_{i^*}}, b_{a_{i^*+1}}]$ such that $f(q^*) = 1$.

Both our preprocessing step and the Eisner-Severance method can be performed in $\mathcal{O}(n \cdot r_{waw})$ -time. Doing so for all k iterations yields an overall running time of $\mathcal{O}(kn \cdot r_{waw})$. □

4.2 Non-sequential Thiele rules

In Sect. 4.2.1, we show that finding a winning committee in a general matching election is NP-hard for most Thiele rules. By contrast, as shown subsequently, this task becomes polynomial-time solvable for bipartite (Sect. 4.2.2) as well as for symmetric (Sect. 4.2.3) matching elections.

4.2.1 General matching elections

In the party-approval setting, computing a winning committee of non-constant size under PAV is NP-hard (Brill et al. 2022a). 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 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 (Holyer 1981).

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 $k = 2$, even if each agent approves at most three agents.*

Proof We reduce from the problem of deciding whether a 3-regular graph $G = (V, E)$ contains two edge-disjoint perfect matchings M_1 and M_2 .⁷ Let $V = \{v_1, \dots, v_\eta\}$ and $E = \{e_1, \dots, e_m\}$ and observe that 3-regularity of G implies that η is even and $m = 3\eta/2$. From G , we construct a matching election (N, A, k) as follows. We introduce one *node agent* a_i for each $v_i \in V$. Moreover, for each edge $\{v_i, v_j\} \in E$ with $i < j$, we add an edge gadget consisting of one *happy edge agent* a_{ij} and one *sad edge agent* a'_{ij} , where the node agent a_i approves the happy edge agent a_{ij} , the happy edge agent a_{ij} approves the sad edge agent a'_{ij} , and the node agent a_j approves the sad edge agent a'_{ij} . We set $k = 2$ and $\alpha = (5/2w_1 + 3/2w_2)\eta$ and refer to the two matchings to be found as M'_1 and M'_2 .

We call a matching M' of the agents N a *proper matching* if M' is approved by all node agents and, for each edge $\{v_i, v_j\} \in E$ with $i < j$, it either holds that $\{a_{ij}, a'_{ij}\} \in M'$ or that both $\{a_i, a_{ij}\} \in M'$ and $\{a'_{ij}, a_j\} \in M'$: A proper matching M' matches all η node agents to $\frac{\eta}{2}$ happy and $\frac{\eta}{2}$ sad edge agents and the remaining $m - \frac{\eta}{2}$ sad and $m - \frac{\eta}{2}$ happy edge agents to each other. We show later that every two matchings M'_1 and M'_2 with $sc_w(\{M'_1, M'_2\}) \geq \alpha$ need to be proper matchings. There exists a one-to-one correspondence between perfect matchings M in G and proper matchings M' of the agents N by including an edge $\{v_i, v_j\} \in E$ in M if and only if $\{a_{ij}, a'_{ij}\} \notin M'$. A visualization of the construction is depicted in Fig. 4. One example for the described correspondence are the two matchings marked by dashed red edges.

Before we prove the correctness of the forward and backward direction of the reduction, we will compute $sc_w(\{M'_1, M'_2\})$ assuming that M'_1 and M'_2 are proper matchings. It is possible to calculate $sc_w(\{M'_1, M'_2\})$ by summing up the w -Thiele score of M'_1 , i.e., $sc_w(\{M'_1\})$, and the marginal contribution of M'_2 given M'_1 , i.e., $sc_w(\{M'_1, M'_2\}) - sc_w(\{M'_1\})$. Since we have assumed that M'_1 is a proper matching, it is approved by all node agents and by all happy edge agents not matched to node agents. Thus, $sc_w(M'_1) = (\eta + m - \frac{\eta}{2})w_1 = 2\eta w_1$, as the graph is 3-regular. Turning

⁷ As this problem is equivalent to deciding whether G is 3-edge-colorable, NP-hardness follows from the work of Holyer et al. (1981)




Fig. 4 Example for the reduction from Theorem 3. The left side shows parts of a 3-regular graph and the right side the constructed matching election. Red dashed arcs indicate how a matching in the graph is transformed to a matching in the matching election

to the marginal contribution of M'_2 , as all node agents approve both matchings, they contribute ηw_2 . Sad edge agents do not approve any matching. For the happy edge agents a_{ij} , it is possible to distinguish four different cases:

- Case 1: $\{a_{ij}, a'_{ij}\} \notin M'_1$ and $\{a_{ij}, a'_{ij}\} \in M'_2$. In this case, the marginal contribution of a_{ij} is w_1 .
- Case 2: $\{a_{ij}, a'_{ij}\} \in M'_1$ and $\{a_{ij}, a'_{ij}\} \in M'_2$. In this case, the marginal contribution of a_{ij} is w_2 .
- Case 3: $\{a_{ij}, a'_{ij}\} \in M'_1$ and $\{a_{ij}, a'_{ij}\} \notin M'_2$. In this case, the marginal contribution of a_{ij} is 0.
- Case 4: $\{a_{ij}, a'_{ij}\} \notin M'_1$ and $\{a_{ij}, a'_{ij}\} \notin M'_2$. In this case, the marginal contribution of a_{ij} is 0.

By the assumption that M'_1 is proper, exactly $\frac{\eta}{2}$ happy edge agents are matched to node agents in M'_1 . Thus, Case 1 can occur at most $\frac{\eta}{2}$ times. Moreover, as there exist $\frac{3}{2}\eta$ happy edge agents and $\frac{\eta}{2}$ of them need to be matched to node agents in M'_2 , M'_2 can only be approved by η happy edge agents and thus Cases 1 and 2 combined can occur at most η times. Thus, as $w_1 > w_2$, the marginal contribution of M'_2 can be upper bounded by $\frac{\eta}{2} \cdot w_1 + \frac{\eta}{2} \cdot w_2$, leading to an upper bound for the combined score of any two proper matchings of $(5/2w_1 + 3/2w_2)\eta$. Note that we set α exactly to match this upper bound and that it is only possible to achieve it if the second matching is chosen in a way such that Case 1 occurs $\frac{\eta}{2}$ times (directly implying that the two matchings corresponding to M'_1 and M'_2 in G need to be edge-disjoint). We are now ready to show that there exist two edge-disjoint perfect matchings in G if and only if there exist two matching M'_1 and M'_2 with $sc_w(\{M'_1, M'_2\}) \geq (5/2w_1 + 3/2w_2)\eta$ in the constructed matching election.

(\Rightarrow) Let M_1 and M_2 be two edge-disjoint perfect matchings in G . Let M'_1 and M'_2 be the corresponding proper matchings of agents from N , i.e., for $t \in \{1, 2\}$:

$$M'_t := \bigcup_{\{v_i, v_j\} \in M_t: i < j} \{\{a_i, a_{ij}\}, \{a'_{ij}, a_j\}\} \cup \bigcup_{\{v_i, v_j\} \in E \setminus M_t: i < j} \{\{a_{ij}, a'_{ij}\}\}.$$

As M_1 is a perfect matching, M'_1 has a score of $2\eta w_1$. Moreover, as M_2 is perfect and edge-disjoint from M_1 , the first three cases concerning the marginal distribution of M'_2 from above all occur exactly $\frac{\eta}{2}$ times, while the last case does not appear at all. Hence, $sc_w(\{M'_1, M'_2\}) = (5/2w_1 + 3/2w_2)\eta$.

(\Leftarrow) Let M'_1 and M'_2 be two matchings such that $sc_w(\{M'_1, M'_2\}) \geq (5/2w_1 + 3/2w_2)\eta$, which implies that they are both proper matchings (we show this at the end of this proof). Let M_1 and M_2 be the corresponding perfect matchings in G , i.e., for $t \in \{1, 2\}$:

$$M_t := \{\{v_i, v_j\} \in E \mid \{a_i, a_{ij}\}, \{a'_{ij}, a_j\} \in M'_t\}$$

From $sc_w(\{M'_1, M'_2\}) \geq (5/2w_1 + 3/2w_2)\eta$, it follows by our previous observations and as $w_1 > w_2$ that Case 1 appears exactly $\frac{\eta}{2}$ times. Therefore, there exist $\frac{\eta}{2}$ happy edge agents that approve M'_2 but not M'_1 . This implies that the corresponding edges are not included in M_2 but included in M_1 . Thus, M_1 and M_2 are edge-disjoint.

It remains to be proven that for every two matching M'_1 and M'_2 in the constructed matching election with $sc_w(\{M'_1, M'_2\}) \geq (5/2w_1 + 3/2w_2)\eta$ it needs to hold that both of them are proper matchings. For every matching M' of agents in N and each edge $\{v_i, v_j\} \in E$ with $i < j$, one of the following four cases has to hold:

- Case 1: $\{a_i, a_{ij}\}$ and $\{a_j, a'_{ij}\} \in M'$
- Case 2: $(\{a_i, a_{ij}\} \in M'$ and $\{a_j, a'_{ij}\} \notin M')$ or $(\{a_i, a_{ij}\} \notin M'$ and $\{a_j, a'_{ij}\} \in M')$
- Case 3: $\{a_{ij}, a'_{ij}\} \in M'$
- Case 4: None of the three edges $\{a_i, a_{ij}\}$, $\{a_j, a'_{ij}\}$, and $\{a_{ij}, a'_{ij}\}$ is part of M'

Note that the last case never occurs, as M' cannot be Pareto optimal in this case (we can additionally match the respective happy and sad edge agent which leads to a strict extension of $N_{M'}$). This implies that the first three cases together happen $\frac{3}{2}\eta$ times. Let y denote the frequency of the first and z the frequency of the second case in M'_1 and \tilde{y} and \tilde{z} their frequencies in M'_2 .

We now bound the score of M'_1 and the marginal contribution of M'_2 in these four variables. The number of agents that approve M'_1 is $2y$ plus z plus the number of times $\frac{3}{2}\eta - y - z$ the third case appears:

$$sc_w(M'_1) = \left(\frac{3}{2}\eta + y\right)w_1.$$

Turning to the marginal contribution of the second matching M'_2 , we first consider the contribution of the node agents. We know that M'_2 is approved by $2\tilde{y} + \tilde{z}$

node agents. Since M'_1 is approved by $2y + z$ node agents, at most $\eta - 2y - z$ node agents can contribute with w_1 to the marginal score of M'_2 , while the remaining $2\tilde{y} + \tilde{z} - (\eta - 2y - z)$ contribute with w_2 . Turning to the edge agents, M'_2 is approved by $\frac{3}{2}\eta - \tilde{y} - \tilde{z}$ happy edge agents. Since the first matching is approved by all but $y + z$ happy edge agents, the number of happy edge agents contributing w_1 to the marginal score of M'_2 can be upper bounded by $y + z$, while the remaining $\frac{3}{2}\eta - \tilde{y} - \tilde{z} - y - z$ happy edge agents approving M'_2 contribute w_2 . Thus, the marginal contribution of M'_2 can be upper bounded as:

$$sc_w(\{M'_1, M'_2\}) - sc_w(\{M'_1\}) \leq (\eta - y)w_1 + \left(\frac{\eta}{2} + \tilde{y} + y\right)w_2.$$

Combining the two bound yields:

$$sc_w(\{M'_1, M'_2\}) \leq \frac{5}{2}\eta w_1 + \left(\frac{\eta}{2} + \tilde{y} + y\right)w_2.$$

Recall that $y, \tilde{y} \leq \frac{\eta}{2}$. Thus, as we have assumed that $sc_w(\{M'_1, M'_2\}) \geq (5/2w_1 + 3/2w_2)\eta$, it needs to hold that $y = \tilde{y} = \frac{\eta}{2}$, which implies that $z = \tilde{z} = 0$. From this it directly follows that M'_1 and M'_2 are proper matchings. \square

4.2.2 Bipartite matching elections

In contrast to the NP-hardness on general matching elections from the previous Sect. 4.2.1, all Thiele rules are tractable in bipartite matching elections.

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.*

Proof Let $(N = N_1 \dot{\cup} N_2, A, k)$ be a bipartite matching election and w a weight sequence. We assume without loss of generality that $|N_1| = |N_2|$. If this is not the case, we add $||N_1| - |N_2||$ dummy agents to the smaller side, which are neither approved by any of the original agents nor approve any of them. Clearly, every matching in this new instance can be mapped to a matching in the original instance of equal w -Thiele score, and vice versa.

We reduce our problem to solving one Weighted Approval Winner instance of a meta matching election. In the meta-election, each agent a is replaced by k copies $a^{(1)}, \dots, a^{(k)}$ and for all agents $a \neq b$ and all $i, j \in [k]$, copy $a^{(i)}$ approves copy $b^{(j)}$ if and only if agent a approves agent b . We construct a Weighted Approval Winner instance for the meta-election by setting $\omega(a^{(i)}) = w_i$ for all agents $a \in N$ and $i \in [k]$. Let \tilde{M} be a solution of the constructed Weighted Approval Winner instance. Because of the special structure of the meta-instance and the fact that the weight sequence w is non-increasing, we can assume without loss of generality that, for every agent a , there exists a threshold $i_a \in [k]$ such that her first i_a copies are exactly those that are matched to partners she approves in \tilde{M} . Hence, the contribution of the copies of an

agent $a \in N$ to the weight of \tilde{M} under ω equals $\sum_{j=1}^{i_a} w_j$. In the original instance, this is exactly the contribution of an agent to the w -Thiele score of a committee if she approves i_a of the matchings in the committee. In the following, we show that, indeed, we can find a committee \mathcal{M} of k matchings in the original instance such that every agent $a \in N$ approves i_a matchings in \mathcal{M} , i.e., $h_a(\mathcal{M}) = i_a$.

In order to do so, we extend the matching \tilde{M} to a perfect matching in the meta-instance respecting the bipartition. Recall that we can do so since we assumed that $|N_1| = |N_2|$. From that, we construct a “small” bipartite graph G which may contain parallel edges. More precisely, we define the multiset of edges R of G in the following, straightforward way: For every edge $\{a^{(i)}, b^{(j)}\} \in \tilde{M}$ for some $i, j \in [k]$, add one copy of the edge $\{a, b\}$ to R . Then, the multiset R induces a bipartite graph $G = (N_1 \dot{\cup} N_2, R)$ which is in particular k -regular. We extract k perfect matchings from G by a simply greedy procedure: From Hall’s Theorem (Hall 1935) it follows that every regular bipartite graph contains a perfect matching. To extract our matchings, we start by selecting some perfect matching in G and set it to be M_1 . Subsequently, we delete the edges contained in M_1 from G . Again, the obtained graph is regular and hence contains a perfect matching. By induction, we can proceed until we have selected k perfect matchings. Lastly, we modify the extracted perfect matchings by deleting pairs that are not approved by any of the two endpoints, or in other words, we make the matchings minimal. Note that Pareto optimality of the constructed matchings is guaranteed by the Pareto optimality of \tilde{M} .

Let $\mathcal{M} = \{M_1, \dots, M_k\}$ be the committee obtained from the above procedure. By construction, the number of matchings that are approved by some agent a is exactly i_a and hence $sc_w(\mathcal{M}) = \sum_{a \in N} \sum_{j=1}^{h_a(\mathcal{M})} w_j = \sum_{a \in N} \sum_{j=1}^{i_a} w_j$. It remains to show that this is also optimal. Assume for contradiction that there exists a committee $\mathcal{M}' = \{M'_1, \dots, M'_k\}$ with $sc_w(\mathcal{M}') > sc_w(\mathcal{M})$. From \mathcal{M}' , we construct a matching \tilde{M}' in the meta-instance as follows. For all $i \in [k]$ and every $\{a, b\} \in M'_i$, add the pair $\{a^{(i)}, b^{(i)}\}$ to the matching \tilde{M}' . Now, for every agent a , it holds that the number of partners that its copies $a^{(1)}, \dots, a^{(k)}$ approve in the meta-instance matching \tilde{M}' equals $h_a(\mathcal{M}')$. However, so far, it is not guaranteed that the satisfied copies of a are a prefix of $a^{(1)}, \dots, a^{(k)}$, which we need to ensure that the weight of \tilde{M}' under ω is maximal. We can ensure this by a simple exchange argument: Whenever there exists a copy $a^{(i)}$ matched to an unapproved agent, say $b^{(j)}$, while there exists another copy $a^{(i')}$ with $i' > i$ matched to an approved agent, say $c^{(j')}$, we replace the pairs $\{a^{(i)}, b^{(j)}\}$ and $\{a^{(i')}, c^{(j')}\}$ by the pairs $\{a^{(i)}, c^{(j')}\}$ and $\{a^{(i')}, b^{(j)}\}$. After doing this exhaustively, the contribution of the copies of any agent $a \in N$ to the weight of the matching \tilde{M}' under ω is exactly $\sum_{i=1}^{h_a(\mathcal{M}')} w_i$. Hence, under ω , the weight of the constructed matching \tilde{M}' is $sc_w(\mathcal{M}')$, which is strictly larger than the weight of the matching \tilde{M} , a contradiction to \tilde{M} being a weighted approval winner under ω in the meta-election.

We conclude the proof by showing the claimed running time. Solving an instance of Weighted Approval Winner in the meta-instance can be done in $\mathcal{O}((kn)^3)$ -time. Computing k perfect matchings in the “small” bipartite graph can be done in

$\mathcal{O}((kn)^2)$ -time. Lastly, transforming the resulting matchings to minimal matchings can be done in $\mathcal{O}(kn^2)$ -time. In total, we obtain a running time of $\mathcal{O}((kn)^3)$. \square

4.2.3 Symmetric matching elections

Unfortunately, the algorithm from the proof of Theorem 4 does not directly 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 (Gallai 1964; Edmonds 1965) (see Sect. 3.2) 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. We prove this formally in the following lemma:

Lemma 2 *There exists a function ψ mapping every symmetric matching election (N, A, k) to a bipartite and symmetric matching election $\psi((N, A, k))$ and a function φ mapping every committee in $\psi((N, A, k))$ to a committee in (N, A, k) such that, if a committee \mathcal{M} is winning under w -Thiele in $\psi((N, A, k))$, then $\varphi(\mathcal{M})$ is winning under w -Thiele in (N, A, k) . Both ψ and φ can be computed in $\mathcal{O}(n^3)$ -time.*

Proof Let (N, A, k) be a symmetric matching election. Applying the Gallai–Edmonds decomposition to the approval graph G of (N, A, k) , we can partition the set of agents into three sets W, X, Y with the above described properties. For any minimal Pareto optimal matching M in (N, A, k) , the only relevant information to determine N_M is the matching between the agents in X and the agents in Y . Following this idea, we construct the bipartite and symmetric matching election $\psi((N, A, k)) = (N' = N'_1 \dot{\cup} N'_2, A', k)$ that provides this information as follows. We set $N'_1 = Y$ and $N'_2 = X \cup D$, where D is a set of dummy agents. More precisely, D is constructed as follows: Let Y_1, \dots, Y_ℓ be the subgroups of agents corresponding to the connected components in $G[Y]$. For each group Y_i , we add $|Y_i| - 1$ dummy agents $d_i^{(1)}, \dots, d_i^{(|Y_i|-1)}$ to D . These agents approve the agents of Y_i and vice versa. Lastly, agents from Y and X approve each other in the new preference profile A' iff they




Fig. 5 The figure on the left depicts a symmetric matching election (N, A, k) with a partition of the agents into sets X, Y , and W as provided by the Gallai–Edmonds decomposition. The figure on the right depicts an (essentially equivalent) symmetric and bipartite matching election $\psi((N, A, k))$ as returned by the mapping defined in the proof of Lemma 2. The blue dashed edges indicate matchings that correspond to each other (in the mappings μ and φ)

approve each other in the original preference profile A . See Fig. 5 for an illustration of the mapping.

We further define two transformations μ and φ that, given a Pareto optimal matching in the symmetric instance (N, A, k) , return a Pareto optimal matching in $\psi((N, A, k)) = (N' = N'_1 \dot{\cup} N'_2, A', k)$ and vice versa. For a Pareto optimal matching M in (N, A, k) , we define $\mu(M)$ as follows: For each pair between an agent from X and an agent from Y in M , we add the same pair to $\mu(M)$. For all groups Y_i which now already have one matched agent, we arbitrarily match the remaining agents from Y_i to the dummy agents corresponding to this component. For all other groups Y_i , we leave exactly the agent unmatched which is unmatched in M and match the remaining agents to the dummy agents corresponding to this component in an arbitrary way. Observe that this transformation maintains the set of agents in Y that are matched and thus approve the matching, i.e., $N_M \cap Y = N'_{\mu(M)} \cap Y$. For the opposite direction, $\varphi(\cdot)$, let M' be a Pareto optimal matching in the constructed bipartite graph. By using Pareto optimality, all dummy agents need to be matched to agents they approve and at most one agent from each group Y_i is matched to an agent from X . Moreover, all agents in X need to be matched to agents they approve. We define the matching $\varphi(M')$ by first adding all pairs between agents from X and agents from Y that are part of M' and some perfect matching of the agents in W . Lastly, for all groups Y_i , we add a matching leaving exactly one agent in Y_i unmatched. More precisely, for those groups having an agent matched to an agent from X , we leave this agent unmatched and for a group Y_i not having any agent matched to an agent from

X , we leave the same agent unmatched which is unmatched in M' . Similarly to before, this transformation maintains the set of agents in Y that are matched and thus approve the matching, i.e., $N'_{M'} \cap Y = N_{\varphi(M')} \cap Y$.

We straightforwardly extend the two transformations from matchings to committees of matchings. More precisely, for a given multiset $\mathcal{S} = \{M_1, \dots, M_k\}$ in (N, A, k) , we define $\mu(\mathcal{S}) := \{\mu(M_1), \dots, \mu(M_k)\}$ and for a multiset $\mathcal{M} = \{M'_1, \dots, M'_k\}$ in $\psi((N, A, k))$, we define $\varphi(\mathcal{M}) := \{\varphi(M'_1), \dots, \varphi(M'_k)\}$. In order to clearly distinguish both instances, we write $h_a(\mathcal{S})$ for the number of matchings agent a from instance (N, A, k) approves in \mathcal{S} and $h'_a(\mathcal{M})$ for the number of matchings in \mathcal{M} an agent a from instance $\psi((N, A, k))$ approves. Observe that for some committee \mathcal{S} in (N, A, k) it holds that $h_a(\mathcal{S}) = h'_a(\mu(\mathcal{S}))$ for all $a \in Y$. Symmetrically, for some committee \mathcal{M} in $\psi((N, A, k))$ it holds that $h'_a(\mathcal{M}) = h_a(\varphi(\mathcal{M}))$ for all agents $a \in Y$.

We now turn to proving that φ and ψ fulfill the property stated in the theorem. Assume for contradiction that \mathcal{M} is winning under w -Thiele in $\psi((N, A, k))$, but $\varphi(\mathcal{M})$ is not winning under w -Thiele in (N, A, k) . Hence, there exists a size- k committee \mathcal{S} in (N, A, k) with $sc_w(\mathcal{S}) > sc_w(\varphi(\mathcal{M}))$. In particular, this implies that

$$\sum_{a \in Y} \sum_{i=1}^{h_a(\mathcal{S})} w_i > \sum_{a \in Y} \sum_{i=1}^{h_a(\varphi(\mathcal{M}))} w_i. \tag{1}$$

Now, using one of our transformations, from \mathcal{S} , we get a committee $\mu(\mathcal{S})$ in the bipartite instance $\psi((N, A, k))$ such that all agents in $X \cup D$ are matched k times and $h_a(\mathcal{S}) = h'_a(\mu(\mathcal{S}))$ for all agents in $a \in Y$, i.e., an agent from Y in the bipartite instance approves the same number of matchings from $\mu(\mathcal{S})$ as the corresponding agent from the symmetric instance approves in \mathcal{S} . We get

$$\begin{aligned} sc_w(\mu(\mathcal{S})) &= \sum_{a \in Y} \sum_{i=1}^{h'_a(\mu(\mathcal{S}))} w_i + \sum_{i=1}^k w_i \cdot (|X| + |D|) \\ &= \sum_{a \in Y} \sum_{i=1}^{h_a(\mathcal{S})} w_i + \sum_{i=1}^k w_i \cdot (|X| + |D|) \\ &> \sum_{a \in Y} \sum_{i=1}^{h_a(\varphi(\mathcal{M}))} w_i + \sum_{i=1}^k w_i \cdot (|X| + |D|) \\ &= \sum_{a \in Y} \sum_{i=1}^{h'_a(\mathcal{M})} w_i + \sum_{i=1}^k w_i \cdot (|X| + |D|) \\ &= sc_w(\mathcal{M}), \end{aligned}$$

where the inequality follows from (1). This yields a contradiction to the optimality of \mathcal{M} .

Concerning the running time of $\psi(\cdot)$, note that a Gallai–Edmonds decomposition can be computed by running Edmond’s blossom algorithm (Edmonds 1965) once which needs $\mathcal{O}(n^3)$ -time. Given such a decomposition, constructing $\psi(\cdot)$ can be done in $\mathcal{O}(n^2)$ -time. On the other hand, applying the transformation $\varphi(\cdot)$, we have

to compute one maximum cardinality matching of the vertices Y_i for each $i \in [\ell]$. Since the groups Y_i correspond to the connected components of $G[Y]$, this can be done by computing one maximum cardinality matching in $G[Y]$ (where some nodes were deleted). This can be done in $\mathcal{O}(n^3)$ -time. \square

Using the above lemma, 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, Equal Shares satisfies EJR, and seq-Phragmén satisfies PJR. Below, we study whether stronger axiomatic guarantees are obtainable for our subdomain (see Table 2 for an overview of our results). 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 Thiele rules (including seq-PAV, which fails all considered and even some weaker axioms in general party-approval elections) 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 Sect. 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 (where seq- w -Thiele selected matching M_i in iteration i) and let $\mathcal{M}_{<i} = \{M_1, \dots, M_{i-1}\}$ be the set of matchings selected in the first $i-1$ rounds. 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 \tilde{M} with $S \subseteq N_{\tilde{M}}$ 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 Fix some $i \in [k]$. Note that for each $a \in S$ and $i \in [k]$ it holds that $h_a(\mathcal{M}_{<i}) \leq \ell - 1$ by the definition of S . Thus, if all agents in S are matched in M_i ,

Table 2 Summary of results on the axiomatic properties of several multiwinner voting rules

Rules	Party-approval elections	Symmetric matching elections
PAV	Core stability (Brill et al. 2022a)	Core stability
seq- w -Thiele	Not PJR (Aziz et al. 2017)	EJR (Thm. 5), not core stability (Prop. 1)
Equal Shares	EJR, not core stability (Brill et al. 2022a)	EJR, not core stability (Prop. 3)
seq-Phragmén	PJR, not EJR (Brill et al. 2017)	PJR, not EJR (Prop. 3)

If a rule satisfies an axiom in the party-approval setting, then this also holds in the setting of general and symmetric matching elections. Our positive result in Theorem 5 holds for seq- w -Thiele rules satisfying $w_i > w_{i+1}$ for all $i \in \mathbb{N}$

setting $S_i = S$, the claim holds. Thus, consider some $a \in S$ which is not matched in M_i . Since M_i and \tilde{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 \tilde{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 \tilde{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. □

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 $i \in [k]$, for all $a \in S$, we have $g_a(\mathcal{S}) \leq h_a(\mathcal{M}) \leq \ell - 1$. 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 second-to-last step holds as $|S| \geq \ell n/k$ and the last step holds as $|Y| \leq n$):


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

This concludes the proof. □

However, sequential Thiele rules do not satisfy core stability in symmetric instances.

Proposition 1 *Let w be a weight sequence. In symmetric and bipartite matching elections, there exist committees that are winning under seq- w -Thiele but not core stable.*

Fig. 6 Approval graph of counterexample for core stability for sequential Thiele rules from Proposition 1



Proof To show the proposition, we present a symmetric matching election and construct a committee which is winning under seq- w -Thiele but fails to be core stable.⁸ The instance consists of three groups of *dummy agents* $A = \{a_1, \dots, a_{27}\}$, $B = \{b_1, \dots, b_{27}\}$, and $C = \{c_1, \dots, c_{41}\}$ and three *special agents* $x, y,$ and z . Approvals are symmetric. The special agent x approves all agents from A , the special agent y approves all agents from B , and the special agent z approves all dummy agents. See Fig. 6 for a visualization. Note that this instance consists of 98 agents. We set $k = n = 98$. Thus, every agent deserves to be represented by one matching.

We now construct a committee \mathcal{M} that is winning under seq- w -Thiele and argue that it is not core stable. In the first nine matchings, we match x and z to distinct agents from A and y to distinct agents from B . In the matchings ten to eighteen, we match y and z to previously unmatched agents from B and x to a previously unmatched agent from A . Note that the selected matchings are winning under seq- w -Thiele in their respective round, as we match only so-far unmatched dummy agents and assume $w_1 \geq w_2$. Overall, all agents from A and B are matched in exactly one of the first eighteen matchings. In the remaining 80 matchings, we match x to an agent from A , y to an agent from B , and z to an agent from C such that approvals within $A, B,$ and C are distributed as equally as possible. We can do so by constructing the matchings sequentially and always matching each special agent to the so far unhappiest agent from the respective group. Note that every matching in the constructed sequence is winning under seq- w -Thiele in the respective round, as approvals are distributed as equally as possible within each set and we have assumed that $w_i \geq w_{i+1}$ for all $i \in \mathbb{N}$. Moreover, after matching eighteen, it is always possible to match z to an agent of C in a winning matching, as over the whole construction process, each node from A and B approves the same or more of the already added matchings than a node from C ($|B|, |A| < |C|$).

To summarize, the summed happiness score of the agents from the three different sets are as follows: $\sum_{a \in A} h_a(\mathcal{M}) = \sum_{a \in B} h_a(\mathcal{M}) = 98 + 9 = 107$ and $\sum_{a \in C} h_a(\mathcal{M}) = 80$. Note that it holds that $\frac{107+1}{27} = 4$ and $\frac{80+2}{41} = 2$. By the pigeonhole principle, this implies that there exists at least one agent a from A that approves only three matchings from \mathcal{M} , at least one agent b from B that approves only three

⁸ To make some of the calculations easier, we construct the instance in a way such that $n = k$. Thus, the example should not be understood as a minimal counterexample.




Fig. 7 Approval graph of counterexample for core stability for Equal Shares from Proposition 2

matchings, and, as happiness scores are distributed as equally as possible, two agents c and c' from C which only approve one matching. We claim that the group $\{a, b, c, c'\}$ blocks \mathcal{M} . Note that this group deserves to be represented by four matchings. Let $\mathcal{M} \simeq$ be a set of four matchings, where a is matched to x and b is matched to y in all four matchings, while in two matchings, c is matched to z and in the other two, c' is matched to z . Since all four agents approve strictly more matchings from $\mathcal{M} \simeq$ than from \mathcal{M} , core stability is violated. \square

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

Proposition 2 *In symmetric matching elections, there exist committees that are winning under seq-Phragmén and do not satisfy EJR.*

Proof Consider a symmetric matching election consisting of three agents a_1, a_2 , and a_3 all approving each other. We set $k = 6$ and claim that the committee \mathcal{M} consisting of three times matching $\{\{a_1, a_2\}\}$ and three times matching $\{\{a_2, a_3\}\}$ is a winning committee under seq-Phragmén. In the first step, all possible non-empty matchings become affordable at $t = 0.5$. Breaking ties, we select $\{\{a_1, a_2\}\}$. Now, a_3 has 0.5 dollars left and thus needs to be included in the next matching. Again breaking ties, we select $\{\{a_2, a_3\}\}$. Continuing this way of breaking ties, we alternate between adding $\{\{a_1, a_2\}\}$ and $\{\{a_2, a_3\}\}$ until \mathcal{M} is constructed. However, \mathcal{M} violates EJR, as the group $\{a_1, a_3\}$ is 4-cohesive but $h_{a_1}(\mathcal{M}) = h_{a_3}(\mathcal{M}) = 3$. \square

Proposition 3 *In symmetric and bipartite matching elections, there exist committees that are winning under Equal Shares and do not satisfy core-stability.*

Proof We depict our counterexample in Fig. 7. It consists of 13 agents $\{w, x, y, z, a_1, a_2, a_3, b_1, b_2, c_1, c_2, d_1, d_2\}$. Approvals are symmetric. Agent w approves a_1, a_2, a_3 , and b_1 . Agent x approves b_1 and b_2 . Agent y approves c_1 and c_2 . Agent z approves d_1 and d_2 . We set $k = 13$. Thereby, each agent starts with a budget of one dollar. We now describe a run of Equal Shares on the constructed instance which returns a

committee \mathcal{M} that is not core stable. In the first eight rounds, all matchings which are approved by eight agents are $\frac{1}{8}$ -affordable. Breaking ties, we select the matching $\{\{w, b_1\}, \{x, b_2\}, \{y, c_1\}, \{z, d_1\}\}$ eight times. After that, all agents except $a_1, a_2, a_3, c_2,$ and d_2 have zero budget left. In the next three rounds, every matching which is approved by one of $a_1, a_2, a_3,$ and $c_2,$ and d_2 is $\frac{1}{3}$ -affordable. We select the matching $\{\{w, a_1\}, \{x, b_2\}, \{y, c_2\}, \{z, d_2\}\}$ three times. Subsequently, only a_2 and a_3 have budget, which makes all matchings which are approved by one of them 1-affordable. We select $\{\{w, a_2\}, \{x, b_2\}, \{y, c_1\}, \{z, d_1\}\}$ and $\{\{w, a_3\}, \{x, b_2\}, \{y, c_1\}, \{z, d_1\}\}$ as the last two matchings. Note that a_2 and a_3 both approve one matching from \mathcal{M} , while c_2 and d_2 approve three matchings from \mathcal{M} . Let $\mathcal{M} \simeq$ be a set of four matchings, where all matchings match z to d_2 and y to c_2 , two of the matchings match w to a_2 and the remaining two matchings match w to a_3 . The group $\{a_2, a_3, c_2, d_2\}$ blocks \mathcal{M} , as they deserve to be represented by four matchings and all four agents approve more matchings from $\mathcal{M} \simeq$ than from \mathcal{M} . \square

In the counterexamples for seq- w -Thiele, seq-Phragmén, and Equal Shares, 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 Complexity of checking axioms

In this section, we settle the computational complexity of checking whether a given committee provides any of our three proportionality guarantees. We first consider EJR.

Deciding whether a committee W in a party-approval election provides EJR can be reduced to solving Weighted Approval Winner: For each $\ell \in [k]$, we check whether there exists an ℓ -cohesive group violating EJR by marking all agents that approve less than ℓ matchings from W and checking whether there exists a candidate that is approved by at least $\ell \frac{n}{k}$ of the marked agents. The latter step can be solved by a single call to Weighted Approval Winner by assigning a weight of one to all marked agents and a weight of zero to all other agents.

Observation 6 Given a party-approval election (N, C, A, k) and a committee W , it is possible to check whether W provides EJR in $\mathcal{O}(k \cdot r_{\text{waw}})$ -time.

This approach does not extend to PJR. In fact, it turns out that checking whether a committee of matchings provides PJR is coNP-complete. This is in contrast to general party-approval elections, for which this problem can be solved in polynomial time (Brill et al. 2022a). Notably, this is our only example for a computational

problem that is polynomial-time solvable in the party-approval setting but intractable in the setting of matching elections.

Proposition 4 *Given a matching election (N, A, k) and a committee \mathcal{M} , checking whether \mathcal{M} provides PJR is coNP-complete, even if the given matching election is symmetric and bipartite.*

Proof We reduce from the NP-hard Clique problem on r -regular graphs (Garey and Johnson 1979), where given an undirected r -regular graph $G = (V, E)$ and an integer q the question is whether there exists a set of q pairwise adjacent nodes. We assume without loss of generality that $q > 3$. We construct a matching election and a committee \mathcal{M} as follows.

We insert one *node agent* a_v for each node $v \in V$, q *dummy agents*, and q *good agents*. All node and dummy agents approve all good agents and the other way round. Turning to the construction of \mathcal{M} , for each edge $\{u, v\} \in E$, we add a matching to \mathcal{M} that matches a_u, a_v , and $q - 2$ dummy agents to good agents. Further, we insert $2|E| + 1$ matchings in which each dummy agent is matched to a good agent. Lastly, we modify the instance such that $\frac{k}{n} = r - \frac{q-1}{2} + \frac{1}{q}$ by adding agents with empty approval ballot and matchings that match each dummy agent to a good agent. Note that each node agent approves r matchings from \mathcal{M} and a group of q agents deserves to be represented by $q\frac{k}{n} = qr - \binom{q}{2} + 1$ matchings. Moreover, note that only node agents can be part of a violating group, as good agents approve all matchings and dummy agents approve more than $\frac{2}{3}$ of the matchings. From this it follows that only node agents can be part of a violating group, as every group of agents that all approve the same matching and thus every cohesive group can have size at most $2q \leq \frac{2}{3}n$ (only the q good agents are approved by some agent). In the following, we show that there exists a size- q clique in G if and only if \mathcal{M} does not satisfy PJR. Intuitively, this holds as for a group of node agents $X = \{a_v \mid v \in V'\}$ for some subset of nodes $V' \subseteq V$, the set of matchings approved by some agent from X corresponds to the set of edges that are incident to some node from V' .

(\Rightarrow) Let $V' \subseteq V$ be a clique in G of size q , then exactly $qr - \binom{q}{2}$ different edges are incident to some node from V' (every node is incident to r edges and $\binom{q}{2}$ edges have both endpoints in V'). Since there exist q good agents, a matching where all agents from $\{a_v \mid v \in V'\}$ are matched to some good agent exists and is approved by all agents from $\{a_v \mid v \in V'\}$. Since $q\frac{k}{n} = qr - \binom{q}{2} + 1$, this implies that $\{a_v \mid v \in V'\}$ is $(qr - \binom{q}{2} + 1)$ -cohesive. Since they together approve only $qr - \binom{q}{2}$ different matchings from \mathcal{M} , $\{a_v \mid v \in V'\}$ is a violating group for PJR.

(\Leftarrow) Assume that there exists a violating group of agents X for PJR. Recall that only node agents can be part of X . Moreover, as only the q good agents are approved

by some node agent, it further needs to hold that $|X| \leq q$. For the sake of contradiction, assume that $|X| = x$ for some $x < q$. Each set of vertices of size x in G needs to be adjacent to at least $xr - \binom{x}{2}$ different edges. Thus, agents from X must approve at least $xr - \binom{x}{2}$ different matchings in \mathcal{M} , while they deserve to be represented by $x \cdot (r - \frac{q-1}{2} + \frac{1}{q})$ matchings. However, note that such a group cannot be violating, as, for all $x \in [1, q - 1]$, it holds that

$$\begin{aligned} xr - \binom{x}{2} &> x \cdot (r - \frac{q-1}{2} + \frac{1}{q}) \\ \Leftrightarrow r - \frac{x-1}{2} &> r - \frac{q-1}{2} + \frac{1}{q} \\ \Leftrightarrow -x &> -q + \frac{2}{q}, \end{aligned}$$

where the last inequality holds as $x \in [1, q - 1]$ and $q > 3$. Thus, X needs to have size q . For a group of size q to violate PJR, they need to approve at most $qr - \binom{q}{2}$ matchings together (as such a group can only be $(qr - \binom{q}{2} + 1)$ -cohesive). Thus, the set of vertices $\{v \mid a_v \in X\}$ is incident to at most $qr - \binom{q}{2}$ different edges in G implying that they form a clique in G . □

In our hardness reduction from Proposition 4, the given committee has a non-constant size. In fact, given a party-approval election, the problem whether a size- k committee W provides PJR is solvable in $\mathcal{O}(2^k \cdot r_{\text{waw}})$ -time: For all $\ell \in [k]$, we iterate over all $(\ell - 1)$ -subsets of candidates $W' \subseteq W$ and mark all agents whose approval set is a subset of W' . Subsequently, we check whether there exists a candidate approved by at least $\ell \frac{n}{k}$ of the marked agents. In this case, the group of $\ell \frac{n}{k}$ agents is ℓ -cohesive and by construction all of them approve only candidates from the set of $\ell - 1$ candidates W' . Thus, they constitute a violating group for PJR.

Observation 7 Given a party-approval election (N, C, A, k) and a committee W , checking whether W provides PJR can be done in $\mathcal{O}(2^k \cdot r_{\text{waw}})$ -time.

Finally, we show that checking core stability is computationally intractable, even for a constant committee size.

Proposition 5 (i) *Given a matching election (N, A, k) and a committee \mathcal{M} of not necessarily Pareto optimal matchings, checking whether \mathcal{M} is core stable is coNP-hard, even if $k = 6$ and the given matching election is symmetric and bipartite.* (ii) *Given a matching election (N, A, k) and a committee \mathcal{M} of Pareto optimal matchings,*

checking whether \mathcal{M} is core stable is coNP-hard, even if $k = 6$ and the given matching election is bipartite.

Proof We start by proving the first part of the statement, where we have symmetric approvals and we make use of Pareto dominated matchings in the constructed committee. Afterwards, we describe how we can get rid of the Pareto dominated matchings at the cost of losing symmetry.

In the NP-hard Exact Cover by 3-Sets (X3C) problem, we are given a universe X of size $3q$ and a collection C of 3-element subsets of X and the question is whether there exists an exact cover $C' \subseteq C$ of X . In fact, we reduce from the restricted version where each element appears in exactly three sets from C . Thus, it holds that $|C| = 3q$. We construct a matching election and a committee \mathcal{M} of size $k = 6$ as follows.

We start by describing the central part of the constructed matching election before adding additional agents to cope with some technical details. For each element $x \in X$, we insert one *element agent* a_x and one *dummy element agent* b_x . Moreover, for each set $c \in C$, we add one *set agent* a_c . Approvals are symmetric. For each element $x \in X$, the element agent a_x and the dummy element agent b_x approve each other. Moreover, the element agent a_x approves the three set agents a_c corresponding to sets in which x is contained. We construct \mathcal{M} such that each dummy element agent approves one matching, each element agent approves two matchings, and each set agent approves two matchings. Moreover, we modify the instance such that each possible blocking coalition needs to deserve to be represented by three matchings and needs to contain $7q$ of the so-far introduced agents.

To realize these requirements, we need to introduce several additional agents. That is, we introduce for each set $c \in C$, three *dummy set agents* b_c , d_c , and d'_c . Approvals are again symmetric. Agent b_c approves the set agent a_c and the two dummy agents d_c and d'_c . We construct \mathcal{M} such that b_c approves two matchings, d_c approves one matching and d'_c approves zero matchings. Lastly, to adjust the total number of agents, we add $14q$ *filling agents* with empty approval ballot. In total, the instance consists of $3q$ element agents and $3q$ dummy element agents, $3q$ set agents and $9q$ dummy set agents, and $14q$ filling agents, i.e., $32q$ agents in total. For a visualization of the reduction see Fig. 8.

We are now ready to construct \mathcal{M} realizing the already mentioned happiness scores of the agents. First, we add a matching where for each element $x \in X$, the element agent a_x is matched to the dummy element agent b_x and, for each set $c \in C$, the set agent a_c is matched to the dummy set agent b_c . In the second matching, we match all element agents a_x to a set agent a_c that they approve. (Note that such a perfect matching of element agents and set agents has to exist because these agents form a 3-regular bipartite graph.) Moreover, for each $c \in C$, we match dummy set agents b_c and d_c . Finally, we add four matchings that are not approved by anyone. Thus, as \mathcal{M} consists of six matchings and as the total number of agents is $32q$, each group of $\frac{32q}{6} = \frac{16q}{3}$ agents deserves to be represented by one matching. We now show that the given X3C instance (X, C) admits an exact cover if and only if there exists a group violating core stability in the constructed matching election.




Fig. 8 Example of the hardness reduction from Proposition 5 for Exact Cover By 3-Sets instance: $X = \{1, 2, 3, 4, 5, 6\}$ and $\{\{1, 2, 3\}, \{2, 4, 5\}, \{4, 5, 6\}\} \subseteq C$. Numbers in the superscripts denotes the number of matchings from \mathcal{M} the agent approves

(\Rightarrow) Let us assume that there exists an exact cover $C' \subseteq C$ of X . We claim that the group S consisting of all element and dummy element agents, all set agents corresponding to sets from C' , and all dummy set agents block committee \mathcal{M} . Note that S consists of $6q + 1q + 9q = 16q$ agents and thus deserves to be represented by three matchings. We now describe the three blocking matchings. For each $c \in C$, b_c is matched to d'_c in the first two of the three matchings and to d'_c in the third. For each $c = \{x_i, x_j, x_k\} \in C'$, we match a_c to a_{x_i} in the first matching, to a_{x_j} in the second matching, and to a_{x_k} in the third matching. This is always possible, as C' is an exact cover of X . Thereby, each element agent is matched to a set agent in one of the three matchings. We match each element agent in the remaining two matchings to the corresponding dummy element agent. Note that all element agents and all set agents corresponding to sets from C' approve all three matchings. All dummy element agents approve two matchings. For all $c \in C$, b_c approves all three matchings, d'_c approves two matchings and d'_c one matching. Thus, S is blocking.

(\Leftarrow) Assume that there exists a blocking coalition S for \mathcal{M} because of a multiset of matchings \mathcal{M}' . Observe that S cannot contain any filling agents because these agents do not approve anyone. Note that there exist only $3q$ non-filling agents that do not approve any matching from \mathcal{M} and only $9q$ non-filling agents that approve at most one matching from \mathcal{M} . Since each group of $\frac{16q}{3}$ agents deserves to be represented by one matching and $3q \cdot \frac{3}{16q} < 1$ and $9q \cdot \frac{3}{16q} < 2$, it needs to hold that

$|\mathcal{M}'| \geq 3$ and thus S needs to have size at least $16q$. Moreover, note that there cannot exist a blocking coalition that deserves to be represented by four matchings, as there exist only $18q$ non-filling agents.

To complete the proof, we need the following claim.

Claim *Let $S_{\text{set}} \subseteq S$ be the set of set agents a_c that are part of the blocking coalition S . Then, it holds that $|S_{\text{set}}| = q$.*

Proof As S can only contain non-filling agents, from $|S| \geq 16q$ and the fact that there exist only $18q$ non-filling agents of which $3q$ are set agents, it follows that $|S_{\text{set}}| \geq q$ needs to hold. To prove that $|S_{\text{set}}| = q$, first of all, note that all set agents from S_{set} need to approve all three matchings from \mathcal{M}' . Thus, in total, there exist $3|S_{\text{set}}|$ pairs in \mathcal{M}' each containing exactly one agent from S_{set} . Let w be the number of dummy element agents that are part of S , y the number of dummy set agents of the form d_c and z the number of dummy set agents of the form d'_c . Overall, it needs to hold that $|S_{\text{set}}| + w + y + z + 6q \geq 16q$ (as the only other non-filling agents are the $3q$ element agents and the $3q$ dummy set agents of the form b_c) and thus $|S_{\text{set}}| + w + y + z \geq 10q$. Note that as each dummy element agent needs to approve two matchings from \mathcal{M}' , even if $w = 3q$, each element agent can be matched to an agent from S_{set} in one matching. For the sake of contradiction, let us assume that $t := |S_{\text{set}}| - q > 0$. Then, $3t$ approvals for set agents are needed that do not come from element agents which are matched to the corresponding dummy element agent in the other two matchings. However, for each two of these $3t$ approvals, either an element agent needs to be matched more than once to a set agent or a dummy set agent b_c needs to be matched twice to a set agent. While the former implies that the corresponding dummy element agent cannot be part of the blocking coalition S , the latter implies that either one less dummy set agent of the form d_c or d'_c can be part of the blocking coalition. Thus, $t > 0$ implies that $w + y + z \leq 9q - \frac{3}{2}t$. Overall we get that $|S_{\text{set}}| + w + y + z = q + t + w + y + z \leq q + t + 9q - \frac{3}{2}t = 10q - \frac{1}{2}t$. As we have initially observed that $|S_{\text{set}}| + w + y + z \geq 10q$, it needs to hold that $t = 0$. This directly implies that $|S_{\text{set}}| = q$. \square

From the claim it directly follows that S consists of the agents S_{set} and all non-filling agents that are not set agents.

To ensure that all dummy element agents approve two matchings from \mathcal{M}' , each element agent needs to be matched to the corresponding dummy element agent in two of the three matchings. Moreover, each set agent from S_{set} needs to approve all three matchings from \mathcal{M}' and no dummy set agent b_c can be matched to an agent from S_{set} . Thus, each element agent is matched to a set agent it approves in exactly one of the three matchings. Since each set agent from S_{set} needs to approve all three matchings from \mathcal{M}' , this implies that each set agent from S_{set} needs to be matched to each of the three element agents corresponding to its elements in one of the three matchings. Thus, S_{set} forms an exact cover of X . This concludes the correctness proof of the reduction.

It is possible to slightly modify the reduction to avoid that Pareto-dominated matchings are part of the given committee, at the cost of losing symmetry. We start by modifying the approval ballots of $6q$ arbitrary filling agents and make them approve all element agents a_x for $x \in X$ and all dummy set agents b_c for $c \in C$ (but not the other way round). Constructing \mathcal{M} , instead of adding four matchings not approved by anyone, we add four matchings in which the $6q$ modified filling agents are matched to all element agents a_x and dummy set agents b_c . Note that these matchings are Pareto optimal, as modified filling agents only approve these agents and the remaining non-filling agents also only approve element agents a_x or dummy set agents b_c .

The correctness of the forward direction of the proof remains unaffected, while for the backward direction it is necessary to argue why none of the modified filling agents can be part of a blocking coalition. To see this, note that these agents approve four matchings in \mathcal{M} and thus any blocking coalition S they are part of needs to deserve to be represented by at least five matchings. However, this implies that $|S| \geq 5 \cdot \frac{16q}{3} > 26q$, which cannot be the case, as there exist only $24q$ agents approving some other agent. \square

7 Conclusion

We initiated the study of a multiagent problem at the intersection of social choice and matching theory: Given the 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, while we have focused on symmetric matching elections in our axiomatic study in Sect. 5, it would be interesting to extend this study to bipartite or general matching elections. In particular, it is open whether seq- w -Thiele satisfies any axiom in general (or bipartite) matching elections and whether seq-Phragmén satisfies EJR in bipartite (and symmetric) matching elections. Second, one could consider axioms 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. Note that both sequential Thiele rules and Equal Shares could fulfill this weaker version of core stability, as in the respective counterexamples for core stability from Proposition 1 and Proposition 3 the deviating group is not matched among itself in the deviation. Third, it would be natural to allow agents to rank-order potential matching partners and apply ordinal multiwinner voting procedures. Fourth, it would be interesting to identify other 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 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 (Lackner 2020).

Appendix A: Weighted approval winner for matching elections with stable candidate matchings

The following proposition implies that restricting ourselves to stable matchings in the committee in general matching elections has drastic implications on the computational tractability of the Weighted Approval Winner problem. For a motivation of this question consider our remark in Sect. 3.

Proposition 6 *Let N be a set of agents, $\omega : N \rightarrow \mathbb{Q}_{\geq 0}$ a weight function on the agents, $A_a \subseteq N \setminus \{a\}$ an approval ballot for each agent $a \in N$, and $\ell \in \mathbb{Q}_{\geq 0}$. Then, it is NP-hard to decide whether there is a matching M of the agents whose approvers have summed weight ℓ and where no two agents a and b exist that both disapprove M but approve each other. The hardness holds regardless of whether we require that M is Pareto optimal or not.*

Proof We reduce from a restricted NP-hard version of 3-SAT, where each variable occurs exactly twice positively and twice negatively (Berman et al. 2003). Given an instance of this problem consisting of a set $X = \{x_1, \dots, x_q\}$ of variables and a set $C = \{c_1, \dots, c_p\}$ of clauses, for each $i \in [q]$, let $c_{i,1}^+$ and $c_{i,2}^+$ denote the two clauses where x_i appears positively and $c_{i,1}^-$ and $c_{i,2}^-$ denote the two clauses where x_i appears negatively. We construct an instance of our problem as follows.

For each $j \in [p]$, we add a *clause agent* a_j . For each $i \in [q]$ we add two *positive variable agents* $b_{i,1}^+$ and $b_{i,2}^+$ and two *negative variable agents* $b_{i,1}^-$ and $b_{i,2}^-$. For each variable, the two corresponding positive variable agents approve the two negative variable agents and the other way round. Moreover, $b_{i,1}^+$ is approved by $a_{c_{i,1}^+}$, $b_{i,2}^+$ is approved by $a_{c_{i,2}^+}$, $b_{i,1}^-$ is approved by $a_{c_{i,1}^-}$, and $b_{i,2}^-$ is approved by $a_{c_{i,2}^-}$. Lastly, for each $i \in [q]$, we add two dummy agents $d_{i,1}$ and $d_{i,2}$. The dummy agents approve the respective four variable agents and the other way round. Concerning the weight function, we assign weight one to all clause agents and weight zero to all other agents. We set $\ell = p$. Thus, we search for a matching that is approved by all clause agents.

(\Rightarrow) Assume that there is a satisfying assignment α for (X, C) . We construct a matching M as follows. For each $j \in [p]$, pick some arbitrary literal appearing in c_j that is set to true by α . If x_i for some $i \in [q]$ was picked, then match the clause agent

a_j to the agent from $\{b_{i,1}^+, b_{i,2}^+\}$ which a_j approves. Otherwise, if \bar{x}_i for some $i \in [q]$ was picked, then match the clause agent a_j to the agent from $\{b_{i,1}^-, b_{i,2}^-\}$ which a_j approves. For each variable that is set to true by α , we match $b_{i,1}^-$ to $d_{i,1}$ and $b_{i,2}^-$ to $d_{i,2}$. For each variable that is set to false by α , we match $b_{i,1}^+$ to $d_{i,1}$ and $b_{i,2}^+$ to $d_{i,2}$. As α is a satisfying assignment, after this, each clause c_j is matched (to an agent it approves) and by construction each variable agent is matched at most once. Thus, M has weight $\ell = p$.

Observe that the only agents which disapprove the matching are the two positive variable agents for variables that are set to true by α and the two negative variable agents for variables that are set to false by α . None of these agents approve each other. Thus, as M has weight ℓ , M is a solution to the constructed instance.

(\Leftarrow) Assume that there is a matching M approved by all clause agents where no two agents a and b exist that both disapprove M but approve each other. The absence of such a pair of agents implies in particular that for each variable either only negative or positive variable agents are matched to a clause agent. If there is a variable such that both a positive and a negative variable agent are matched to a clause agent, then both of them disapprove M but approve each other, a contradiction. Let α be the assignment that sets variable x_i to true if at least one corresponding positive variable agent is matched to a clause agent in M , and to false otherwise. Then, by our above observation, as every clause agent is matched to a variable agent corresponding to a literal appearing in the clause, α is a satisfying assignment for (X, C) . \square

Acknowledgements This material is based on work supported by the Deutsche Forschungsgemeinschaft under grants NI 369/19, NI 369/22, and BR 4744/2-1. We thank Ágnes Cseh for helpful discussions.

Funding Open Access funding enabled and organized by Projekt DEAL.

Data Availability Not available.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Aziz H, Gaspers S, Gudmundsson J, Mackenzie S, Mattei N, Walsh T (2015a) Computational aspects of multi-winner approval voting. In: Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), IFAAMAS, pp 107–115
- Aziz H, Gaspers S, Mackenzie S, Walsh T (2015b) Fair assignment of indivisible objects under ordinal preferences. *Artif Intell* 227(2857):71–92
- Aziz H, Brill M, Conitzer V, Elkind E, Freeman R, Walsh T (2017) Justified representation in approval-based committee voting. *Soc Choice Welf* 48(2):461–485

- Aziz H, Elkind E, Huang S, Lackner M, Sánchez-Fernández S, Skowron P (2018) On the complexity of extended and proportional justified representation. In: Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI). AAAI Press, pp 902–909
- Aziz H, Brandl F, Brandt F, Harrenstein P, Olsen M, Peters D (2019) Fractional hedonic games. *ACM Trans Econ Comput* 7(2):6:1-6:29
- Berman P, Karpinski M, Scott AD (2003) Approximation hardness of short symmetric instances of MAX-3SAT. *Electronic Colloquium on Computational Complexity* (049)
- Boehmer N, Elkind E (2020) Stable roommate problem with diversity preferences. In: Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI), pp 96–102. <https://ijcai.org>
- Boehmer N, Brill M, Schmidt-Kraepelin U (2022) Proportional representation in matching markets: selecting multiple matchings under dichotomous preferences. In: Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS). IFAAMAS, pp 136–144
- Bogomolnaia A, Jackson MO (2002) The stability of hedonic coalition structures. *Games Econom Behav* 38(2):201–230
- Bogomolnaia A, Moulin H (2004) Random matching under dichotomous preferences. *Econometrica* 72(1):257–279
- Bouveret S, Endriss U, Lang J (2010) Fair division under ordinal preferences: computing envy-free allocations of indivisible goods. In: Proceedings of the 19th European Conference on Artificial Intelligence (ECAI). IOS Press, pp 387–392
- Brill M, Freeman R, Janson S, Lackner M (2017) Phragmén’s voting methods and justified representation. Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI), AAAI Press, pp 406–413
- Brill M, Laslier J-F, Skowron P (2018) Multiwinner approval rules as apportionment methods. *J Theor Polit* 30(3):358–382
- Brill M, Gözl P, Peters D, Schmidt-Kraepelin U, Wilker K (2022a) Approval-based apportionment. *Math Program*. <https://doi.org/10.1007/s10107-022-01852-1>
- Brill M, Israel J, Micha E, Peters J (2022b) Individual representation in approval-based committee voting. Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI), AAAI Press, pp 4892–4899
- Chamberlin JR, Courant PN (1983) Representative deliberations and representative decisions: proportional representation and the Borda rule. *Am Polit Sci Rev* 77(3):718–733
- Cheng Y, Jiang Z, Munagala K, Wang K (2019) Group fairness in committee selection. Proceedings of the 20th ACM Conference on Economics and Computation (ACM-EC). ACM, pp 263–279
- Chevalyre Y, Endriss U, Lang J, Maudet N (2008) Preference handling in combinatorial domains: from AI to social choice. *AI Mag* 29(4):37–46
- Cseh Á (2017) Popular matchings. In: Endriss U (ed) Trends in computational social choice, chapter 6. AI Access, pp 105–122
- Edmonds J (1965) Paths, trees and flowers. *Can J Math* 17:449–467
- Eisner MJ, Severance DG (1976) Mathematical techniques for efficient record segmentation in large shared databases. *J Assoc Comput Mach* 23(4):619–635
- Faliszewski P, Skowron P, Slinko A, Talmon N (2017) Multiwinner voting: a new challenge for social choice theory. In: Endriss U (ed) Trends in computational social choice, chap 2. AI Access, pp 27–47
- Gale D, Shapley LS (1962) College admissions and the stability of marriage. *Am Math Mon* 69(1):9–15
- Gallai T (1964) Maximale systeme unabhängiger kanten. *Publ Math Inst Hung Acad Sci Ser A* 9:401–413
- García-Soriano D, Bonchi F (2020) Fair-by-design matching. *Data Min Knowl Disc* 34:1291–1335
- Gärdenfors P (1975) Match making: assignments based on bilateral preferences. *Behav Sci* 20(3):166–173
- Garey M, Johnson D (1979) Computers and intractability: a guide to the theory of NP-completeness. W. H. Freeman and Company, New York
- Hall P (1935) On representatives of subsets. *J Lond Math Soc* 10:26–30
- Holyer I (1981) The NP-completeness of edge-colouring. *SIAM J Comput* 10(4):718–720
- Janson S (2016) Phragmén’s and Thiele’s election methods. Technical report, [arXiv:1611.08826](https://arxiv.org/abs/1611.08826) [math.HO]
- Jiang Z, Munagala K, Wang K (2020) Approximately stable committee selection. In: Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC). ACM, pp 463–472
- Kavitha T, Mestre J, Nasre M (2011) Popular mixed matchings. *Theoret Comput Sci* 412(24):2679–2690

- Korte B, Vygen J (2012) *Combinatorial optimization: theory and algorithms*, 5th edn. Springer
- Lackner M (2020) Perpetual voting: fairness in long-term decision making. *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*. AAAI Press, pp 2103–2110
- Lackner M, Skowron P (2022) Multi-winner voting with approval preferences. Technical report, [arXiv:2007.01795](https://arxiv.org/abs/2007.01795) [cs.GT]
- Lang J, Xia L (2016) Voting in combinatorial domains. In: Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (eds) *Handbook of computational social choice*, chap 9. Cambridge University Press, pp 197–222
- Manlove DF (2013) *Algorithmics of matching under preferences*. World Scientific Publishing Company
- Monroe BL (1995) Fully proportional representation. *Am Polit Sci Rev* 89(4):925–940
- Peters D, Skowron P (2020) Proportionality and the limits of welfarism. *Proceedings of the 21st ACM Conference on Economics and Computation (ACM-EC)*. ACM Press, pp 793–794
- Peters D, Pierczynski G, Shah N, Skowron P (2021) Market-based explanations of collective decisions. AAAI Press, pp 5656–5663
- Phragmén E (1894) Sur une méthode nouvelle pour réaliser, dans les élections, la représentation proportionnelle des partis. *Öfversigt af Kongliga Vetenskaps-Akademiens Förhandlingar* 51(3):133–137
- Roth AE, Rothblum UG, Vande Vate JH (1993) Stable matchings, optimal assignments, and linear programming. *Math Oper Res* 18(4):803–828
- Sánchez-Fernández L, Elkind E, Lackner M, Fernández N, Fisteus JA, Basanta Val P, Skowron P (2017) Proportional justified representation. In: *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*. AAAI Press, pp 670–676
- Thiele TN (1895) Om flerfold valg. *Oversigt over det Kongelige Danske Videnskabernes Selskabs Fordhandlinger*
- Valiant Leslie G (1979) The complexity of computing the permanent. *Theoret Comput Sci* 8:189–201

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.