

Algorithms for Structured Elections under Thiele Voting Rules

Alexandra Lassota¹, Krzysztof Sornat²

¹TU Eindhoven, the Netherlands

²AGH University, Poland

a.a.lassota@tue.nl, sornat@agh.edu.pl

Abstract

We study the computational complexity of winner determination problems in approval-based committee elections under Thiele voting rules. These form a class of rules parameterized by a fixed weight vector that specifies how a voter’s satisfaction depends on the number of approved candidates elected. We first analyze the structure of optimal solutions based on the sets of voters who approve each candidate—that is, how voters’ approval ballots induce dependencies between candidates—revealing constraints on a winning committee under any fixed Thiele voting rule. Using this, we design FPT algorithms for Proportional Approval Voting (PAV) and other Thiele rules on a natural restricted domain known as the Voter Interval (VI) domain—that is, after a suitable ordering of voters, each candidate is approved by a consecutive interval of voters. In particular, we show that every Thiele rule on VI is FPT with respect to a parameter for which the problem is NP-hard on general instances, even when the parameter takes constant values. Our results advance the understanding of the computational complexity of PAV on Voter Interval instances, which remains one of the central open questions in this area. We further resolve two open questions from the literature on PAV (and other Thiele voting rules) by providing a polynomial-time algorithm for instances where each candidate is approved by at most two voters, and an FPT algorithm parameterized by the total score of a winning committee.

1 Introduction

Multi-winner elections based on approval ballots are used in many settings, such as recommendation systems, committee selection, and blockchain (Skowron, Faliszewski, and Lang 2016; Lackner and Skowron 2023; Boehmer et al. 2024), where every voter expresses its preferences as a subset of candidates it approves of. A central family of voting rules for such applications are Thiele rules (Thiele 1895), which include Proportional Approval Voting (PAV) as a prominent member (Aziz et al. 2015). These rules balance diversity, proportionality, and excellence in the selected committee, depending on the specific Thiele rule used (Lackner and Skowron 2021). For instance, Chamberlin-Courant Approval Voting (CC) (Chamberlin and Courant 1983) promotes diversity by ensuring broad representation; PAV aims

for proportional representation, satisfying strong proportionality axioms like EJR+ (Brill and Peters 2023); Multi-winner Approval Voting focuses on excellence by selecting the most approved candidates. However, computing a winning committee under Thiele rules is computationally challenging. In fact, finding a winning committee under any non-trivial Thiele rule (i.e., every rule except for Multi-winner Approval Voting) is NP-hard. For CC and PAV, the problems remain NP-hard even under very restricted conditions where each candidate is approved by exactly 3 voters and each voter approves exactly 2 candidates (Aziz et al. 2015; Skowron, Faliszewski, and Lang 2016).

To better understand and overcome this computational hardness, a natural direction is to restrict the input domain. Such restrictions often enable polynomial-time algorithms for otherwise intractable rules, and this approach has been particularly fruitful in approval-based committee elections. Two central restricted domains for approval ballots are the *Candidate Interval (CI)* and the *Voter Interval (VI)* domains (Elkind, Lackner, and Peters 2017, 2025).

In the CI domain, candidates can be ordered so that each voter’s approval set forms a contiguous interval. CI captures scenarios where candidates are linearly ordered, e.g. by ideology or location, and each voter is focused on a specific region of the spectrum. One key advantage of CI is that it allows the winner determination problem for all Thiele rules to be solved efficiently. It is by using an integer linear programming formulation that admits a totally unimodular constraint matrix, which are known to be solvable in polynomial time (Peters 2018; Peters and Lackner 2020).

The VI domain, in contrast, imposes structure on the voters rather than on the candidates. Here, voters can be ordered so that each candidate is approved by a consecutive segment of voters. This domain models scenarios where voters are structured by demographic or socioeconomic features such as age, income, or education level, and candidates appeal to specific groups, e.g. young voters, low-income households, or university-educated individuals. The VI domain has received significant attention (Elkind, Lackner, and Peters 2017, 2025) and, consequently, for several voting rules, such as CC, Monroe’s and Minimax Approval Voting polynomial-time winner-determination algorithms for VI elections are known (Betzler, Slinko, and Uhlmann 2013; Liu and Guo 2016).

Yet, for most Thiele rules, and PAV in particular, the computational complexity of winner determination on VI remains a prominent open question (Elkind, Lackner, and Peters 2017; Peters 2018; Godziszewski et al. 2021; Lackner and Skowron 2023; Elkind, Lackner, and Peters 2025). Despite the dual nature of the VI and CI domains, the techniques used for CI do not translate to VI. In particular, the constraint matrix of the integer linear programming formulation is not totally unimodular anymore under VI preferences and, thus, it is not clear whether this problem can be solved in polynomial time or not. This motivates a search for new structural results and algorithmic techniques.

Structural Results. This work takes a new approach to understand the computational complexity of Thiele rules by studying the structure of winning committees. Instead of considering candidates in isolation, we examine how shared support among voters constrains the possible combinations of candidates in a winning committee. We provide a structural characterization of optimal committees under Thiele rules based on a *dominance* relation among candidates. Specifically, a candidate c is said to dominate another candidate d if the set of supporters of c strictly contains that of d . This induces a hierarchy of *dominancy levels* where candidates within the same level do not dominate one another and each is dominated by some candidate in a higher level. We show that there always exists a winning committee that is *non-dominated*, meaning that no member is dominated by any candidate outside the committee. These structural insights apply to general approval profiles and are of independent interest, offering new theoretical tools for analyzing and determining winners under Thiele rules. Practically, this structure can guide the development of more efficient algorithms and heuristics. Theoretically, it offers a new perspective on the open question of whether winner determination under PAV is polynomial-time solvable on VI.

Algorithmic Results. We show that our structural insights are particularly effective when combined with the VI property: candidates of a VI instance can be partitioned into parts such that candidates within a part influence only a limited number of voters. This allows to design a dynamic program over a sequence of such parts, exploiting the limited interaction. As a result, we obtain an FPT algorithm for PAV and, more generally, for any Thiele rule, when parameterized by two parameters combined: the maximum number of approvals received by a candidate, and the maximum number of approvals in a vote. Notably, the same parameterization is para-NP-hard in the general (unstructured) case—that is, the problem remains NP-hard even when both parameters are constants. This contrast highlights the algorithmic power of our structural approach and provides substantial progress toward resolving the open question on the complexity of PAV in the VI domain (Elkind, Lackner, and Peters 2017; Peters 2018; Lackner and Skowron 2023; Elkind, Lackner, and Peters 2025).

Beyond the VI setting, we further contribute to understanding tractable cases for Thiele rules by resolving two open problems posed by Yang and Wang (2018, 2023), originally asked for PAV: (1) we provide a polynomial-time al-

gorithm for instances where each candidate is approved by at most two voters, based on a proper integer linear programming formulation; and (2) we give an FPT algorithm parameterized by the total score of a winning committee, employing combinatorial tools like color-coding and splitters. It is particularly interesting as the score may not be integer. Moreover, this parameter can be smaller than the number of voters for which FPT algorithms exist. Importantly, our results extend to all Thiele rules, not just PAV, demonstrating the generality of our approach.

Structure of the Paper. Section 2 introduces notation for the approval-based committee election model, structured domains and voting rules studied in this paper. Section 3 presents our structural results on winning committees under Thiele rules. Section 4, using the structural results, develops an FPT algorithm on VI instances. In Section 5, we give an FPT algorithm parameterized by the total score of the optimal committee, and a polynomial-time algorithm for instances where each candidate is approved by at most two voters. Due to space constraints, some proofs are deferred to the full version of the paper.

1.1 Related Work

Recently, another FPT algorithm parameterized by the total score of a winning committee was provided independently by Gupta, Jain, Saha, Saurabh, and Upasana (2025). Even though they also use color-coding at the heart of their algorithm, the approaches differ significantly: while they color candidates and voters, we color the approvals of the voters. This results in a different number, type, and meaning of guesses, as well as a different construction of the overall solution. It also results in differences in running times. While the algorithms outperform each other in some cases, our algorithm runs in time truly linear in the number of voters, which we see as a strong advantage. For more detailed discussion about differences between both algorithms, as well as quantitative comparison of running times, we defer to the full version of the paper.

There is also a substantial body of work on tractability of voting rules under restricted domains in the ordinal setting (in which voters cast votes in a form of linear orders over candidates), such as Single-Peaked (SP) and Single-Crossing (SC) preferences. These domains are the ordinal counterparts of CI and VI, respectively, and have led to efficient algorithms for many voting rules, including CC (in its general version defined on cardinal values of misrepresentation) and Kemeny. In particular, a classic dynamic programming algorithm was proposed for CC by Betzler, Slinko, and Uhlmann (2013). More recently, an algorithm with near linear-time in the input size has been developed for the case where the SC-axis (ordering of voters) or the SP-axis (ordering of candidates) is explicitly given (Constantinescu and Elkind 2021; Sornat, Vassilevska Williams, and Xu 2022).

These works illustrate the general principle that structural properties of preferences can be algorithmically exploited. For more information about these and other restricted domains see, e.g., the works of Elkind, Lackner, and Peters (2017, 2025).

Parameterized Complexity of Thiele Rules. For an overview of the parameterized complexity of Thiele rules, we refer to the recent work by Yang and Wang (2023), who provide a comprehensive summary of known results (see Table 1 therein), along with new findings for several structural parameters and their combinations. Below, we briefly discuss the main results for standard parameters such as the number of voters n , the number of candidates m , and the committee size k , with an emphasis on the techniques used and known lower bounds.

A trivial brute-force algorithm checking total scores of all size- k subsets of candidates runs in FPT time with respect to m , namely $O^*(2^m)$. This is essentially optimal for every non-constant Thiele rule, as under the Exponential Time Hypothesis there is no $O^*(2^{o(m)})$ -time algorithm for this problem (Sornat, Vassilevska Williams, and Xu 2022).

A mixed integer linear program (MILP) presented by Faliszewski et al. (2018, Fig. 2) implies an FPT algorithm with respect to n for every Thiele rule.¹ The main idea is to define an integer variable for each candidate type (defined by its set of supporters; hence, there are at most 2^n types), which encodes how many candidates of a particular type are selected for the solution. Non-integral variables are forced to take integral values in the optimal solution, as first used by Bredebeck et al. (2015). The resulting running time is double-exponential, namely $O^*(2^{2^{O(n)}})$ (Bredebeck et al. 2020a). A single-exponential lower bound of $O^*(2^{o(n)})$ under the Exponential Time Hypothesis is known (Sornat, Vassilevska Williams, and Xu 2022), hence, there is still a significant gap remaining. A similar MILP idea has been applied in many contexts, e.g., for extensions of Thiele rules (Jain, Sornat, and Talmon 2020; Yang and Wang 2023), as well as for problems related to bribery and control in elections (Bredebeck et al. 2020b).

Every non-constant Thiele rule is also W[1]-hard with respect to the committee size (Aziz et al. 2015; Jain, Sornat, and Talmon 2020; Sornat, Vassilevska Williams, and Xu 2022).

Approximability of Thiele Rules. Thiele rules have also been studied from the perspective of approximability, where the goal is to compute a committee whose total score is close to optimal (Skowron, Faliszewski, and Lang 2016; Byrka, Skowron, and Sornat 2018). For a broad class of Thiele rules, tight polynomial-time approximation algorithms have been established (Dudycz et al. 2020; Barman, Fawzi, and Fermé 2021; Barman et al. 2022).

2 Preliminaries

We are given a set of candidates $C = \{c_1, c_2, \dots, c_m\}$ and a set of voters $V = \{v_1, v_2, \dots, v_n\}$. Each voter $v \in V$ expresses its preference in the form of an *approval set* $A_v \subseteq$

¹Even though ordinal ballots and top- k -counting rules are considered in the MILP (Faliszewski et al. 2018, Fig. 2), it is enough to replace $g_{m,k}(j)$ with w_j , where w_j is the j -th element of the Thiele sequence w , and to adjust the definition of $\mathcal{T}(S_i)$ to the set of candidates approved by voters from S_i and no voter from $V \setminus S_i$.

C , and the collection $A = (A_v)_{v \in V}$ is referred to as an *approval profile*. Any subset of C is called a *committee*. We write $\mathcal{W}_k := \{W \subseteq C : |W| = k\}$ to denote the set of all committees of size k . For a candidate $c \in C$, we denote by $V_c := \{v \in V : c \in A_v\}$ the *set of supporters* of c . We extend the notation to sets of candidates $C' \subseteq C$, i.e., $V_{C'} := \bigcup_{c \in C'} V_c$. We define $\Delta_C := \max_{c \in C} |V_c|$ as the maximum number of approvals given to a candidate and $\Delta_V := \max_{v \in V} |A_v|$ as the maximum number of approved candidates by a voter. An *approval-based committee (ABC) election* is a tuple $E = (C, V, A, k)$. A voting rule is a function taking an election as an input and outputs a set of *winning committees* of size k .

The w -Thiele voting rule (Thiele 1895) is parameterized by a non-increasing infinite sequence $w = (w_1, w_2, \dots)$, called a *Thiele sequence*. Given an election (C, V, A, k) , a committee $W \in \mathcal{W}_k$ is *optimal* under the w -Thiele rule if it maximizes the *total score*:

$$d := \text{score}_w(W) = \sum_{v \in V} \sum_{i=1}^{|A_v \cap W|} w_i,$$

over all committees of size k . The w -Thiele rule returns the set of all such optimal committees. The corresponding computational problem w -THIELE requires outputting a single optimal committee. Its decision variant asks whether there exists a committee of size k with a total score at least a given value.

All Thiele rules considered in the literature are defined with $w_1 = 1$ (see examples below). Dividing w by w_1 and obtaining $w_1 = 1$ does not affect either its set of optimal solutions or its approximability (Dudycz et al. 2020), but it affects the parameter d , the total score of an optimal solution that is studied in this paper. In particular, if we allow $w_1 \leq \frac{1}{k}$, then $d \leq \Delta_C \cdot k \cdot \frac{1}{k} = \Delta_C$, but every w -THIELE with non-constant w is NP-hard even if $\Delta_C = 3$ (Aziz et al. 2015; Skowron, Faliszewski, and Lang 2016), so this would imply paraNP-hardness with respect to d . Therefore, we use the standard normalization $w_1 = 1$ in this paper. (However any algorithm with running time dependent on d after such normalization becomes an algorithm with running time dependent on $d + \frac{1}{w_1}$.)

Arguably, three most prominent w -Thiele rules are:

- CHAMBERLIN-COURANT APPROVAL VOTING (CC), that is w -THIELE with $w = (1, 0, 0, \dots)$.
- PROPORTIONAL APPROVAL VOTING (PAV), that is w -THIELE with $w_j = \frac{1}{j}$.
- MULTI-WINNER APPROVAL VOTING (AV), that is w -THIELE with $w = (1, 1, 1, \dots)$.

An interpolation between CC and AV is ℓ -COVERAGE that is w -THIELE with $w_j = 1$ for $j \leq \ell$ and $w_j = 0$ otherwise (Barman et al. 2022).

GENERALIZED THIELE is a voting rule (Sornat, Vassilevska Williams, and Xu 2022) which takes as an input an election (C, V, A, k) and n Thiele sequences represented as $w: V \times \mathbb{N} \rightarrow [0, 1]$, where $(w_i^v)_{i \in \mathbb{N}}$ is a Thiele sequence for a voter $v \in V$, and outputs a committee $W \in \mathcal{W}_k$ which maximizes the total score: $d := \text{score}_w(W) =$

$\sum_{v \in V} \sum_{i=1}^{|A_i \cap W|} w_i^v$, over all committees of size k . Naturally, GENERALIZED THIELE where $w^v = w^{v'}$ for $v, v' \in V$ is equivalent to w^v -THIELE.

Definition 1 (Voter Interval). *An approval profile A has Voter Interval (VI) property if there exists a linear order of voters such that for every candidate $c \in C$, the set V_c is an interval on the linear order.*

An analogous restricted domain is defined for ordering of candidates.

Definition 2 (Candidate Interval). *An approval profile A has Candidate Interval (CI) property if there exists a linear order of candidates such that for every voter $v \in V$, the set A_v is an interval on the linear order.*

If an approval profile is VI, then a corresponding ordering of voters—called a *VI-axis*—can be found in polynomial time; thus, we assume w.l.o.g. that the ordering is (v_1, v_2, \dots, v_n) . An analogous ordering for CI profiles can be found also in polynomial time (Faliszewski et al. 2011; Elkind and Lackner 2015).

We write $[n] = \{1, 2, \dots, n\}$ and adopt standard notation from computational and parameterized complexity theory (Cygan et al. 2015). In particular, a decision problem parameterized by k is *fixed parameter tractable* (FPT) with respect to k if it can be solved in time $f(k) \cdot \text{poly}(|\mathcal{I}|)$ for any instance (\mathcal{I}, k) , where f is a computable function and $|\mathcal{I}|$ denotes the input size. A problem solvable in time $|\mathcal{I}|^{f(k)}$ belongs to the class XP, which implies it is solvable in polynomial time for any fixed value of k . The notation $O^*(\cdot)$ suppresses factors polynomial in the input size.

3 Structure of Winning Committees

For a given approval profile A , we create the *dominancy graph* of A where vertices correspond to candidates and a directed edge from c to c' exists if $V_{c'} \subset V_c$, i.e., a candidate c' is dominated by c (all supporters of c' are also supporters of c , and c has a supporter not supporting c'). This dominancy relation coincides with the one recently introduced independently by Dong et al. (2025).

Non-dominated candidates are those candidates that have no incoming edge in the dominancy graph, we denote this set as L_1 . We partition C into *dominancy levels* L_1, \dots, L_δ such that every candidate $c \in L_i$ is at distance exactly i in G from some candidate in L_1 and there is no candidate $c' \in L_1$ such that the distance in the dominancy graph between c' and c is strictly smaller than i . We denote by δ the *depth* of a dominancy graph. By the definition of dominancy relation involving strict inclusion, we have $\delta \leq \Delta_C + 1$.

Definition 3. *A committee $W \subseteq C$ is non-dominated if every candidate belonging to any directed path from an element of L_1 to an element of W also belongs to W .*

The next theorem characterizes optimal solutions to GENERALIZED THIELE with respect to non-dominancy.

Theorem 4. *There exists an optimal solution to GENERALIZED THIELE that is non-dominated. Furthermore, if $w_i^v > w_{i+1}^v$ for all $v \in V$ and all $i \in \mathbb{N}$, then every optimal solution to GENERALIZED THIELE must be non-dominated.*

Theorem 4 has several useful implications. To solve GENERALIZED THIELE, it suffices to consider only the first k dominancy levels. Indeed, for any $c \in L_{k+1} \cup \dots \cup L_\delta$ and any size- k committee containing c , there exists a candidate $c' \in L_1 \cup \dots \cup L_k$ outside the committee that dominates c .

Proposition 5. *There exists an optimal solution to GENERALIZED THIELE that is a subset of $L_1 \cup \dots \cup L_k$.*

For specific Thiele rules, the instances might be even further restricted. For example, in ℓ -COVERAGE, the score of a voter v from committee W equals $\min\{|A_v \cap W|, \ell\}$. This means that a voter can receive score at most ℓ , which implies that only dominancy levels L_1, \dots, L_ℓ are of interest because removing candidates from levels $L_{\ell+1}, \dots, L_\delta$ from a solution does not change its total score. In the case $|L_1 \cup \dots \cup L_\ell| < k$, by taking all candidates from $L_1 \cup \dots \cup L_\ell$ to the committee, we obtain a solution with the total score equal to $\text{score}_w(C)$ (we fill the remaining seats in the committee with arbitrary candidates). Hence, in the case of ℓ -COVERAGE, we may assume w.l.o.g. that $\delta \leq \ell$. In general, the above discussion implies the following.

Proposition 6. *There exists an optimal solution to GENERALIZED THIELE that is non-dominated and that is either a subset or a superset of $L_1 \cup \dots \cup L_\ell$, where $\ell = \arg \max_{i \in \mathbb{N}} (\exists v \in V (w_i^v > 0))$.*

Further structural results, specialized for the VI domain, are developed in the proof of Theorem 7.

4 FPT Results on Voter Interval

The following result shows that GENERALIZED THIELE on VI is FPT parameterized by $\Delta_C + \Delta_V$. Note that GENERALIZED THIELE on general instances is NP-hard even if $\Delta_C = 3$ and $\Delta_V = 2$, and the hardness persists even for PAV and for CC, which are special cases of GENERALIZED THIELE (Aziz et al. 2015; Skowron, Faliszewski, and Lang 2016). This shows that the problem is easier to solve in the VI domain assuming $P \neq NP$ (a standard assumption in computational complexity theory). This result progresses towards resolving the central open question on the complexity of PAV in the VI domain.

Theorem 7. *GENERALIZED THIELE on the Voter Interval domain is FPT parameterized by $\Delta_C + \Delta_V$ and is XP parameterized by Δ_C .*

Proof. Due to having VI preferences with respect to (v_1, \dots, v_n) , for every candidate $c \in C$, we have $V_c = \{v_i, v_{i+1}, \dots, v_j\}$ for some $1 \leq i \leq j \leq n$ (w.l.o.g., we assumed there are no candidates with an empty set of supporters). We denote by $\min(V_c)$ and $\max(V_c)$ the indices of the first and, respectively, the last supporting voters of candidate c .

We order the candidates (c_1, \dots, c_m) such that for every pair of indices $i < j$, one of the following conditions holds: (1) $\max(V_{c_i}) < \max(V_{c_j})$, or (2) $\max(V_{c_i}) = \max(V_{c_j})$ and $\min(V_{c_i}) < \min(V_{c_j})$.

For each $i \in [n]$, we define the set $C_i = \{c \in C : \max(V_c) = i\}$, called a *triangle*. Each triangle C_i consists of candidates with the same last supporting voter v_i .

Note that $\{C_1, \dots, C_n\}$ is a partition of C . The term *triangle* comes from the visual shape of the approval set of candidates in C_i when the approval profile is displayed as a matrix. In particular, after applying the above candidate ordering (c_1, \dots, c_m) , where candidates in C_i appear consecutively, and arranging the voters according to the VI-axis, the approvals form triangular patterns, as shown in Figure 1. We denote by $\min(C_i)$ and $\max(C_i)$ the indices of the first and, respectively, the last candidate included in C_i .

	c_1	c_2	c_3	c_4	c_5	c_6
v_1	✓					
v_2	✓	✓		✓		
v_3	✓	✓	✓	✓		
v_4				✓	✓	✓

Figure 1: An example of a VI approval profile, where the voters are ordered according to the VI-axis, and the candidates are ordered as described in the proof of Theorem 7. In this example, there are two non-empty triangles: $C_3 = \{c_1, c_2, c_3\}$ and $C_4 = \{c_4, c_5, c_6\}$. Triangle C_3 is associated with voter v_3 , meaning that each candidate in C_3 has v_3 as their last supporter. For instance, $c_1 \in C_3$ because $\max(V_{c_1}) = 3$. The collective set of supporters of C_4 is $V_{C_4} = \{v_2, v_3, v_4\}$. We have $\min(C_4) = 4$ and $\max(C_4) = 6$.

Due to Theorem 4, let us consider a fixed non-dominated solution W_{OPT} . It holds that $W_{\text{OPT}} \cap C_i$ is a subset of candidates from C_i having indices exactly $\{\min(C_i), \dots, j\}$ where $j = \min(C_i) + |W_{\text{OPT}} \cap C_i| - 1 \leq \max(C_i)$. In other words, any optimal solution contains a (possibly empty) *prefix* of every triangle C_i .

The number of all possible committees consisting of prefixes of the triangles might still be exponential in n . We solve the problem via a dynamic program. It iterates over some order of the triangles C_1, \dots, C_n and considers prefixes of C_i based on prefixes taken so far to a solution of a limited number (in terms of the parameters) of preceding triangles. A crucial observation is that candidates from triangle C_i have common supporters with at most $O(\Delta_V \cdot \Delta_C)$ many candidates outside of C_i . This allows for making “local” decisions when considering candidates from a triangle C_i , where the size of locality is defined in terms of the parameters.

More formally, let V_{C_i} be the set of voters supporting any candidate from C_i (hence, $|V_{C_i}|$ is the height of a triangle C_i). We observe that $|V_{C_i}| \leq \Delta_C$, as all members of C_i have the i -th voter as the last supporter and the maximum number of approvals received by a candidate is Δ_C . Hence, we can write $V_{C_i} \subseteq \{\max(V_{C_i}) - \Delta_C, \dots, \max(V_{C_i})\}$. We have $|C_i| \leq \Delta_V$, as every candidate in C_i is supported by v_i . Thus, overall, the size of every triangle is bounded in terms of the parameters of our instance.

We base our dynamic program table $T[a, b, d_1, \dots, d_{|V_{C_b}|}]$ on these observations: the table saves the maximum total score of GENERALIZED THIELE of a (non-dominated) committee of size a taken from the first b triangles (that is, from $C_1 \cup \dots \cup C_b$) such that the i -th

voter from V_{C_b} approves exactly d_i committee members. As our goal is finding a committee, we will also store (only) one committee that achieves a particular score in table T .

We initialize the table with the first triangle, that is, for $b = 1$, we fill exactly $\min\{|C_1|, k\} + 1$ entries of $T[a, 1, d_1, \dots, d_{|V_{C_1}|}]$ by considering committees $C_1(a)$ being a prefix of C_1 of size a . Formally, $C_1(0) = \emptyset$ and $C_1(a) = \{c_{\min(C_1)}, \dots, c_{\min(C_1)+a-1}\}$ for $a \in \mathbb{N} : a \leq \min\{|C_1|, k\}$. By $d_i(a)$ we denote the number of approved candidates by i -th voter from V_{C_1} in a committee $C_1(a)$. We store the total score achieved by a committee $C_1(a)$ in a corresponding entry of $T[a, 1, d_1(a), \dots, d_{|V_{C_1}|}(a)]$.

When iterating over $b > 1$, we will consider committees $C_b(a)$ for $0 \leq a \leq \min\{|C_b|, k\}$ being merged with every committee stored in non-empty entries created in the previous step, i.e., when considering solution up to $(b - 1)$ -th triangle. Formally, we take a committee W_{temp} from every non-empty cell of $T[|W_{\text{temp}}|, b - 1, d_1^{\text{temp}}, \dots, d_{|V_{C_{b-1}}|}^{\text{temp}}]$ and

we create a new committee $W_{\text{temp}} \cup C_b(a)$ for every a such that $0 \leq a \leq \min\{|C_b|, k - |W_{\text{temp}}|\}$. Note that in this way $|W_{\text{temp}} \cup C_b(a)| \leq k$. Now, we compare the score of $W_{\text{temp}} \cup C_b(a)$ with a score of a committee of respective entry of T , i.e., $T[|W_{\text{temp}}| + a, b, d_1(a, W_{\text{temp}}), \dots, d_{|V_{C_b}|}(a, W_{\text{temp}})]$, where $d_i(a, W_{\text{temp}})$ is the number of approved candidates by the i -th voter from V_{C_b} in the committee $W_{\text{temp}} \cup C_b(a)$. If $W_{\text{temp}} \cup C_b(a)$ achieves a strictly higher score, i.e. $\text{score}(W_{\text{temp}} \cup C_b(a)) > T[|W_{\text{temp}}| + a, b, d_1(a, W_{\text{temp}}), \dots, d_{|V_{C_b}|}(a, W_{\text{temp}})]$, then we update the entry with this higher score, and we store $W_{\text{temp}} \cup C_b(a)$ as a committee achieving this score. We note that the procedure is well-defined also in the case of $C_b = \emptyset$ because the only prefix of C_b considered will be an empty set. The formal recurrence is provided in the full version of the paper.

To obtain the maximum total score committee, we search for the largest value among non-empty entries of $T[k, n, d_1, \dots, d_{|V_{C_n}|}]$. The correctness of the dynamic program follows from the fact that every such entry stores a valid committee of size k , and, as argued before, the optimal committee W_{OPT} satisfies $W_{\text{OPT}} \cap C_i = C_i(|W_{\text{OPT}} \cap C_i|)$, i.e., W_{OPT} consists of prefixes of the triangles C_i . This implies that, for every $b \in [n]$, the table entry $T[|W_{\text{OPT}} \cap (\cup_{i \in [b]} C_i)|, b, d_1(b), \dots, d_{|V_{C_b}|}(b)]$ stores a committee whose total score is at least $\text{score}_w(W_{\text{OPT}} \cap \cup_{i \in [b]} C_i)$ because the relevant triangle prefixes $C_i(|W_{\text{OPT}} \cap C_i|)$ (or any other prefixes yielding the same score and satisfying each voter in V_{C_b} by the same number of approved committee members) were explicitly considered during the construction of this entry. Together with the base case, i.e., filling the entries $T[a, 1, d_1, \dots, d_{|V_{C_1}|}]$, this completes an inductive argument for the correctness of the algorithm.

The size of T is at most $O(k \cdot n \cdot \Delta_V^{\Delta_C})$. In order to fill all entries for particular $b \in [n]$, we consider at most $k + 1$ prefixes of b -th triangle merged with every committee stored in a non-empty cell of T with $b - 1$. Therefore, we consider at most $O(k^2 \cdot \Delta_V^{\Delta_C})$ many committees, each in polynomial time. In total, this gives a running time of $O^*(\Delta_V^{\Delta_C})$ which is FPT with respect to $\Delta_C + \Delta_V$. \square

A consequence from the above presented proof is an FPT algorithm parameterized by $\Delta_C + k$. That is, the indices of T , which encode how many times every voter involved is represented, never exceed the committee size k .

Corollary 8. GENERALIZED THIELE *on the Voter Interval domain can be solved in $O^*(k^{\Delta_C})$ time.*

We note that Theorem 14 provides an FPT algorithm parameterized by $\Delta_C + k$ for every instance, but its running time is double-exponential on k in contrast to the non-exponential dependence on k in the case of VI (Corollary 8). It is another example of parameterization for which much more efficient algorithms exist for the VI structured domain.

5 FPT Results on General Instances

In the next two subsections, we present two FPT algorithms for GENERALIZED THIELE which answer affirmatively two open questions known in the literature (Yang and Wang 2018, 2023), which were asked for a special case of PAV. In Section 5.1 we provide a polynomial-time algorithm for instances with $\Delta_C = 2$. In Section 5.2 we give an FPT algorithm parameterized by the total score of an optimal committee.

5.1 Polynomial-Time Algorithm when $\Delta_C = 2$

Our polynomial-time algorithm for GENERALIZED THIELE with $\Delta_C = 2$ is based on a generalization of an integer linear program (ILP) studied by Peters (2018). The ILP formulation of Peters (2018) is defined for any w -THIELE, but it can be easily adjusted to GENERALIZED THIELE by modifying the objective function (Sornat, Vassilevska Williams, and Xu 2022). For a given election (C, V, A, k) the ILP for GENERALIZED THIELE (ILP-GT) is defined as follows, where roughly speaking, the y_c variables mark selected candidates, the $x_{v,i}$ variables track voter satisfaction levels and the objective ensures that the highest available values in a Thiele sequence are always chosen:

$$\text{maximize } \sum_{v \in V} \sum_{i \in [k]} w_i^v \cdot x_{v,i} \quad (\text{ILP-GT})$$

$$\text{subject to } \sum_{c \in C} y_c = k \quad (1)$$

$$\sum_{i \in [k]} x_{v,i} \leq \sum_{c \in A_v} y_c \quad \forall v \in V \quad (2)$$

$$\begin{aligned} x_{v,i} &\in \{0, 1\} & \forall v \in V, i \in [k] \\ y_c &\in \{0, 1\} & \forall c \in C \end{aligned}$$

Peters (2018) showed that the constraint matrix of ILP-GT is totally unimodular (TU) when an approval profile is CI, hence an optimal solution can be found in polynomial time. He actually argued that if an approval profile A captured as a matrix, one can find an optimal solution in polynomial time even when an additional row with all-1s is appended (corresponding to Constraint (1)), as the resulting matrix is still TU. In the case of A being CI, total unimodularity was an immediate implication from the fact that A

being CI has consecutive one property, and because the additional row consists of only 1s (hence, it is consistent with the consecutive 1s property).

In the case where A is VI, total unimodularity is not necessarily preserved. While the transpose A^T satisfies the consecutive 1s property and is thus TU, this property may be lost when appending an all-ones row to A . Specifically, the matrix $\begin{bmatrix} \mathbf{1}_m \\ A \end{bmatrix}$ obtained by adding a row of 1s of length m to A —does not, in general, have the consecutive 1s property in its transpose. There exists a small VI approval profile that yields a non-TU matrix; see Example (3), where a VI profile with 4 candidates and 3 voters together with the cardinality constraint (the first row) is not TU (its determinant is -2).

$$\begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 \end{bmatrix} \quad (3)$$

Example (3) extends to Example (4) with $2m$ candidates and $2m - 1$ voters for any $m \geq 2$, whose determinant is $2 - 2m \leq -2$. Hence, these matrices are not TU either.

$$\begin{bmatrix} 1 & \mathbf{1}_{2m-1} \\ (\mathbf{1}_{2m-1})^T & I_{2m-1} \end{bmatrix} \quad (4)$$

However, there are more classes of integer programs that can be solved in polynomial time. In particular, if the coefficients of the constraint matrix are in $\{-2, -1, 0, 1, 2\}$ and the sum of absolute values is at most 2 for each column, then the problem is polynomial time solvable (Schrijver 2003). Such matrices are called *generalized matching matrices*.

Theorem 9 (Schrijver 2003). *ILPs with a generalized matching matrix can be solved in strongly polynomial time.*

These matrices correspond to problems called (generalized) matching problems, hence the name. The corresponding ILPs capture a variety of well-known problems in polynomial time such as minimum cost flow, minimum cost (b -)matching and certain graph factor problems (Schrijver 2003). These structures even remain FPT time solvable parameterized by the number p of additional columns (Lassota and Ligthart 2025). However, we are facing the problem of having an additional row, which is shown to be, in general, not FPT assuming FPT is not equal to $W[1]$ (a common hypothesis in parameterized complexity similar to P versus NP) (Lassota and Ligthart 2025). Fortunately, we can alter ILP-GT slightly and obtain a generalized matching matrix.

Before we show the proof, note that if there is a candidate that is not approved by any voter, we can delete this candidate from the instance. If k is larger than the remaining candidates, we take all those remaining candidates and greedily fill up our committee with any candidate from the discarded list to obtain an optimal solution. If there is at least one candidate that is approved by only one voter, then we add a new dummy voter v_d that approves exactly all of candidates with only one approval. This voter v_d will have a zero contribution to the objective function. Hence, in the following, we assume w.l.o.g. that all candidates are approved by exactly two voters.

Theorem 10. GENERALIZED THIELE with $\Delta_C = 2$ can be solved in polynomial time.

Proof. We modify the formulation of ILP-GT as follows. First, we replace Constraint (1) by the following equality:

$$\sum_{v \in V} \sum_{i \in [k]} x_{v,i} = 2k. \quad (5)$$

Second, we strengthen Constraint (2) by replacing the inequality with equality:

$$\sum_{i \in [k]} x_{v,i} = \sum_{c \in A_v} y_c \quad \forall v \in V. \quad (6)$$

Constraint (1) previously ensured that exactly k candidates are selected. As argued above, each candidate has exactly two supporters, so choosing any candidate (represented by variable y_c) implies that two variables $x_{v,i}$ must be set to 1 (one for each approving voter) for valid solutions. Thus, replacing (1) by (5) and enforcing equality in Constraint (6) shifts the responsibility for enforcing the committee size from the y_c variables to the $x_{v,i}$ variables.

After these modifications, each variable $x_{v,i}$ and y_c (each corresponding to one column of the constraint matrix) appears in exactly 2 constraints with coefficients in $\{-1, 1\}$ (and 0 everywhere else). This implies that the resulting constraint matrix is a generalized matching matrix. By applying Theorem 9, we obtain a solution in polynomial time. \square

5.2 FPT Algorithm Parameterized by the Score

Next, we prove that GENERALIZED THIELE is FPT parameterized by $k + \Delta_C$. The result relies on reducing the problem to a set cover variant called p -partial set cover that is solvable efficiently.

In the p -partial set cover problem, we are given a universe U of t elements, and a set of subsets \mathcal{S} of size s . The goal is to cover at least p different elements of U using the minimal number of sets in \mathcal{S} . This problem is FPT time solvable parameterized by p .

Theorem 11 (Bläser 2003). *A minimum p -partial set cover can be computed in time $2^{O(p)} \cdot s \cdot t$.*

Note that this algorithm also solves the weighted p -partial set cover problem where each set in \mathcal{S} is assigned a weight and the goal is to find the minimum weight set cover hitting at least p distinct elements of U .

Our algorithm reduces GENERALIZED THIELE to p -partial set cover with p bounded by a function of k and Δ_C . We view each voter as an element of the universe and every candidate as the set of supporters. Contrary to the set cover problem, in GENERALIZED THIELE, it can be beneficial to “cover” a voter multiple times. To implement this in our reduction, we use color coding, a powerful tool to design FPT time algorithms (see, e.g., the book of Cygan et al. (2015)). Color coding is used to color solution candidates such that, with high probability, an optimal solution only takes one solution candidate per color. We use the colors to color all approvals where each color indicates the additive contribution to the valuation function. This allows us to create different

elements not just for every approval, but also the way it contributes to the objective function, e.g., the number of times a voter has been “covered”.

To derandomize color coding, splitters have been introduced (Naor, Schulman, and Srinivasan 1995). Splitters compute a number of colorings instead of a single one such that at least one coloring has the desired property that an optimal solution only takes one solution candidate per color. This can be expressed in terms of hash functions:

Definition 12. *An (n, k, ℓ) splitter is a family of hash functions F from $\{1, 2, \dots, n\}$ to $\{1, 2, \dots, \ell\}$ such that for every $S \subseteq \{1, 2, \dots, n\}$ with $|S| = k$, there exists a function $f \in F$ that splits S evenly; that is, for every $j, j' \leq \ell$, we have $|f^{-1}(j) \cap S|$ and $|f^{-1}(j') \cap S|$ differ by at most 1.*

Lemma 13 (Naor, Schulman, and Srinivasan 1995). *There exists an (n, k, k) splitter of size $e^k k^{O(\log(k))} \log(n)$ which is computable in time $e^k k^{O(\log(k))} \cdot n \log(n)$.*

Equipped with splitters, we can finally present the main result of this section.

Theorem 14. GENERALIZED THIELE is FPT parameterized by $k + \Delta_C$.

This yields the following FPT algorithm.

Proposition 15. *For every Thiele sequence w , w -THIELE can be solved in time $2^{d^{O(d)}} \cdot nm^{O(1)}$, hence it is FPT parameterized by d , the total score of an optimal solution.*

6 Conclusion and Future Work

We presented new algorithms for computing optimal committees under Thiele rules. We identified structural properties of optimal solutions, which on Voter Interval instances enable a dynamic programming approach over a chain of subsets of candidates. We also resolved an open problem by showing that winner determination under any Thiele rule is polynomial-time solvable when each candidate is approved by at most two voters, using an ILP-based approach. Furthermore, we provided an FPT algorithm parameterized by $k + \Delta_C$ using color-coding technique, which we apply to obtain an FPT algorithm parameterized by the total score d .

Several directions remain open. The most prominent question is whether winner determination under PAV on Voter Interval instances is polynomial-time solvable or NP-hard. Another question is whether the ILP-based result for $\Delta_C = 2$ can be replaced by a purely combinatorial algorithm—a question that also arises in the context of Thiele rules on Candidate Interval profiles (Peters 2018).

It is natural to ask whether Theorem 7 extends to more general preference domains such as Voter-Candidate Interval (Dong et al. 2025; Elkind et al. 2024; Godziszewski et al. 2021). In contrast to (Dong et al. 2025), whose objective allows discarding dominated candidates, dominated candidates crucially affect Thiele scores.

Finally, our structural results and algorithms could be used to reason about tied committees, for example by analyzing possible and necessary winners under Thiele rules, since our methods that compute optimal scores and can handle preselected candidates (by adjusting Thiele sequences in GENERALIZED THIELE).

Acknowledgements

We thank Andrei Constantinescu and Piotr Faliszewski for discussions of our results that helped us clarify and improve their presentation. We also thank the anonymous reviewers for their valuable feedback.

Alexandra Lassota was supported by the Dutch Research Council (NWO) under project number VI.Veni.242.293. Krzysztof Sornat was supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 101002854).

References

- Aziz, H.; Gaspers, S.; Gudmundsson, J.; Mackenzie, S.; Mattei, N.; and Walsh, T. 2015. Computational Aspects of Multi-Winner Approval Voting. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multi-agent Systems (AAMAS 2015)*, 107–115.
- Barman, S.; Fawzi, O.; and Fermé, P. 2021. Tight Approximation Guarantees for Concave Coverage Problems. In *Proceedings of the 38th International Symposium on Theoretical Aspects of Computer Science (STACS 2021)*, 9:1–9:17.
- Barman, S.; Fawzi, O.; Ghoshal, S.; and Gürpınar, E. 2022. Tight Approximation Bounds for Maximum Multi-coverage. *Math. Program.*, 192(1): 443–476.
- Betzler, N.; Slinko, A.; and Uhlmann, J. 2013. On the Computation of Fully Proportional Representation. *J. Artif. Intell. Res.*, 47: 475–519.
- Bläser, M. 2003. Computing Small Partial Coverings. *Inf. Process. Lett.*, 85(6): 327–331.
- Boehmer, N.; Brill, M.; Cevallos, A.; Gehrlein, J.; Fernández, L. S.; and Schmidt-Kraepelin, U. 2024. Approval-Based Committee Voting in Practice: A Case Study of (over)Representation in the Polkadot Blockchain. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI 2024)*, 9519–9527.
- Bredereck, R.; Faliszewski, P.; Knop, A. K. R. D.; and Niedermeier, R. 2020a. Parameterized Algorithms for Finding a Collective Set of Items. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI 2020)*, 1838–1845.
- Bredereck, R.; Faliszewski, P.; Niedermeier, R.; Skowron, P.; and Talmon, N. 2015. Elections with Few Candidates: Prices, Weights, and Covering Problems. In *Proceedings of the 4th Conference on Algorithmic Decision Theory (ADT 2015)*, 414–431.
- Bredereck, R.; Faliszewski, P.; Niedermeier, R.; Skowron, P.; and Talmon, N. 2020b. Mixed Integer Programming with Convex/Concave Constraints: Fixed-Parameter Tractability and Applications to Multicovering and Voting. *Theor. Comput. Sci.*, 814: 86–105.
- Brill, M.; and Peters, J. 2023. Robust and Verifiable Proportionality Axioms for Multiwinner Voting. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC 2023)*, 301.
- Byrka, J.; Skowron, P.; and Sornat, K. 2018. Proportional Approval Voting, Harmonic k-Median, and Negative Association. In *Proceedings of the 45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*, 26:1–26:14.
- Chamberlin, J. R.; and Courant, P. N. 1983. Representative Deliberations and Representative Decisions: Proportional Representation and the Borda Rule. *Am. Political Sci. Rev.*, 77: 718–733.
- Constantinescu, A. C.; and Elkind, E. 2021. Proportional Representation under Single-Crossing Preferences Revisited. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI 2021)*, 5286–5293.
- Cygan, M.; Fomin, F. V.; Kowalik, Ł.; Lokshantov, D.; Marx, D.; Pilipczuk, M.; Pilipczuk, M.; and Saurabh, S. 2015. *Parameterized Algorithms*. Springer.
- Dong, C.; Bullinger, M.; Waş, T.; Birnbaum, L.; and Elkind, E. 2025. Selecting Interlacing Committees. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025)*, 630–638.
- Dudycz, S.; Manurangsi, P.; Marcinkowski, J.; and Sornat, K. 2020. Tight Approximation for Proportional Approval Voting. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI 2020)*, 276–282.
- Elkind, E.; Faliszewski, P.; Igarashi, A.; Manurangsi, P.; Schmidt-Kraepelin, U.; and Suksompong, W. 2024. The Price of Justified Representation. *ACM Trans. Economics and Comput.*, 12(3): 11:1–11:27.
- Elkind, E.; and Lackner, M. 2015. Structure in Dichotomous Preferences. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI 2015)*, 2019–2025.
- Elkind, E.; Lackner, M.; and Peters, D. 2017. Structured Preferences. In Endriss, U., ed., *Trends in Computational Social Choice*, chapter 10, 187–207. AI Access.
- Elkind, E.; Lackner, M.; and Peters, D. 2025. Preference Restrictions in Computational Social Choice: A Survey. *CoRR*, abs/2205.09092v2.
- Faliszewski, P.; Hemaspaandra, E.; Hemaspaandra, L. A.; and Rothe, J. 2011. The Shield that Never was: Societies with Single-peaked Preferences are More Open to Manipulation and Control. *Inf. Comput.*, 209(2): 89–107.
- Faliszewski, P.; Skowron, P.; Slinko, A.; and Talmon, N. 2018. Multiwinner Analogues of the Plurality Rule: Axiomatic and Algorithmic Perspectives. *Soc. Choice Welf.*, 51(3): 513–550.
- Godziszewski, M. T.; Batko, P.; Skowron, P.; and Faliszewski, P. 2021. An Analysis of Approval-Based Committee Rules for 2D-Euclidean Elections. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI 2021)*, 5448–5455.
- Gupta, S.; Jain, P.; Saha, S.; Saurabh, S.; and Upasana, A. 2025. More Efforts Towards Fixed-Parameter Approximability of Multiwinner Rules. *CoRR*, abs/2505.12699.
- Jain, P.; Sornat, K.; and Talmon, N. 2020. Participatory Budgeting with Project Interactions. In *Proceedings of the 29th*

International Joint Conference on Artificial Intelligence (IJ-CAI 2020), 386–392.

Lackner, M.; and Skowron, P. 2021. Consistent Approval-based Multi-winner Rules. *J. Econ. Theory*, 192: 105173.

Lackner, M.; and Skowron, P. 2023. *Multi-Winner Voting with Approval Preferences*. Springer.

Lassota, A.; and Lighthart, K. 2025. Parameterized Algorithms for Matching Integer Programs with Additional Rows and Columns. In *Proceedings of the 52nd International Colloquium on Automata, Languages, and Programming (ICALP 2025)*, 112:1–112:18.

Liu, H.; and Guo, J. 2016. Parameterized Complexity of Winner Determination in Minimax Committee Elections. In *Proceedings of the 2016 International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016)*, 341–349.

Naor, M.; Schulman, L. J.; and Srinivasan, A. 1995. Splitters and Near-Optimal Derandomization. In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science (FOCS 1995)*, 182–191.

Peters, D. 2018. Single-Peakedness and Total Unimodularity: New Polynomial-Time Algorithms for Multi-Winner Elections. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI 2018)*, 1169–1176.

Peters, D.; and Lackner, M. 2020. Preferences Single-Peaked on a Circle. *J. Artif. Intell. Res.*, 68: 463–502.

Schrijver, A. 2003. *Combinatorial Optimization: Polyhedra and Efficiency*, volume 24 of *Algorithms and Combinatorics*. Springer.

Skowron, P.; Faliszewski, P.; and Lang, J. 2016. Finding a Collective Set of Items: From Proportional Multirepresentation to Group Recommendation. *Artif. Intell.*, 241: 191–216.

Sornat, K.; Vassilevska Williams, V.; and Xu, Y. 2022. Near-Tight Algorithms for the Chamberlin-Courant and Thiele Voting Rules. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI 2022)*, 482–488.

Thiele, T. 1895. Om Flerfoldsvalg. In *Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandlinger (in Danish)*, 415–441. København: A.F. Høst.

Yang, Y.; and Wang, J. 2018. Parameterized Complexity of Multi-winner Determination: More Effort Towards Fixed-Parameter Tractability. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2018)*, 2142–2144.

Yang, Y.; and Wang, J. 2023. Parameterized Complexity of Multiwinner Determination: More Effort Towards Fixed-Parameter Tractability. *Auton. Agents Multi Agent Syst.*, 37(2): 28.