

Tight Approximation for Proportional Approval Voting

Szymon Dudycz¹, Pasin Manurangsi², Jan Marcinkowski¹ and Krzysztof Sornat^{1,3}

¹University of Wrocław, Poland

²Google Research, US

³Ben-Gurion University, Israel

szymon.dudycz@cs.uni.wroc.pl, pasin@google.com, {jan.marcinkowski, krzysztof.sornat}@cs.uni.wroc.pl

Abstract

In approval-based multiwinner elections, we are given a set of voters, a set of candidates, and, for each voter, a set of candidates approved by the voter. The goal is to find a committee of size k that maximizes the total utility of the voters. In this paper, we study approximability of Thiele rules, which are known to be NP-hard to solve exactly. We provide a tight polynomial time approximation algorithm for a natural class of geometrically dominant weights that includes such voting rules as Proportional Approval Voting or p-Geometric. The algorithm is relatively simple: first we solve a linear program and then we round a solution by employing a framework called pipage rounding due to Ageev and Sviridenko (2004) and Calinescu et al. (2011). We provide a matching lower bound via a reduction from the Label Cover problem. Moreover, assuming a conjecture called Gap-ETH, we show that better approximation ratio cannot be obtained even in time $f(k) \cdot \text{poly}(n, o(k))$.

1 Introduction

We study multiwinner approval-based voting rules called Thiele rules [Brill et al., 2018] that were introduced by Thorvald Thiele [1895].¹ Versions of Thiele rules have been used for distribution of seats within parties in Swedish parliamentary elections [Janson, 2016]. In recent years, Thiele rules regained interest in social choice community as a class of natural voting rules generalizing such well known multiwinner voting rules as Approval Chamberlin-Courant [Chamberlin and Courant, 1983; Procaccia et al., 2008; Skowron and Faliszewski, 2017] or Proportional Approval Voting [Thiele, 1895; Kilgour, 2010; Aziz et al., 2015; Byrka et al., 2018].

Thiele rules. w -THIELE rule is parameterized by a non-increasing non-negative sequence $\mathbf{w} = (w_1, w_2, \dots)$. Utility of a voter from a committee is increasing with the number of approved winners and it obeys the law of diminishing marginal utility known in economics (also called Gossen’s First Law). The increase is defined by the weights in \mathbf{w} .

¹The original paper was written in Danish, but there is a detailed survey in English [Janson, 2016].

Specifically, utility of a voter from a committee in which he or she approves j candidates is equal to $\sum_{i=1}^j w_i$. w -THIELE rule is purely utilitarian: the goal is to find a committee of k candidates that maximizes the total utility of the voters under the sequence of weights \mathbf{w} .

When \mathbf{w} is a constant sequence, the problem trivially admits a polynomial time algorithm. Such a voting rule is called Single Non-Transferable Vote (SNTV) [Faliszewski et al., 2018] and it is a generalization of a single-winner voting rule called Plurality Voting into multiwinner setting (every voter approves only the most favorite candidate, and the winners are k candidates with the highest support). All other Thiele rules are NP-hard [Skowron et al., 2016, Theorem 5]. Hence, they have been studied in terms of approximability [Skowron et al., 2016], parameterized complexity [Bredereck et al., 2020] and specific sequences of weights were considered [Kilgour, 2010; Aziz et al., 2015; Skowron and Faliszewski, 2017; Lackner and Skowron, 2019].

Related works. The simplest w -THIELE rule has $\mathbf{w} = (1, 0, 0, \dots)$ which is known as the Max k -Coverage problem (or Approval Chamberlin-Courant in social choice theory [Procaccia et al., 2008; Skowron and Faliszewski, 2017]). Nemhauser et al. [1978] provided a greedy algorithm that finds $(1 - 1/e)$ -approximate solution for Max k -Coverage. This has been shown to be tight unless $P = NP$ [Feige, 1998]. Recently, Cohen-Addad et al. [2019] and Manurangsi [2020] showed this is tight even in FPT time with respect to k . In particular, the latter showed that, assuming the Gap Exponential Time Hypothesis (Gap-ETH)², any $(1 - 1/e + \epsilon)$ -approximation algorithm for Max k -Coverage must run in time at least $(n+m)^{\Omega_\epsilon(k)}$. This means that a brute-force algorithm checking all possible subsets of size k (runs in $O^*(m^k)$) is essentially best possible, if one insists on achieving any approximation ratio better than $(1 - 1/e)$.

Recently, Barman et al. [2020] gave an $(1 - \frac{\ell^\ell}{e^{\ell \ell}})$ -approximation algorithm for the Multi ℓ -Coverage problem which is tight assuming the Unique Games Conjecture. Multi ℓ -Coverage is equivalent to w -THIELE rule with $w_\ell(i) =$

²Gap-ETH states that, for some constant $\delta > 0$, there is no $2^{o(n)}$ -time algorithm that, given n -variable 3-SAT formula, can distinguish whether the formula is fully satisfiable or that it is not even $(1 - \delta)$ -satisfiable. Gap-ETH is a standard assumption in proving FPT hardness of approximation (see e.g. [Chalermsook et al., 2017]).

$(1, 1, \dots, 1, 0, 0, \dots)$ (i.e., ℓ ones followed by zeros). In fact, Barman et al.’s result holds also for general utilities (not only approval based). This means they provide a tight improved approximation for the ℓ -best OWA problem [Skowron et al., 2016] (improving previous best $(1 - 1/e)$ -approximation).

Our contribution. Our main result (Theorem 1) is a tight $\alpha_{\mathbf{w}}$ -approximation algorithm for a natural class of “geometrically dominant” \mathbf{w} -THIELE rules, where $\alpha_{\mathbf{w}}$ is a sequence-dependent constant defined in (1). Resulting approximation ratios for selected \mathbf{w} -THIELE rules are presented in Table 1.

Theorem 1. *Let \mathbf{w} be any geometrically dominant sequence such that $\lim_{i \rightarrow \infty} w_i = 0$. There exists a polynomial time $\alpha_{\mathbf{w}}$ -approximation algorithm for \mathbf{w} -THIELE rule, and it is NP-hard to compute $(\alpha_{\mathbf{w}} + \epsilon)$ -approximate solution for any constant $\epsilon > 0$. Furthermore, assuming Gap-ETH, we cannot achieve $(\alpha_{\mathbf{w}} + \epsilon)$ -approximation even in $f(k) \cdot n^{o(k)}$ -time.*

The algorithm is a simple application of a well known linear programming rounding technique called *pipage rounding* due to Ageev and Sviridenko [2004]. Our contribution is an analysis of this algorithm for \mathbf{w} -THIELE rules (Section 3). Pipage rounding gives a better ratio than $(1 - 1/e)$ for a large class of \mathbf{w} -THIELE rules with \mathbf{w} being “geometrically dominant” (intuitively, this means that \mathbf{w} “decreases slower than” a geometric sequence; for a formal definition, please refer to Section 2). Our algorithm can be extended by replacing cardinality constraint (choosing k winners) into matroid constraint due to Theorem 2. Our proof of the hardness of approximation (Section 4) largely follows the proof of $(1 - 1/e + o(1))$ -inapproximability of Max k -Coverage of Feige [1998] by reducing from the Label Cover problem. We remark that our result is based on the $P \neq NP$ assumption in contrast to Barman et al.’s result [2020] which assumes the Unique Games Conjecture. As a consequence of our result, \mathbf{w} -THIELE rule is APX-hard whenever marginal utility goes to 0 and, assuming Gap-ETH, no FPT-approximation scheme exists (wrt. parameter k). Additionally, we provide a family of hard instances for the standard greedy algorithm on which our LP-based algorithm works better (Theorem 18).

Applications. Multiwinner elections has practical implications not only in committee [Skowron et al., 2016] or parliamentary elections (apportionment seats in parliament) [Brill et al., 2018]. Thiele rules with a knapsack constraint (instead of cardinality constraint k) can be considered as participatory budgeting model [Fluschnik et al., 2019; Talmon and Faliszewski, 2019]. Multiwinner elections can model many real-world issues in which we need to choose some number of elements. In such scenarios, approximation algorithms seem to have clear motivation. For more discussion on using approximation and randomized algorithms in collective decision making, we refer the reader to, e.g., [Skowron et al., 2016; Byrka et al., 2018; Cygan et al., 2018].

Organisation of the paper. In Section 2 we define additional notations that will be used throughout. In Section 3 we describe the algorithm and provide the proof of its approximation guarantee. In Section 4 we prove our hardness of approximation result. All the proofs omitted due to space restrictions will be available in the full version of this paper.

2 Preliminaries

An approval-based election $\mathcal{E} = (V, C, \{A_v\}_{v \in V})$ consists of a set V of n voters, a set C of m candidates, and, for each voter $v \in V$, a set $A_v \subseteq C$ of candidates approved by v . The \mathbf{w} -THIELE rule is parameterized by an infinite-dimensional non-increasing non-negative sequence $\mathbf{w} = (w_1, w_2, \dots)$; the \mathbf{w} -THIELE score (total utility) of a subset $W \subseteq C$ is defined as $\text{scr}_{\mathbf{w}}^{\mathcal{E}}(W) := \sum_{v \in V} \sum_{i=1}^{|A_v \cap W|} w_i$. The goal of an optimization problem called \mathbf{w} -THIELE rule is to find a set $W \subseteq C$ of k candidates that maximizes its \mathbf{w} -THIELE score.

For notational convenience, we assume throughout that $w_1 = 1$.³ Furthermore, for brevity, the sequence \mathbf{w} is henceforth assumed to be non-negative and non-increasing.

We say that a sequence \mathbf{w} (and its corresponding \mathbf{w} -THIELE rule) is *geometrically dominant* if for every $i \in \mathbb{N}$, $w_i \cdot w_{i+2} \geq w_{i+1}^2$. It follows that if \mathbf{w} is geometrically dominant then either $\mathbf{w} = (1, 0, 0, \dots)$ or $\frac{w_i}{w_{i+1}} \geq \frac{w_{i+1}}{w_{i+2}}$ for all $i \in \mathbb{N}$. All Thiele rules in Table 1 satisfy these properties.

We use $\text{Poi}(\cdot)$, $\text{Ber}(\cdot)$, $\text{Bin}(\cdot, \cdot)$ to denote the Poisson, Bernoulli and Binomial random variables respectively.

For a sequence \mathbf{w} , we let $\alpha_{\mathbf{w}}$ denote

$$\alpha_{\mathbf{w}} := \mathbb{E}_{x \sim \text{Poi}(1)} \left[\sum_{i=1}^x w_i \right] = \sum_{x=1}^{\infty} \frac{1}{e \cdot x!} \cdot \left(\sum_{i=1}^x w_i \right). \quad (1)$$

In words, $\alpha_{\mathbf{w}}$ is the expected utility of a voter if the number of elected candidates that he/she approves, is distributed as Poisson distribution with mean one.

3 Approximation Algorithm

In this section, we describe our algorithm and argue that it gives $\alpha_{\mathbf{w}}$ -approximate solution, thereby proving the algorithmic part of Theorem 1. Our algorithm is based on an LP rounding technique, which requires us to first define $\text{scr}_{\mathbf{w}}^{\mathcal{E}}$ also on a fractional solution $x \in [0, 1]^C$ specifying for each candidate fractionally, how much the candidate is selected.

$$\text{scr}_{\mathbf{w}}^{\mathcal{E}}(x) = \sum_{v \in V} \sum_{l=1}^{|A_v|} w_l \cdot \min \{1, \max \{0, x_{A_v} - l + 1\}\},$$

where $x_{A_v} = \sum_{c \in A_v} x_c$. Intuitively, for each voter v this function adds the first x_{A_v} elements of the sequence \mathbf{w} . The last element of the sequence may be counted in only fractionally. With this definition, we now compute the optimum fractional solution x^* of the following linear program.

$$\begin{aligned} & \text{maximize} && \text{scr}_{\mathbf{w}}^{\mathcal{E}}(x) \\ & \text{subject to} && \sum_{c \in C} x_c = k \\ & && x \in [0, 1]^C \end{aligned}$$

To round x^* , we will employ a framework called *pipage rounding* [Ageev and Sviridenko, 2004]. Specifically, we resort to the following result of Călinescu et al.⁴

³This is w.l.o.g. since, as long as $w_1 \neq 0$, we may scale sequence \mathbf{w} by any constant factor without effecting the approximation ratio.

⁴Applying the theorem directly results in a randomized algorithm. A way to construct a deterministic algorithm is described in the full version of this paper.

Problem Name	Sequence \mathbf{w}	References	Our Ratio
APPROVAL CHAMBERLIN-COURANT	$\{1, 0, 0, \dots\}$	[Chamberlin and Courant, 1983]	$0.6321 \dots = 1 - 1/e$
PROPORTIONAL APPROVAL VOTING	$w_i = 1/i$	[Thiele, 1895; Byrka <i>et al.</i> , 2018]	$0.7965 \dots$
Saint-Laguë method	$w_i = 1/(2i - 1)$	[Brill <i>et al.</i> , 2018]	$0.7394 \dots$
Penrose apportionment method	$w_i = 1/i^2$	[Brill <i>et al.</i> , 2018]	$0.7084 \dots$
p -HARMONIC	$w_i = 1/i^p$	[Faliszewski <i>et al.</i> , 2017]	$\sum_{x=1}^{\infty} \frac{1}{e \cdot x!} \cdot \left(\sum_{i=1}^x \frac{1}{i^p} \right)$
p -GEOMETRIC	$w_i = p^i$	[Skowron <i>et al.</i> , 2016]	$\frac{1}{1-p} \cdot \left(1 - \frac{1}{e^{1-p}} \right)$

Table 1: Our approximation ratios for \mathbf{w} -THIELE rules. They are tight unless $\mathbf{P} = \mathbf{NP}$. All listed \mathbf{w} -THIELE rules were known to have $(1 - 1/e)$ -approximation algorithm [Nemhauser *et al.*, 1978; Skowron *et al.*, 2016]. To the best of our knowledge, no better than $(1 - 1/e)$ -approximation algorithm was known for any of these rules.

Theorem 2 (Călinescu *et al.* [2011, Lemma 3.5]). *Let f be a monotone, submodular function and a polytope $B(\mathcal{M})$ be described by matroid constraints. Given vector $y^* = \arg \max \{f(x) \mid x \in B(\mathcal{M})\}$ the procedure **PipageRounding**(\mathcal{M}, y^*) will return a solution $S \in \mathcal{M}$ of value $\mathbb{E}[f(S)] \geq \mathbb{E}_{\hat{y} \sim y^*}[f(\hat{y})]$ where $\hat{y} \sim y^*$ means a random vector with each coordinate \hat{y}_i independently chosen to be 1 with probability y_i^* or 0 otherwise.*

The function $\text{scr}_{\mathbf{w}}^{\mathcal{E}}$ is submodular (because \mathbf{w} is non-increasing) and monotone (because \mathbf{w} is non-negative). We can thus use Theorem 2 to round the optimum LP solution x^* . The resulting approximation ratio is then equal to

$$\rho = \frac{\mathbb{E}_{\hat{x} \sim x^*}[\text{scr}_{\mathbf{w}}^{\mathcal{E}}(\hat{x})]}{\text{scr}_{\mathbf{w}}^{\mathcal{E}}(x^*)}.$$

The remainder of this section is dedicated to bounding ρ .

Let the approximation ratio of a voter v be $\rho_v := \frac{\mathbb{E}_{\hat{x} \sim x^*}[\text{scr}_{\mathbf{w}}^{\mathcal{E}}(\hat{x}, v)]}{\text{scr}_{\mathbf{w}}^{\mathcal{E}}(x^*, v)}$ where $\text{scr}_{\mathbf{w}}^{\mathcal{E}}(x, v)$ is the score of v , i.e., $\sum_{l=1}^{|A_v|} w_l \cdot \min\{1, \max\{0, x_{A_v} - l + 1\}\}$. Henceforth, we will focus on a single voter v and show that $\rho_v \geq \alpha_{\mathbf{w}}$. This straightforwardly implies that $\rho \geq \alpha_{\mathbf{w}}$ as desired.

We next characterize x^* that minimizes the ratio ρ_v , which in the end will lead us to the claimed bound.

Step 1: Three-Valued x^*

Let $\tau = x_{A_v}^* = \sum_{c \in A_v} x_c^*$ be a sum of x^* values for candidates approved by v . The score of the fractional solution x^* is now $\text{scr}_{\mathbf{w}}^{\mathcal{E}}(x^*, v) = w_1 + \dots + w_{\lfloor \tau \rfloor} + \{\tau\} \cdot w_{\lfloor \tau \rfloor + 1}$.⁵

The numerator $\mathbb{E}_{\hat{x} \sim x^*}[\text{scr}_{\mathbf{w}}^{\mathcal{E}}(\hat{x}, v)]$ of ρ_v can be written as

$$\sum_{l=0}^{\infty} \left((w_1 + \dots + w_l) \mathbb{P}_{\hat{x} \sim x^*} \left[\sum_{c \in A_v} \hat{x}_c = l \right] \right).$$

We recall the following lemma due to Barman *et al.*

Lemma 3 (Barman *et al.* [2020]⁶). *Let $x \in [0, 1]^m$ be such that $x_1 + \dots + x_m = \tau$. Then, for any non-negative sequence (a_1, a_2, \dots) , there exist $q \in (0, 1)$ and $x' \in \{0, q, 1\}^m$ such that $\sum_i x'_i = \tau$ and $\sum_l (a_l \mathbb{P}_{\hat{x} \sim x'}[\sum_{i=1}^m \hat{x}_i = l]) \geq \sum_l (a_l \mathbb{P}_{\hat{x} \sim x}[\sum_{i=1}^m \hat{x}_i = l])$.*

⁵ $\{x\} := x - \lfloor x \rfloor$ is a fractional part of a number x .

⁶In [Barman *et al.*, 2020, Lemma 2.5], τ is required to be an integer. It is however obvious from the proof that this is not necessary.

Now, notice that the denominator $\text{scr}_{\mathbf{w}}^{\mathcal{E}}(x^*, v)$ in our ratio ρ_v only depends on the sum $\tau = \sum_{c \in A_v} x_c^*$ and not the way values are distributed among x_c^* s. Hence, from Lemma 3, ρ_v is minimised by x^* that only takes three values—0, q and 1, although we do not yet know what the value of q is.

Step 2: Getting Rid of Ones

Let $w(x)$ be a function $w : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ extending \mathbf{w} , defined by $w(x) := w_{\lfloor x \rfloor} \cdot (1 - \{x\}) + w_{\lfloor x \rfloor + 1} \cdot \{x\}$. Clearly, when x is integral $w(x) = w_x$. For fractional x , the function just takes the weighted average of $w_{\lfloor x \rfloor}, w_{\lfloor x \rfloor + 1}$. The following observation is straightforward (as w is geometrically dominant).

Observation 4. *The function w is convex.*

Using this observation, we can prove the following lemma, which allows us to only consider x^* that does not contain 1.

Lemma 5. *Suppose $x_c^* = 1$ for a candidate $\tilde{c} \in A_v$. Then the approximation ratio for the voter v decreases after removing \tilde{c} from the fractional solution x^* . Formally, if $x_{\tilde{c}}^* = 1$, then*

$$\begin{aligned} \rho_v &= \frac{\sum_{l=0}^{\infty} (w_l \mathbb{P}_{\hat{x} \sim x^*}[\sum_c \hat{x}_c \geq l])}{w_1 + \dots + w_{\lfloor \tau \rfloor} + \{\tau\} \cdot w_{\lfloor \tau \rfloor + 1}} \\ &\geq \frac{\sum_{l=0}^{\infty} (w_l \mathbb{P}_{\hat{x} \sim x^*}[\sum_{c \neq \tilde{c}} \hat{x}_c \geq l])}{w_1 + \dots + w_{\lfloor \tau - 1 \rfloor} + \{\tau\} \cdot w_{\lfloor \tau \rfloor}}. \end{aligned}$$

Hence, the fractional solution x^* that minimizes the ratio ρ_v of a single voter only has values 0 and q where $q \in (0, 1)$.

Step 3: Comparing Binomial and Poisson

The values $\{x_c^*\}_{c \in A_v}$ are now all either 0 or q and sum up to τ . It means that, when we sample $\hat{x} \sim x^*$, we now are performing τ/q coin tosses with an identical coin with bias q . The value of our integral solution which is equal to $\sum_l (w_l \cdot \mathbb{P}_{\hat{x} \sim x^*}[\sum_c \hat{x}_c \geq l])$ can be now rewritten as $\sum_l (w_l \cdot \mathbb{P}[\text{Bin}(\tau/q, q) \geq l]) = \mathbb{E}[w_1 + \dots + w_{\text{Bin}(\tau/q, q)}]$.

Let $s_w(n) = w_1 + \dots + w_n$. The function s_w is monotone (as \mathbf{w} is non-negative) and concave (as \mathbf{w} is non-increasing). We recall another lemma by Barman *et al.*

Lemma 6 (Barman *et al.* [2020, Lemma 2.3]). *For any convex function f , any integer $N \geq 1$ and parameter $p \in [0, 1]$ we have $\mathbb{E}[f(\text{Bin}(N, p))] \leq \mathbb{E}[f(\text{Poi}(Np))]$.*

By applying the above lemma with the function $s_{-w}(x) = -s_w(x)$ which is convex, we get the following inequality

$\mathbb{E}[w_1 + \dots + w_{\text{Bin}(\tau/q, q)}] \geq \mathbb{E}[w_1 + \dots + w_{\text{Poi}(\tau)}]$ and hence

$$\rho_v \geq \frac{\mathbb{E}[w_1 + \dots + w_{\text{Poi}(\tau)}]}{w_1 + \dots + w_{\lfloor \tau \rfloor} + \{\tau\} \cdot w_{\lfloor \tau \rfloor + 1}}. \quad (2)$$

Step 4: Changing the Mean of Poisson

A property of Poisson distributions is that sampling $X_i \sim \text{Poi}(\lambda_i)$ for $i = 1, \dots, m$ and adding them is equivalent to sampling $Y \sim \text{Poi}(\sum_i \lambda_i)$. Hence, $\mathbb{E}[w_1 + \dots + w_{\text{Poi}(\tau)}]$

$$= \left(\sum_{l=1}^{\lfloor \tau \rfloor} \mathbb{E}_{y \sim \text{Poi}(l-1)} [w_{y+1} + \dots + w_{y+\text{Poi}(1)}] \right) + \mathbb{E}_{y \sim \text{Poi}(\lfloor \tau \rfloor)} [w_{y+1} + \dots + w_{y+\text{Poi}(\{\tau\})}]. \quad (3)$$

We can bound each term in (3) as follows.

Lemma 7. For any $l \in \{1, \dots, \lfloor \tau \rfloor\}$, we have

$$\mathbb{E}_{y \sim \text{Poi}(l-1)} [w_{y+1} + \dots + w_{y+\text{Poi}(1)}] \geq \mathbb{E}[w_1 + \dots + w_{\text{Poi}(1)}] \frac{w_l}{w_1}.$$

Lemma 8. $\mathbb{E}_{y \sim \text{Poi}(\lfloor \tau \rfloor)} [w_{y+1} + \dots + w_{y+\text{Poi}(\{\tau\})}] \geq \mathbb{E}[w_1 + \dots + w_{\text{Poi}(1)}] \cdot \frac{w_{\lfloor \tau \rfloor + 1}}{w_1} \cdot \{\tau\}.$

Putting together (2), (3), Lemmas 7 and 8 allows us to conclude that the ratio ρ_v is at least

$$\mathbb{E}[w_1 + \dots + w_{\text{Poi}(1)}] \cdot \frac{w_1 + \dots + w_{\lfloor \tau \rfloor} + \{\tau\} \cdot w_{\lfloor \tau \rfloor + 1}}{w_1 (w_1 + \dots + w_{\lfloor \tau \rfloor} + \{\tau\} \cdot w_{\lfloor \tau \rfloor + 1})}$$

which is equal to $\mathbb{E}[w_1 + \dots + w_{\text{Poi}(1)}] = \alpha_{\mathbf{w}}$ since $w_1 = 1$. Hence, our algorithm yields an $\alpha_{\mathbf{w}}$ -approximate solution.

4 Hardness of Approximation

We next prove the hardness of approximation (second part of Theorem 1). Our proof largely follows the proof of $(1 - 1/e + o(1))$ -inapproximability of Max k -Coverage of Feige [1998]. Specifically, we reduce from the Label Cover problem:

Definition 9. A Label Cover instance $\mathcal{L} = (A, B, E, [L], [R], \{\pi_e\}_{e \in E})$ consists of a bi-regular bipartite graph (A, B, E) , alphabet sets $[L], [R]$ and, for every edge $e \in E$, a constraint $\pi_e : [L] \rightarrow [R]$.

A labelling of \mathcal{L} is a function $\phi : A \rightarrow [L]$. We say that ϕ satisfies a right vertex $v \in B$ if, for every two neighbours u, u' of v , we have $\pi_{(u,v)}(\phi(u)) = \pi_{(u',v)}(\phi(u'))$. The value of ϕ , denoted by $\text{val}(\phi)$, is defined as the fraction of right vertices satisfied by ϕ . Moreover, we say that ϕ weakly satisfies a right vertex $v \in B$ if there exist two distinct neighbours u, u' of v such that $\pi_{(u,v)}(\phi(u)) = \pi_{(u',v)}(\phi(u'))$. Similarly, we define the weak value of ϕ , denoted by $\text{weak-val}(\phi)$, as the fraction of right vertices weakly satisfied by ϕ . We use $\text{val}(\mathcal{L})$ (resp. $\text{weak-val}(\mathcal{L})$) to denote the maximum value (resp. weak value) among all assignments ϕ , i.e., $\text{val}(\mathcal{L}) := \max_{\phi: A \rightarrow [L]} \text{val}(\phi)$ (resp. $\text{weak-val}(\mathcal{L}) := \max_{\phi: A \rightarrow [L]} \text{weak-val}(\phi)$).

For any $\delta > 0$ and $t \in \mathbb{N}$, let δ -Gap-Label-Cover(t, R) denote the following problem: given a Label Cover instance \mathcal{L} with right degree t and right alphabet set $[R]$, distinguish between $\text{val}(\mathcal{L}) = 1$ and $\text{weak-val}(\mathcal{L}) < \delta$.

The hardness of δ -Gap-Label-Cover(t, R) is well-known and is often used as a starting point for reductions in hardness of approximation. Below we summarize the two hardness of approximation of Label Cover relevant to our work.

Theorem 10 (Feige [1998]). For every $\delta > 0$, every $t \in \mathbb{N}$, $t \geq 2$, and for any sufficiently large $R \in \mathbb{N}$ (depending on δ, t), δ -Gap-Label-Cover(t, R) is NP-hard.

Theorem 11 (Manurangsi [2020]). Assuming Gap-ETH, for every $\delta > 0$, every $t \in \mathbb{N}$, $t \geq 2$, and any sufficiently large $R \in \mathbb{N}$ (depending on δ, t), no $f(k) \cdot N^{o(k)}$ -time algorithm can solve δ -Gap-Label-Cover(t, R), where k denote the number of left vertices in Label Cover and f can be any function.

We may now state the main properties of our reductions:

Theorem 12. Let \mathbf{w} be such that $\lim_{i \rightarrow \infty} w_i = 0$. For any $\epsilon > 0$, there exist $\delta > 0, t \in \mathbb{N}$ and a reduction that takes in an instance $\mathcal{L} = (A, B, E, [L], [R], \{\pi_e\}_{e \in E})$ of δ -Gap-Label-Cover(t, R) and produces an election $E = (V, C, \{A_v\}_{v \in V})$ in $\text{poly}(|C|, |V|)$ time such that:

- (Completeness) If $\text{val}(\mathcal{L}) = 1$, then there exists $W \subseteq C$ such that $\text{scr}_{\mathbf{w}}^E(W) = |V|$.
- (Soundness) If $\text{weak-val}(\mathcal{L}) < \delta$, then for all $W \subseteq C$, we have $\text{scr}_{\mathbf{w}}^E(W) = (\alpha_{\mathbf{w}} + \epsilon) \cdot |V|$.
- (Size bound) $|C| = |A| \cdot L$ and $|V| = |B| \cdot t^R$.
- (Parameter) $k = |A|$.

Plugging in Theorem 12 to Theorems 10, 11 immediately yields the desired hardness of approximation in Theorem 1.

The rest of this section is devoted to the proof of Theorem 12. Our reduction is in fact the same as Feige's [1998], which can be interpreted as the inapproximability for APPROVAL CHAMBERLIN-COURANT. Hence, on the hardness front, our main contribution is in extending the analysis to work with more general sequences of weights. To this end, we provide several helpful auxiliary lemmas in Section 4.1. Then, in Section 4.2, we describe Feige's reduction together with our generalized analysis, and prove Theorem 12.

4.1 Auxiliary Lemmas

Similar to the analysis of our algorithm, we will need to move from (generalizations of) Binomial random variables to Poisson random variables. The difference between here and the algorithmic counterpart (Section 3) is that the sign is flipped. Specifically, in Section 3, we need $\mathbb{E}[w_1 + \dots + w_{\text{Bin}(n, p)}] \geq \mathbb{E}[w_1 + \dots + w_{\text{Poi}(np)}]$. Below, we show that in fact the left-hand side is not much more than the right-hand side. Our lemma is presented in its full generality below.

Lemma 13. Let \mathbf{w} be such that $\lim_{i \rightarrow \infty} w_i = 0$, $\gamma, \lambda > 0$ be any positive real numbers. Then, there exists $p_0 = p_0(\mathbf{w}, \gamma, \lambda) > 0$ such that, for any positive real number $p < p_0$, any positive integer $n < \lambda/p$ and any non-negative integers m_1, \dots, m_ℓ with $\sum_{i=1}^{\ell} m_i = n$, we have

$$\mathbb{E}_{z_1, \dots, z_\ell \sim \text{Ber}(p)} [w_1 + \dots + w_{\sum_{i=1}^{\ell} m_i z_i}] < \mathbb{E}[w_1 + \dots + w_{\text{Poi}(np)}] + \gamma.$$

Another useful fact is stated below.

Lemma 14. The function $f : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ given by $f(\lambda) = \mathbb{E}_{z \sim \text{Poi}(\lambda)} [w_1 + \dots + w_z]$ is concave.

4.2 The Reduction and Proof of Theorem 12

We can now describe the reduction and prove Theorem 12. As stated earlier, our reduction is the same as that in [Feige, 1998], which employs the following gadget:

Definition 15 (Hypercube Set System [Feige, 1998]). For any integer $t \in \mathbb{N}, R \in \mathbb{N}$, the (t, R) -hypercube set system is defined as $\mathcal{H} = (\mathcal{U}, \mathcal{P}_1, \dots, \mathcal{P}_R)$ where the universe \mathcal{U} is $[t]^R$ and each $\mathcal{P}_j = (P_{j,1}, \dots, P_{j,t})$ is a partition of \mathcal{U} into t parts where $P_{j,i} = \{\mathbf{u} = (u_1, \dots, u_R) \in [t]^R \mid u_j = i\}$.

In our reduction, we first pick the parameters as follows.

- Let i^* denote the smallest positive integer such that $w_{i^*} < 0.1\epsilon$, and let $\vartheta = 10i^*/\epsilon$.
- Let $\gamma = 0.7\epsilon$, $p_0 = p_0(\mathbf{w}, \gamma, 10\vartheta)$ be as in Lemma 13.
- We select $t = \lceil 2/p_0 \rceil$ and $\delta = \frac{0.0001\epsilon}{\vartheta^3 t^2}$.

Given a Label Cover instance $\mathcal{L} = (A, B, E, [L], [R], \{\pi_e\}_{e \in E})$ where the right degree is t , we create an election $\mathcal{E} = (V, C, \{A_v\}_{v \in V})$ as follows.

- Let $(\mathcal{U}, \mathcal{P}_1, \dots, \mathcal{P}_R)$ be a (t, R) -hypercube set system.
- Let the set of voters be $V := B \times \mathcal{U}$.
- Let the set of candidates be $C := A \times [L]$.
- For each $b \in B$, order its neighbours arbitrarily and, for a neighbour a of b , let $\iota(a) \in [t]$ denote its index in this ordering. A voter $(b, \mathbf{u}) \in B \times \mathcal{U}$ approves a candidate $(a, \sigma) \in A \times [L]$ iff a is a neighbour of b and $\mathbf{u} \in P_{\pi_{(a,b)}(\sigma), \iota(a)}$.
- Let $k = |A|$.

The reduction obviously runs in $\text{poly}(|C|, |V|)$ time and the sizes of C, V are as stated in Theorem 12.

Completeness. Suppose that there is an assignment $\phi : A \rightarrow [L]$ that satisfies every $b \in B$. Pick $W = \{(a, \phi(a)) \mid a \in A\}$. We claim that $\text{scr}_w^{\mathcal{E}}(W) \geq |V|$; specifically, we will argue that every voter approves a candidate in W . Consider any voter $(b, \mathbf{u}) \in B \times \mathcal{U}$ and let a_1, \dots, a_t be b 's (ordered) neighbours. Since ϕ satisfies b , we have $\pi_{(a_1,b)}(\phi(a_1)) = \dots = \pi_{(a_t,b)}(\phi(a_t))$. Since $\mathcal{P}_{\pi_{a_1,b}(\phi(a_1))}$ is a partition of \mathbf{u} , there exists $i \in [t]$ such that $\mathbf{u} \in P_{\pi_{a_1,b}(\phi(a_1)), i} = P_{\pi_{a_i,b}(\phi(a_i)), i}$. This implies that (b, \mathbf{u}) approves $(a_i, \phi(a_i))$ as desired.

Soundness. Assume contrapositively that there exists a subset $W \subseteq C$ of k candidates, such that $\text{scr}_w^{\mathcal{E}}(W) \geq (\alpha_w + \epsilon) \cdot |V|$. We will show that $\text{weak-val}(\phi) \geq \delta$.

For every $b \in B$, let $V_b = \{b\} \times \mathcal{U}$, $\mathcal{E}_b = (V_b, C, \{A_v\}_{v \in V_b})$ denote the subelection on the voters V_b , and let $\deg_W(b) := |W \cap (N(b) \times [L])|$ where $N(b)$ denote the set of neighbours of b . For every $a \in A$, let $W_a \subseteq [L]$ denote $\{\sigma \in [L] \mid (a, \sigma) \in W\}$. Finally, let $D = 10\vartheta t$.

We divide the set B into three sets:

- Let B_1 denote the set of $b \in B$ s.t. $\deg_W(b) > D$.
- Let B_2 denote the set of $b \in B$ such that $\deg_W(b) \leq D$ and, for all distinct neighbours a_1, a_2 of b and all $\sigma_1 \in W_{a_1}, \sigma_2 \in W_{a_2}$, we have $\pi_{(a_1,b)}(\sigma_1) \neq \pi_{(a_2,b)}(\sigma_2)$.
- Let $B_3 = B \setminus (B_1 \cup B_2)$.

We will next argue that $|B_3| \geq \Omega_{t,\epsilon}(|B|)$, which will then allow us to “decode” back the assignment to the Label Cover instance \mathcal{L} . To do so, we start by bounding the contribution of B_1 to $\text{scr}_w^{\mathcal{E}}(W)$. The main idea is that when $\deg_W(b)$ is large,

the “average score per candidate” becomes small, which gives the following upper bound on the desired quantity:

Proposition 16. $\sum_{b \in B_1} \text{scr}_w^{\mathcal{E}_b}(W) \leq 0.2\epsilon \cdot |V|$.

Next, we bound the score contribution from B_2 , using the auxiliary lemmas shown in the previous subsection.

Proposition 17. $\sum_{b \in B_2} \text{scr}_w^{\mathcal{E}_b}(W) \leq (\alpha_w + 0.7\epsilon) \cdot |V|$.

Proof. Let us fix $b \in B_2$. Let a_1, \dots, a_t be its (ordered) neighbours. For each $\kappa \in [R]$, let $\tau^\kappa = |\{(a, \sigma) \in W \mid a \in N(b), \pi_{(a,b)}(\sigma) = \kappa\}|$. Moreover, if $\tau^\kappa \neq 0$, let i^κ denote the index $i \in [t]$ of the neighbour s.t. $\pi_{(a_i,b)}(\sigma_i) = \kappa$ for some $\sigma_i \in W_{a_i}$. (Due to our definition of B_2 , there is a unique such i .) If $\tau^\kappa = 0$, let $i^\kappa = \perp$.

Recall that a voter (b, \mathbf{u}) approves a candidate (a_i, σ) iff $\mathbf{u} \in P_{\pi_{(a_i,b)}(\sigma), i}$, or equivalently $u_{\pi_{(a_i,b)}(\sigma)} = i$. Hence, the number of candidates in W approved by $(b, \mathbf{u}) \in V_b$ is

$$\sum_{i \in [t]} \sum_{\sigma \in W_{a_i}} \mathbf{1}[u_{\pi_{(a_i,b)}(\sigma)} = i] = \sum_{\kappa \in [R]} \tau^\kappa \cdot \mathbf{1}[u_{\kappa} = i^\kappa],$$

where $\mathbf{1}$ is an indicator function. Hence, we may rearrange $\text{scr}_w^{\mathcal{E}_b}(W)$ as

$$|\mathcal{U}| \cdot \mathbb{E}_{u_1, \dots, u_R \sim [t]} \left[w_1 + \dots + w_{\sum_{\kappa \in [R]} \tau^\kappa \cdot \mathbf{1}[u_\kappa = i^\kappa]} \right].$$

Observe that $\mathbf{1}[u_1 = i^1], \dots, \mathbf{1}[u_R = i^R]$ are i.i.d. $\text{Ber}(1/t)$ random variables. Moreover, $\sum_{\kappa \in [R]} \tau^\kappa = \deg_W(b) \leq D = 10\vartheta t$. By our choice of t , we may apply Lemma 13 to get

$$\text{scr}_w^{\mathcal{E}_b}(W) \leq |\mathcal{U}| \cdot \mathbb{E} \left[w_1 + \dots + w_{\text{Poi}(\deg_W(b)/t)} \right] + \gamma.$$

By summing this over $b \in B_2$, $\sum_{b \in B_2} \text{scr}_w^{\mathcal{E}_b}(W)$ is at most

$$\begin{aligned} & |\mathcal{U}| \cdot \sum_{b \in B_2} \left(\mathbb{E} \left[w_1 + \dots + w_{\text{Poi}(\deg_W(b)/t)} \right] + \gamma \right) \\ & \leq |\mathcal{U}| \cdot |B| \cdot \left(\mathbb{E} \left[w_1 + \dots + w_{\text{Poi}\left(\frac{\sum_{b \in B} \deg_W(b)}{|B|t}\right)} \right] + \gamma \right) \\ & = (\alpha_w + 0.7\epsilon) \cdot |V|, \end{aligned}$$

where the inequality comes from Lemma 14, and the equality from $\sum_{b \in B} \deg_W(b) = |B| \cdot t$, $|\mathcal{U}| \cdot |B| = |V|$ and from our choice of γ . \square

From the two propositions above and our assumption that $\text{scr}_w^{\mathcal{E}}(W) \geq (\alpha_w + \epsilon) \cdot |V|$, we must have $\sum_{b \in B_3} \text{scr}_w^{\mathcal{E}_b}(W) \geq 0.1\epsilon \cdot |V|$. It is obvious to see that $\text{scr}_w^{\mathcal{E}_b}(W) \leq \deg_W(b) \cdot |V_b|/t$, which is at most $D \cdot |V|/(t \cdot |B|)$ for $b \in B_3$. Thus,

$$|B_3| \geq \frac{0.1\epsilon \cdot |V|}{D \cdot |V|/(t \cdot |B|)} = \frac{0.1\epsilon t |B|}{D} = \frac{0.01\epsilon |B|}{\vartheta}. \quad (4)$$

Now, consider the following (random) assignment: for every $a \in A$ such that $W_a \neq \emptyset$, let $\phi(a)$ be a random element of W_a . By definition, each $b \in B_3$ has distinct neighbours a_1, a_2 and $\sigma_1 \in W_{a_1}, \sigma_2 \in W_{a_2}$ such that $\pi_{(a_1,b)}(\sigma_1) = \pi_{(a_2,b)}(\sigma_2)$. Hence, the probability that ϕ weakly satisfies such b is at least $\frac{1}{|W_{a_1}| \cdot |W_{a_2}|} \geq \frac{1}{D^2}$, where the inequality follows from $\deg_W(b) \leq D$. Thus,

$$\mathbb{E}_\phi[\text{weak-val}(\phi)] \geq \frac{|B_3|}{|B| \cdot D^2} \stackrel{(4)}{\geq} \frac{0.01\epsilon}{\vartheta \cdot D^2} = \frac{0.0001\epsilon}{\vartheta^3 t^2} = \delta.$$

This completes our proof of Theorem 12.

5 Conclusion

In this work, we provide an approximation algorithm and hardness of approximation for w -THIELE rules, which are tight for the natural class of *geometrically dominant* rules. Many popular Thiele rules, such as PROPORTIONAL APPROVAL VOTING and p -GEOMETRIC, belong to this class. While our results resolve the approximability for these Thiele rules, many interesting questions still remain open.

Tight (in)approximability for other Thiele rules. An immediate open question from our work is whether there is a simple tight characterization of the approximability of w -THIELE rules for any (not necessarily geometrically dominant) w . Regarding this, recall that Barman et al. [2020] gives such a tight bound for $w = (1, \dots, 1, 0, \dots)$. Their ratio is *not* of the form (1). Furthermore, we note that the hardness results of Barman et al. [2020] need to rely on the Unique Games Conjecture (UGC), not just $P \neq NP$ as in our work. These highlight the challenges in generalizing our results to larger families of w .

Performance of the greedy algorithm. As touched upon earlier, a well-studied algorithm for w -THIELE (and its generalization) is the greedy algorithm, in which at each step the candidate whose inclusion maximizes the total utility increase is added to the committee. This is repeated until k candidates are selected. The greedy algorithm is very efficient and hence maybe preferred over LP-based algorithms in practice. Theoretically, the greedy algorithm is known to provide $(1 - 1/e)$ -approximate solution for w -THIELE for any w [Nemhauser et al., 1978; Skowron et al., 2016]. An interesting direction is to determine the exact approximation guarantees of the greedy algorithm for specific w , such as PROPORTIONAL APPROVAL VOTING and p -GEOMETRIC. On this front, we manage to show the following:

Theorem 18. *For any Thiele sequence $w = (1, p, w_3, \dots)$, the approximation ratio of the greedy algorithm for w -THIELE is at most $\frac{1-p}{1-p} \cdot (1 - \frac{1}{e^{1-p}})$.*

An important consequence of the above theorem is that, for any geometrically dominant $w = (1, p, w_3, \dots)$ such that w is *not* p -GEOMETRIC, the approximation ratio guarantee by our LP-based algorithm is *strictly better* than that of the greedy algorithm. It remains open whether this also holds for p -GEOMETRIC and other non-geometrically dominant w .

Extensions to OWA-Winner and Participatory Budgeting. Skowron et al. [2016] studied the w -OWA-Winner problem, which is a generalization of w -THIELE rule that allows each voter to assign any utility to a candidate (in approval ballots the utility was 0 or 1). They, using a result of Nemhauser et al. [1978], gave an $(1 - 1/e)$ -approximation algorithm for all non-increasing OWA vectors w . We note that the greedy algorithm of Khuller et al. [1999] also gives an $(1 - 1/e)$ -approximation for the OWA-Winner problem with a knapsack constraint (instead of a cardinality constraint) due to analysis of Sviridenko [2004]. Knapsack constraint models an issue of Participatory Budgeting (PB) where candidates are projects with assigned costs. Only specific sequences of weights were considered, e.g., $(1, 0, \dots)$, $(1, 1, \dots)$ and $(1, 1/2, 1/3, \dots)$ [Talmon and Faliszewski, 2019; Fluschnik et

al., 2019]. It is interesting whether we can get better than $(1 - 1/e)$ -approximation algorithm for OWA-Winner using our results or if they can be extended by replacing cardinality constraint with a knapsack constraint (but keeping approval ballots as this is the most popular way of balloting—indeed, general non-increasing sequences of weights with approval ballots model projects substitution [Jain et al., 2020]). We note here that our hardness results in Theorems 1 and 18 also hold for OWA-Winner and PB.

Minimization variants. One can consider the minimization version of the problem, i.e., minimizing dissatisfaction instead of maximizing utility. Such a minimization variant, coming from computational social choice, has already been considered for several Thiele rules. Most of them assume that voters and candidates lie in a metric space, and dissatisfaction of a voter from a candidate is equal to the distance between them (not necessarily 0 or 1). The most popular objective function is minimizing the sum of distances of the voters to the closest chosen candidate, i.e., $w = (1, 0, \dots)$. This problem is called k -Median and can be approximated within 2.675-ratio [Byrka et al., 2017]. Hardness of approximation for k -Median is known to be $1 + 2/e$ [Guha and Khuller, 1999], even in FPT wrt. k time, in which the gap has been closed [Cohen-Addad et al., 2019]. The minimization version of w -THIELE rules is called OWA k -Median for which 93-approximation was shown [Byrka et al., 2018]. For some specific w , better approximation ratios have been achieved. For instance, when $w_\ell(i) = (1, \dots, 1, 0, \dots)$ (i.e., ℓ ones followed by zeros), the problem is called ℓ -Fault Tolerant k -Median and a 4-approximation algorithm is known [Swamy and Shmoys, 2008]. When $w(i) = 1/i$, the problem is called Harmonic k -Median and a 2.36-approximation algorithm exists [Byrka et al., 2018]. Surprisingly, this last result does not assume that the function of dissatisfaction is a metric. There are still many approximation gaps for the minimization variants of w -THIELE rules that should be studied.

Acknowledgements

This research is partially supported by the National Science Centre, Poland (grant numbers 2015/18/E/ST6/00456, 2018/28/T/ST6/00366, 2018/29/B/ST6/02633) and the Israel Science Foundation (grant number 630/19). Part of this work was done while Pasin was visiting University of Wrocław. Krzysztof Sornat is partially supported by the Foundation for Polish Science (FNP) within the START programme.

We would like to thank Mateusz Lewandowski for initial discussion and the anonymous reviewers for their helpful comments.

References

- [Ageev and Sviridenko, 2004] Alexander A. Ageev and Maxim Sviridenko. Pipage Rounding: A New Method of Constructing Algorithms with Proven Performance Guarantee. *J. Comb. Optim.*, 8(3):307–328, 2004.
- [Aziz et al., 2015] Haris Aziz, Serge Gaspers, Joachim Gudmundsson, Simon Mackenzie, Nicholas Mattei, and Toby Walsh. Computational Aspects of Multi-Winner Approval Voting. In *AAMAS*, pages 107–115, 2015.

- [Barman *et al.*, 2020] Siddharth Barman, Omar Fawzi, Suprovat Ghoshal, and Emirhan Gürpınar. Tight Approximation Bounds for Maximum Multi-Coverage. In *IPCO*, pages 66–77, 2020.
- [Bredereck *et al.*, 2020] Robert Bredereck, Piotr Faliszewski, Andrzej Kaczmarczyk, Dusan Knop, and Rolf Niedermeier. Parameterized Algorithms for Finding a Collective Set of Items. In *AAAI*, 2020.
- [Brill *et al.*, 2018] Markus Brill, Jean-François Laslier, and Piotr Skowron. Multiwinner Approval Rules as Apportionment Methods. *J. Theor. Politics*, 30(3):358–382, 2018.
- [Byrka *et al.*, 2017] Jarosław Byrka, Thomas Pensyl, Bartosz Rybicki, Aravind Srinivasan, and Khoa Trinh. An Improved Approximation for k -Median and Positive Correlation in Budgeted Optimization. *ACM Trans. Algorithms*, 13(2):23:1–23:31, 2017.
- [Byrka *et al.*, 2018] Jarosław Byrka, Piotr Skowron, and Krzysztof Sornat. Proportional Approval Voting, Harmonic k -Median, and Negative Association. In *ICALP*, pages 26:1–26:14, 2018.
- [Călinescu *et al.*, 2011] Grăia Călinescu, Chandra Chekuri, Martin Pál, and Jan Vondrák. Maximizing a Monotone Submodular Function Subject to a Matroid Constraint. *SIAM J. Comput.*, 40(6):1740–1766, 2011.
- [Chalermsook *et al.*, 2017] Parinya Chalermsook, Marek Cygan, Guy Kortsarz, Bundit Laekhanukit, Pasin Manurangsi, Danupon Nanongkai, and Luca Trevisan. From Gap-ETH to FPT-Inapproximability: Clique, Dominating Set, and More. In *FOCS*, pages 743–754, 2017.
- [Chamberlin and Courant, 1983] John R. Chamberlin and Paul N. Courant. Representative Deliberations and Representative Decisions: Proportional Representation and the Borda Rule. *Am. Political Sci. Rev.*, 77:718–733, 9 1983.
- [Cohen-Addad *et al.*, 2019] Vincent Cohen-Addad, Anupam Gupta, Amit Kumar, Euiwoong Lee, and Jason Li. Tight FPT Approximations for k -Median and k -Means. In *ICALP*, pages 42:1–42:14, 2019.
- [Cygan *et al.*, 2018] Marek Cygan, Łukasz Kowalik, Arkadiusz Socała, and Krzysztof Sornat. Approximation and Parameterized Complexity of Minimax Approval Voting. *J. Artif. Intell. Res.*, 63:495–513, 2018.
- [Faliszewski *et al.*, 2017] Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. Multiwinner Rules on Paths From k -Borda to Chamberlin-Courant. In *IJCAI*, pages 192–198, 2017.
- [Faliszewski *et al.*, 2018] Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. Multiwinner Analogues of the Plurality Rule: Axiomatic and Algorithmic Perspectives. *Soc. Choice Welf.*, 51(3):513–550, 2018.
- [Feige, 1998] Uriel Feige. A Threshold of $\ln(n)$ for Approximating Set Cover. *J. ACM*, 45(4):634–652, 1998.
- [Fluschnik *et al.*, 2019] Till Fluschnik, Piotr Skowron, Mervin Triphaus, and Kai Wilker. Fair Knapsack. In *AAAI*, pages 1941–1948, 2019.
- [Guha and Khuller, 1999] Sudipto Guha and Samir Khuller. Greedy Strikes Back: Improved Facility Location Algorithms. *J. Algorithms*, 31(1):228–248, 1999.
- [Jain *et al.*, 2020] Pallavi Jain, Krzysztof Sornat, and Nimrod Talmon. Participatory Budgeting with Project Interactions. In *IJCAI*, 2020.
- [Janson, 2016] Svante Janson. Phragmén’s and Thiele’s Election Methods. *CoRR*, abs/1611.08826, 2016.
- [Khuller *et al.*, 1999] Samir Khuller, Anna Moss, and Joseph Naor. The Budgeted Maximum Coverage Problem. *Inf. Process. Lett.*, 70(1):39–45, 1999.
- [Kilgour, 2010] D. Marc Kilgour. Approval Balloting for Multi-Winner Elections. In Jean-François Laslier and Remzi M. Sanver, editors, *Handbook on Approval Voting*, pages 105–124. Springer, 2010.
- [Lackner and Skowron, 2019] Martin Lackner and Piotr Skowron. A Quantitative Analysis of Multi-Winner Rules. In *IJCAI*, pages 407–413, 2019.
- [Manurangsi, 2020] Pasin Manurangsi. Tight Running Time Lower Bounds for Strong Inapproximability of Maximum k -Coverage, Unique Set Cover and Related Problems (via t -Wise Agreement Testing Theorem). In *SODA*, pages 62–81, 2020.
- [Nemhauser *et al.*, 1978] George L. Nemhauser, Laurence A. Wolsey, and Marshall L. Fisher. An Analysis of Approximations for Maximizing Submodular Set Functions - I. *Math. Program.*, 14(1):265–294, 1978.
- [Procaccia *et al.*, 2008] Ariel D. Procaccia, Jeffrey S. Rosenschein, and Aviv Zohar. On the Complexity of Achieving Proportional Representation. *Soc. Choice Welf.*, 30(3):353–362, 2008.
- [Skowron and Faliszewski, 2017] Piotr Skowron and Piotr Faliszewski. Chamberlin-Courant Rule with Approval Ballots: Approximating the MaxCover Problem with Bounded Frequencies in FPT Time. *J. Artif. Intell. Res.*, 60:687–716, 2017.
- [Skowron *et al.*, 2016] Piotr Skowron, Piotr Faliszewski, and Jérôme Lang. Finding a Collective Set of Items: From Proportional Multirepresentation to Group Recommendation. *Artif. Intell.*, 241:191–216, 2016.
- [Sviridenko, 2004] Maxim Sviridenko. A Note on Maximizing a Submodular Set Function Subject to a Knapsack Constraint. *Oper. Res. Lett.*, 32(1):41–43, 2004.
- [Swamy and Shmoys, 2008] Chaitanya Swamy and David B. Shmoys. Fault-Tolerant Facility Location. *ACM Trans. Algorithms*, 4(4):51:1–51:27, 2008.
- [Talmon and Faliszewski, 2019] Nimrod Talmon and Piotr Faliszewski. A Framework for Approval-Based Budgeting Methods. In *AAAI*, pages 2181–2188, 2019.
- [Thiele, 1895] Thorvald Thiele. Om Flerfoldsvalg. In *Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandlinger (in Danish)*, pages 415–441. København: A.F. Høst, 1895.