

Optimal Algorithms for Multiwinner Elections and the Chamberlin-Courant Rule

KAMESH MUNAGALA, Computer Science Department, Duke University

ZEYU SHEN, Computer Science Department, Duke University

KANGNING WANG, Computer Science Department, Duke University

We consider the algorithmic question of choosing a subset of candidates of a given size k from a set of m candidates, with knowledge of voters' ordinal rankings over all candidates. We consider the well-known and classic scoring rule for achieving diverse representation: the Chamberlin-Courant (CC) or 1-Borda rule, where the score of a committee is the average over the voters, of the rank of the best candidate in the committee for that voter; and its generalization to the average of the top s best candidates, called the s -Borda rule.

Our first result is an improved analysis of the natural and well-studied greedy heuristic. We show that greedy achieves a $\left(1 - \frac{2}{k+1}\right)$ -approximation to the maximization (or satisfaction) version of CC rule, and a $\left(1 - \frac{2s}{k+1}\right)$ -approximation to the s -Borda score. This significantly improves the existing submodularity-based analysis of the greedy algorithm that only shows a $(1 - 1/e)$ -approximation. Our result also improves on the best known approximation algorithm for this problem. We achieve this result by showing that the average dissatisfaction score for the greedy algorithm is at most $2\frac{m+1}{k+1}$ for the CC rule, and at most $2s^2\frac{m+1}{k+1}$ for s -Borda. We show these dissatisfaction score bounds are tight up to constants, and even the constant factor of 2 in the case of the CC rule is almost tight.

For the dissatisfaction (or minimization) version of the problem, it is known that the average dissatisfaction score of the best committee cannot be approximated in polynomial time to within any constant factor when s is a constant (under standard computational complexity assumptions). As our next result, we strengthen this to show that the score of $\frac{m+1}{k+1}$ can be viewed as an *optimal benchmark* for the CC rule, in the sense that it is essentially the best achievable score of any polynomial-time algorithm even when the optimal score is a polynomial factor smaller. We show that another well-studied algorithm for this problem, called the Banzhaf rule, attains this benchmark.

We finally show that for the s -Borda rule, when the optimal value is small, these algorithms can be improved by a factor of $\Omega(\sqrt{s})$ via LP rounding. Our upper and lower bounds are a significant improvement over previous results, and taken together, not only enable us to perform a finer comparison of greedy algorithms for these problems, but also provide analytic justification for using such algorithms in practice.

CCS Concepts: • **Theory of computation** → **Algorithmic game theory**; *Pseudorandomness and derandomization*; Problems, reductions and completeness; Rounding techniques; • **Mathematics of computing** → *Combinatoric problems*; *Probabilistic algorithms*.

Additional Key Words and Phrases: multiwinner elections, Chamberlin-Courant rule, approximation algorithms

ACM Reference Format:

Kamesh Munagala, Zeyu Shen, and Kangning Wang. 2021. Optimal Algorithms for Multiwinner Elections and the Chamberlin-Courant Rule. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC '21), July 18–23, 2021, Budapest, Hungary*. ACM, New York, NY, USA, 21 pages. <https://doi.org/10.1145/3465456.3467624>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC '21, July 18–23, 2021, Budapest, Hungary

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8554-1/21/07...\$15.00

<https://doi.org/10.1145/3465456.3467624>

1 INTRODUCTION

Multiwinner elections are a classical problem in social choice. In this problem, the goal is to find a set of candidates (or winning committee) of fixed size from voter preferences over the candidates. Indeed, some of the earliest work on the design of voting rules that map individual preferences to a winning committee dates back at least a century [26].

Multiwinner elections clearly arise in choosing a winning parliament in representative democracies. They have also recently found applications in design of systems for making procurement or hiring decisions [21, 24], and in participatory budgeting [16, 23]. In these settings, the candidates are products or public projects that provide shared utility to individuals. An entity such as a company or a city government has to decide, based on individual preferences, which of these projects or products to produce subject to a cardinality constraint.

Much of the work on multiwinner elections has focused on the question of *proportional* or *diverse representation*: How can we choose a winning committee where every voter feels they have some representation? Indeed, classic voting rules such as Proportional Approval Voting (PAV) [26] or Single Transferable Voting (STV) [27] explicitly attempt to enforce such representation.

In this paper, we consider the question of choosing a committee of fixed cardinality k from a set C of m candidates, when voters express *ordinal rankings* over these candidates. In many applications, including parliamentary democracies or participatory budgeting, it is reasonable to assume voters can compare candidates or projects and hence can rank them ordinally, while they may not be able to articulate cardinal utilities for the same.

A classic set of objectives for ensuring diverse representation [5] based on ordinal preferences uses the so-called *Borda score*. In the minimization (or dissatisfaction) version, the Borda score of candidate c for voter v , denoted $r_v(c)$, is the ordinal rank of c in v 's ranking. Here, the top-ranked candidate has Borda score 1, and the bottom-ranked candidate has score m .¹ Let $\mathcal{V}$ denote the set of all voters, with $n = |\mathcal{V}|$. Given a committee (that is, a set of candidates) T of size k , the s -Borda score of this committee (for $s \leq k$) is given by

$$r_{\mathcal{V}}(T) = \frac{1}{n} \sum_{v \in \mathcal{V}} \left(\min_{Q \subseteq T, |Q|=s} \sum_{c \in Q} r_v(c) \right). \quad (1)$$

Throughout the paper, we will denote the minimum possible score as $\text{OPT} = \min_{T \subseteq C, |T|=k} r_{\mathcal{V}}(T)$.

To interpret the above score, for each voter, consider the s candidates in T whose Borda score is the smallest. Now, take the sum of these scores, and average it over all the voters. Therefore, the s -Borda score assumes each voter is represented by the s best candidates in T according to her ranking, so that optimizing this score implies a form of proportional representation, where each voter on average has s “good” candidates representing her.

Our goal in this paper is to study the computational complexity (in m, n) of finding good committees according to the s -Borda score function. In particular (though not exclusively), we focus on the analysis of greedy algorithms, which are appealing for their simplicity and ease of use, especially in settings involving human decision making, such as parliamentary elections or participatory budgeting with ordinal preferences.

1.1 Results for 1-Borda Score (Chamberlin-Courant Rule)

Our main results focus on the canonical case where $s = 1$. This case has been extensively studied in computational social choice [5, 12, 13, 21, 24, 25], starting with the work of Chamberlin and

¹Existing literature also uses a score of 0 for the best ranked and $m - 1$ for the worst ranked candidates. Since our results concern absolute scores, they carry over to this setting by simply subtracting 1 from the bounds. We use a minimum score of 1 since it is the more challenging setting for showing hardness results.

Courant [7]. Here, each voter v is represented by candidate $\arg \min_{c \in T} r_v(c)$, that is, the most preferred candidate from T in v 's ordering. The score of the voter is the rank of its representative, and the goal is to minimize the average of this score over the voters. This rule is also called the *Chamberlin-Courant* voting rule, though we will henceforth call it the 1-Borda score for consistency with the generalizations we study later.

The 1-Borda score is an ordinal version of the celebrated k -medians problem [1]. Unfortunately, for the ordinal version, it is not possible to approximate the minimum score, OPT , to any constant factor in polynomial time unless $P = NP$ [25].

The GREEDY Algorithm. A natural algorithm for the 1-Borda score is the GREEDY algorithm that iteratively adds the candidate that decreases the 1-Borda score the most. This algorithm was analyzed in [21] as follows. Consider the maximization (or satisfaction) version where the score of a candidate c for voter v is $m + 1 - r_v(c)$, so that the score of a committee T is $m + 1 - r_V(T)$. Clearly, the maximization and minimization versions have the same optimum solutions, though they are very different from an approximation perspective. It is easy to check that the maximization objective is submodular [21], so that GREEDY is a $(1 - \frac{1}{e})$ -approximation by the classic result of [22]. However, this analysis only shows that GREEDY yields a solution of score at most m/e for the minimization objective.

Our first, and technically most challenging, contribution is an almost-tight analysis of this GREEDY heuristic for the minimization version. In Section 3, we show that it achieves a score (given by Eq. (1)) of at most $2 \cdot \frac{m+1}{k+1}$ for any instance with m candidates from which we need to choose a committee of size k . We complement this analysis by exhibiting an instance where GREEDY has score at least $1.962 \cdot \frac{m+1}{k+1}$.

For the maximization (or satisfaction) version, since the maximum possible score is m , the above result directly implies the following theorem.

THEOREM 1.1. *GREEDY is a $(1 - \frac{2}{k+1})$ -approximation for the maximization version of 1-Borda score.*

For k larger than a small constant, this *significantly improves* the submodularity-based analysis [21] that only yields a $(1 - 1/e)$ -approximation. Furthermore, it also improves on the best known approximation algorithm for this problem, Algorithm P in [25], which achieves an approximation factor of $(1 - O(\frac{\ln k}{k}))$.

At a technical level, the standard analysis of GREEDY for maximizing submodular functions shows that the next candidate yields an improvement in objective that is at least $1/k$ fraction of the gap between the current solution and the optimum. We use Cauchy-Schwarz inequality on per-voter improvements to show an overall improvement per step that has a *quadratic* dependence on the gap. This yields a significant improvement when the gap is large, and is the crux of why we are able to improve the upper-bound analysis of the maximization version significantly. Our lower bound instance works by carefully choosing per-voter improvements that make Cauchy-Schwarz inequality almost tight. This requires a non-trivial construction where the candidates chosen by GREEDY have ranks that lie on a carefully chosen spiral, and these are interspersed with candidates for whom voters' preferences are random. The ranks in each subsequent layer of the spiral decrease by a factor equal to the golden ratio.

A Benchmark and an Optimal Algorithm. The next natural question we ask is: How much better can we do in polynomial time? In Section 4, we show some hardness results for the minimization version (Eq (1)). Our main result *significantly improves* the constant factor hardness of approximation result of [25] and shows the following:

THEOREM 1.2. *Unless $ZPP = NP$, no polynomial-time algorithm can distinguish between instances with $OPT \geq (1 - o(1)) \cdot \frac{m+1}{k+1}$ from those with either:*

- (1) $OPT \leq \left(\frac{m+1}{k+1}\right)^\delta$, where $\delta \in (0, 1)$ is a constant; or
- (2) $OPT \leq \frac{1}{k^\alpha} \cdot \frac{m+1}{k+1}$, where $\alpha > 0$ is a constant.

This construction yielding this theorem is delicate. We require the full power of Feige’s hardness proof of MAX COVER [14], in particular, that it works on “regular” instances where each set has the same size, and where a collection of disjoint sets cover the instance completely in the “YES” case.

Theorem 1.2 motivates us to define the score $\frac{m+1}{k+1}$ as a reasonable benchmark for this problem, and we call any efficient algorithm achieving this score as an “optimal algorithm”. Such a benchmark is appealing in that it helps us analyze other simple and natural algorithms that have been proposed in literature, and perform a more fine-grained comparison. As we have already seen, the GREEDY algorithm is always within a factor of 2 of this benchmark.

We now observe that if we pick a subset of k candidates at random from C , the expected score is exactly the benchmark $\frac{m+1}{k+1}$. We therefore denote the score $\frac{m+1}{k+1}$ as RAND. Now, we can design a deterministic optimal algorithm via derandomizing this randomized algorithm. Interestingly, we show that this derandomization yields a greedy algorithm that is exactly the same as the BANZHAF algorithm proposed in [13] as a polynomial time heuristic for this problem. In that work, the BANZHAF algorithm was derived by viewing the problem as a cooperative game where players are candidates, and coalitions are committees, and adapting the notion of Banzhaf score of coalitions [3, 11]. It was empirically shown to be a very effective heuristic for this problem, beating GREEDY on most instances. We justify this empirical observation by viewing the BANZHAF algorithm instead as a derandomization of an optimal randomized algorithm.

In summary, we show the following theorem.

THEOREM 1.3. *The BANZHAF algorithm achieves a minimization objective of at most $RAND = \frac{m+1}{k+1}$ in polynomial time, and is a $(1 - \frac{1}{k+1})$ -approximation to the maximization objective of 1-Borda.*

To complete the picture, we show that an easy consequence of Theorem 1.2 is that the approximation factor of $(1 - \frac{1}{k+1})$ is best possible for the maximization version unless $NP = ZPP$.

Committee Monotonicity. One appealing property of GREEDY is that it is *committee-monotone* [12]: The committee found for a smaller k is always a subset of a committee found for larger k ’s. This is immediate because GREEDY adds the next candidate to the committee based on the improvement in the 1-Borda score, and this improvement does not depend on k . On the other hand, the BANZHAF algorithm requires knowledge of k at each greedy step, and is therefore not committee-monotone. We therefore ask: Is there a committee-monotone algorithm that can achieve the benchmark RAND? In Section 5, we answer this question in the negative: There exist instances where any committee-monotone algorithm has score at least $1.015 \cdot RAND$. This shows a separation between committee-monotone algorithms and an optimal algorithm such as the BANZHAF algorithm.

Connection to the Core. The notion of core from cooperative game theory is appealing as a notion of fairness, and provides a strong notion of proportionality. Informally, in a core solution, every reasonably large subgroup of voters is happy in the sense that they do not all prefer the same candidate outside the chosen committee. Formally, the work of [9, 19] defines an α -approximate core as follows. Fix some $\alpha \geq 1$. Given a committee T of size k , a candidate c is blocking if at least $\alpha \cdot \frac{n}{k}$ voters prefer c to any candidate in T , that is,

$$\left| \left\{ v \in \mathcal{V} \mid r_v(c) < \min_{c' \in T} r_v(c') \right\} \right| \geq \alpha \cdot \frac{n}{k}. \quad (2)$$

A committee T is in the α -approximate core if it does not admit a blocking candidate. The work of [9, 19] shows that a 16-approximate core always exists and can be computed in polynomial time, while a $(2 - \epsilon)$ -approximate core is not guaranteed to exist for any constant $\epsilon > 0$.

In the full version, we show that the core indeed achieves a stronger notion of proportionality than the 1-Borda score in the following sense: Any α -approximate core solution has 1-Borda score at most $\alpha(1 + 1/k) \cdot \text{RAND}$. The converse of this statement is however false: None of OPT, GREEDY, or Banzhaf lies in an α -approximate core for any constant α .

1.2 Results for s -Borda Score

We next consider the s -Borda score for $1 < s \leq k$. We start with an analysis of the natural extensions to the greedy algorithms considered above for $s = 1$. It is easy to show that choosing a random committee of size k yields expected score $\frac{s(s+1)}{2} \cdot \text{RAND}$, where as before, $\text{RAND} = \frac{m+1}{k+1}$. This implies its derandomization – the Banzhaf algorithm – has score at most $\frac{s(s+1)}{2} \cdot \text{RAND}$. Furthermore, there are instances where the best possible score $\text{OPT} \geq \frac{s(s+1)}{2} \cdot \text{RAND}$.

Analysis of GREEDY. The GREEDY algorithm extends naturally to this setting. In Section 6.1, we extend the result in Section 3 to show that GREEDY achieves an s -Borda score of at most $2s^2 \cdot \text{RAND}$, which is within a factor of $\frac{4s}{s+1}$ of the upper bound for the Banzhaf algorithm.

For the maximization version, recall that the score of candidate c for voter v is $m + 1 - r_c(v)$, and the voter's score for a committee is the sum of top s candidate scores. Since the maximum possible score is at most ms , this directly implies the following theorem. For $s = o(k)$, this again *significantly improves* on the classic submodularity-based analysis that only shows a $(1 - \frac{1}{e})$ -approximation.

THEOREM 1.4. *GREEDY is a $(1 - \frac{2s}{k+1})$ -approximation for the maximization version of s -Borda.*

Note that for the related maximum multi-cover problem [4], the approximation factor of $(1 - 1/e)$ is actually tight for GREEDY. Therefore, our analysis of GREEDY points to fundamental algorithmic differences between MAX MULTI-COVER (resp. MAX COVER [14]) and s -Borda (resp. 1-Borda), since we obtain significantly better factors for the latter.

Improved Algorithm. In contrast with the $s = 1$ case, for larger values of s , we can obtain a non-trivial improvement over these greedy algorithms (for the minimization version of s -Borda) by using the natural LP relaxation for this problem [6]. In Section 6.2, we devise a randomized algorithm that is based on carefully combining dependent rounding of this LP solution with choosing a committee by uniform random sampling. We show that this algorithm achieves expected score

$$\text{ALG} \leq 3 \cdot \text{OPT} + O\left(s^{3/2} \log s\right) \cdot \text{RAND},$$

where $\text{OPT} = \min_{T \subseteq C, |T|=k} r_V(T)$ is the minimum possible s -Borda score of any committee.

This result improves on the aforementioned bounds for greedy algorithms when OPT is small. For instance, if $\text{OPT} = O(s^{3/2}) \cdot \text{RAND}$, the improvement is $\tilde{\Omega}(\sqrt{s})$. We note that such an improved bound cannot be achieved by GREEDY when OPT is small: In Section 6.1, we show instances where $\text{OPT} = o(1) \cdot \text{RAND}$, while the score of GREEDY is $\Omega(s^2) \cdot \text{RAND}$. On the flip side, our improved bound is based on solving and rounding an LP, and is therefore not as simple or intuitive as the GREEDY or Banzhaf algorithms.

1.3 Related Work

The literature on multiwinner elections is too vast to survey here. We present a survey of computational results in this space to place our work in context.

Suppose the candidates and voters are embedded in a metric space, and suppose the score $r_v(c)$ is not the Borda score, but instead the metric distance between voter v and candidate c . Then the objective for the $s = 1$ case is precisely the celebrated k -medians objective [1, 8, 18, 20], while the general- s case has been studied as fault-tolerant k -medians [17]. For both these problems, constant-factor approximation algorithms are known. The versions we consider can therefore be viewed as ordinal versions of k -medians and fault tolerant k -medians respectively. Towards showing better bounds for the ordinal versions, it is tempting to impose a condition such as the ordinal preferences of voters should correspond to distances in some underlying metric space. However, it is easy to show that given any set of ordinal preferences, there is a metric space that can realize these preferences, which means this assumption does not help. Nevertheless, the LP relaxation we use to derive improved bounds for the s -Borda score is the same as the standard LP relaxation for the (fault-tolerant) k -medians objective [17]. It is an interesting open question to explore what other natural assumptions on voter preferences will lead to improved upper bounds for the 1-Borda and s -Borda objectives.

The work of [12] considers generalizations of the 1-Borda and s -Borda scores to *committee scoring rules*. A committee scoring rule is a function that for each voter v and committee T of size k , maps the set of ordinal ranks $\{r_v(c), c \in T\}$ to a score. The s -Borda score we consider is an example of a decomposable rule, meaning that the score can be written as a sum of contributions from the committee members. The work of [24, 25] defines a special case of committee scoring rules where the score is a weighted sum of ranks of the committee members. They call these Ordered Weighted Average (OWA) rules. Again, it is easy to see that the s -Borda rule is an OWA rule. For the maximization version of committee scoring rules, the GREEDY algorithm continues to be a $(1 - \frac{1}{e})$ -approximation via submodularity. It is an interesting open question to extend the results in this paper to other rules that achieve diverse or proportional representation [2, 5].

Finally, the work of [6] considers the variant where the score $r_v(c)$ is an arbitrary cardinal value, which is different from our focus on ordinal preferences. They consider the “harmonic” OWA rule where a voter assigns weight 1 to candidate in the committee with lowest score, weight $\frac{1}{2}$ to the candidate with second lowest score, and so on, till weight $\frac{1}{k}$ to the candidate with highest score. They show this version has a constant-factor approximation algorithm by randomized rounding of the natural LP relaxation, first used in [10]. The difficulty with the s -Borda rule is that the weight jumps discretely from 1 to 0 when we move from the top s candidates for a voter to the $(s + 1)^{\text{st}}$ candidate. This discontinuity is most pronounced for $s = 1$, and leads to our strong impossibility result. In essence, this discontinuity is what motivates us to consider an alternate benchmark to analyse the performance of natural greedy algorithms assuming ordinal preferences.

2 PRELIMINARIES

We consider the problem of selecting a subset of cardinality k from a set C of m candidates. We call this subset a *committee*. A set $\mathcal{V}$ of n voters express their preferences on the candidates ordinally. Each voter v has a bijective ranking function $r_v : C \rightarrow \{1, 2, \dots, m\}$, and v prefers those c 's with smaller $r_v(c)$. For example, the top-ranked candidate of v , denoted by $c_{\text{top}(v)}$, satisfies $r_v(c_{\text{top}(v)}) = 1$, and the bottom-ranked $c_{\text{bot}(v)}$ satisfies $r_v(c_{\text{bot}(v)}) = m$.

In s -Borda score, the cost for a voter v of a committee T is the sum of her ranks of the top s candidates in T : $r_v(T) = \min_{Q \subseteq T, |Q|=s} \sum_{c \in Q} r_v(c)$. Further, the s -Borda score ($s \leq k$) of a committee T is the average cost for all voters:

$$r_{\mathcal{V}}(T) = \frac{1}{n} \sum_{v \in \mathcal{V}} r_v(T) = \frac{1}{n} \sum_{v \in \mathcal{V}} \left(\min_{Q \subseteq T, |Q|=s} \sum_{c \in Q} r_v(c) \right).$$

In particular, when $s = 1$, $r_{\mathcal{V}}(T) = \frac{1}{n} \sum_{v \in \mathcal{V}} \min_{c \in T} r_v(c)$.

Fix a voter v and look at her ranking on C . If we pick a random size- k subset of C , the t^{th} smallest rank is $t \cdot \frac{m+1}{k+1}$ in expectation. (See the full version for a proof of this well-known fact.) Therefore, the expected performance of a random committee is

$$\mathbb{E}_{T \subseteq C} [r_{\mathcal{V}}(T)] = \frac{1}{n} \sum_{v \in \mathcal{V}} \mathbb{E}_{T \subseteq C} [r_v(T)] = \sum_{t=1}^s t \cdot \frac{m+1}{k+1} = \frac{s(s+1)}{2} \cdot \frac{m+1}{k+1}.$$

Define the benchmark RAND to be the expected performance of a random committee when $s = 1$ as $\text{RAND} = \frac{m+1}{k+1}$. We will justify this benchmark in the subsequent sections.

We consider two simple committee-selection rules: GREEDY and BANZHAF. These algorithms run in k iterations, during which they build sets $\emptyset = T_0 \subsetneq T_1 \subsetneq \dots \subsetneq T_k$, and declare T_k as the selected committee.

In the j^{th} iteration, GREEDY picks candidate $c_j \in C \setminus T_{j-1}$ that minimizes $r_{\mathcal{V}}(T_{j-1} \cup \{c_j\})$, and let $T_j = T_{j-1} \cup \{c_j\}$. BANZHAF [3, 13], on the other hand, picks candidate $c_j \in C \setminus T_{j-1}$ to minimize

$$\sum_{\substack{S \subseteq C: |S|=k \\ S \supseteq T_{j-1} \cup \{c_j\}}} r_{\mathcal{V}}(S)$$

in the j^{th} iteration, and then sets $T_j = T_{j-1} \cup \{c_j\}$. In other words, it greedily picks the candidate that minimizes the final score if the rest of the committee is chosen uniformly at random. Both GREEDY and BANZHAF can run in polynomial time [13].

Throughout the paper, we use RAND, GREEDY and BANZHAF to denote either the algorithms or their performances, which should be clear from the context.

3 ANALYSIS OF GREEDY FOR 1-BORDA

In this section, we analyze the performance of GREEDY, evaluated with respect to the benchmark RAND. Throughout this section, we only consider the 1-Borda score, i.e., $s = 1$. We first show an upper bound that $\text{GREEDY} \leq 2 \cdot \text{RAND}$, and then present an almost-matching lower-bound instance where $\text{GREEDY} > 1.962 \cdot \text{RAND}$.

3.1 Upper Bound

Now we show $\text{GREEDY} \leq 2 \cdot \text{RAND}$ as an upper bound. We first present the following lemma, which gives a lower bound on the improvement at each iteration.

LEMMA 3.1. *Let T_t and T_{t+1} be the set of candidates produced by GREEDY in the t^{th} and $(t+1)^{\text{st}}$ iterations, and $r_{\mathcal{V}}(T_t)$, $r_{\mathcal{V}}(T_{t+1})$ be their respective score. We have:*

$$r_{\mathcal{V}}(T_t) - r_{\mathcal{V}}(T_{t+1}) \geq \frac{\sum_{v \in \mathcal{V}} r_v(T_t)(r_v(T_t) - 1)}{2n(m-t)}.$$

PROOF. For a candidate $c \notin T_t$, define $\Delta_c := r_{\mathcal{V}}(T_t) - r_{\mathcal{V}}(T_t \cup \{c\})$, i.e., the current marginal contribution of c to the 1-Borda score. Taking the sum of Δ_c over $c \notin T_t$:

$$\sum_{c \in C \setminus T_t} \Delta_c = \frac{1}{n} \sum_{v \in \mathcal{V}} \sum_{j=1}^{r_v(T_t)-1} j = \frac{\sum_{v \in \mathcal{V}} r_v(T_t)(r_v(T_t) - 1)}{2n}.$$

GREEDY chooses $c^* = \arg \max_c \Delta_c$ at the $(t+1)^{\text{st}}$ iteration, giving us

$$r_{\mathcal{V}}(T_t) - r_{\mathcal{V}}(T_{t+1}) = \Delta_{c^*} \geq \frac{1}{m-t} \sum_{c \in C \setminus T_t} \Delta_c = \frac{\sum_{v \in \mathcal{V}} r_v(T_t)(r_v(T_t) - 1)}{2n(m-t)}. \quad \square$$

Now we prove our upper bound of 2.

THEOREM 3.2. $\text{GREEDY} \leq 2 \cdot \text{RAND}$.

PROOF. We prove by induction. As the base case where $k = 1$, $\text{GREEDY} \leq m < m + 1 = 2 \cdot \text{RAND}$. Now suppose that the claim holds for some $k - 1$ and we will prove that it also holds for k . By induction hypothesis, we have:

$$r_{\mathcal{V}}(T_{k-1}) \leq 2 \cdot \frac{m+1}{k}.$$

If $r_{\mathcal{V}}(T_{k-1}) \leq 2 \cdot \frac{m+1}{k+1}$, then $r_{\mathcal{V}}(T_k) \leq r_{\mathcal{V}}(T_{k-1}) \leq 2 \cdot \frac{m+1}{k+1}$ finishes the proof. Thus, we only need to consider the following case:

$$2 \cdot \frac{m+1}{k+1} < r_{\mathcal{V}}(T_{k-1}) \leq 2 \cdot \frac{m+1}{k}.$$

We now have the following, where the first inequality is by Lemma 3.1 and second by Cauchy-Schwarz inequality:

$$\begin{aligned} r_{\mathcal{V}}(T_{k-1}) - r_{\mathcal{V}}(T_k) &\geq \frac{\sum_{v \in \mathcal{V}} r_v(T_{k-1})(r_v(T_{k-1}) - 1)}{2n(m - k + 1)} \\ &\geq \frac{\frac{1}{n}(\sum_{v \in \mathcal{V}} r_v(T_{k-1}))^2 - \sum_{v \in \mathcal{V}} r_v(T_{k-1})}{2n(m - k + 1)} \\ &= \frac{(\sum_{v \in \mathcal{V}} r_v(T_{k-1}))^2}{2n^2(m + 1)} \cdot \frac{m + 1}{m - k + 1} \cdot \frac{\sum_{v \in \mathcal{V}} (r_v(T_{k-1}) - 1)}{\sum_{v \in \mathcal{V}} r_v(T_{k-1})}. \end{aligned}$$

Since $r_{\mathcal{V}}(T_{k-1}) \geq 2 \cdot \frac{m+1}{k+1}$ by assumption, we have:

$$\begin{aligned} \frac{m + 1}{m - k + 1} \cdot \frac{\sum_{v \in \mathcal{V}} (r_v(T_{k-1}) - 1)}{\sum_{v \in \mathcal{V}} r_v(T_{k-1})} &\geq \frac{m + 1}{m - k + 1} \cdot \frac{2 \cdot \frac{m+1}{k+1} - 1}{2 \cdot \frac{m+1}{k+1}} \\ &= \frac{2(m + 1) - k - 1}{2(m + 1) - 2k} \geq 1. \end{aligned}$$

Combining the previous two inequalities, we therefore have:

$$r_{\mathcal{V}}(T_{k-1}) - r_{\mathcal{V}}(T_k) \geq \frac{(\sum_{v \in \mathcal{V}} r_v(T_{k-1}))^2}{2n^2(m + 1)} = \frac{r_{\mathcal{V}}^2(T_{k-1})}{2(m + 1)},$$

which is equivalent to:

$$r_{\mathcal{V}}(T_k) \leq -\frac{1}{2(m + 1)} r_{\mathcal{V}}^2(T_{k-1}) + r_{\mathcal{V}}(T_{k-1}).$$

Notice that the right hand side is a quadratic function in $r_{\mathcal{V}}(T_{k-1})$, which is monotonically increasing for $r_{\mathcal{V}}(T_{k-1}) \leq m + 1$. Since $r_{\mathcal{V}}(T_{k-1}) \leq 2 \cdot \frac{m+1}{k} \leq m + 1$, the right hand side reaches its maximum at $2 \cdot \frac{m+1}{k}$. Thus, we have:

$$r_{\mathcal{V}}(T_k) \leq -\frac{1}{2(m + 1)} \cdot \left(\frac{2(m + 1)}{k} \right)^2 + \frac{2(m + 1)}{k} \leq \frac{2(m + 1)}{k + 1},$$

which concludes our induction. $\square$

Proof of Theorem 1.1. For the maximization version, the above result implies GREEDY achieves score at least $(m + 1) \cdot (1 - \frac{2}{k+1})$. Since the maximum possible score is m , this implies that GREEDY is a $(1 - \frac{2}{k+1})$ -approximation.

3.2 Lower Bound

Now we complement our result with a lower-bound example for GREEDY.

THEOREM 3.3. *There exists an instance in which $r_V(T_k) > 1.962 \cdot \text{RAND}$.*

Construction. In the sequel, we will prove the above theorem. In the instance we construct, m , n , and k are all sufficiently large. For convenience of illustration, we scale down the ranks by a factor of m : now the ranks are $\frac{1}{m}, \frac{2}{m}, \dots, \frac{m-1}{m}, 1$. As $m \rightarrow \infty$, $\frac{1}{m} \rightarrow 0$, so the set of ranking $\{\frac{1}{m}, \frac{2}{m}, \dots, 1\}$ will become dense in $[0, 1]$, and thus we regard the ranking as being continuous from 0 to 1. Our goal becomes to construct an instance in which GREEDY gives $r_V(T_k) > 1.962 \cdot \frac{1}{k+1}$.

There are sufficiently many voters, enabling us to view them as a continuum from 0 to 1, forming a circle (the base in Fig. 1) with angular position ranging from 0 to 2π . Imagine that each voter writes down her favorite, her second favorite, $\dots$, her least favorite candidate in that order vertically. The result is the side of a cylinder with height 1, as depicted in Fig. 1. Each point on the side identifies a candidate, whose distance to the top, d , indicates the corresponding voter ranks him as her $(dm)^{\text{th}}$ favorite candidate (i.e., the candidate has a rank of d in the voter's preference after scaling).




Fig. 1. Construction of the Bad Instance for GREEDY

We divide the set of candidates into two types – *critical* and *dummy*. The former set has size $k \ll m$, and the latter has size $m - k$. Our proof will show that GREEDY will choose the critical candidates in a fixed order, and will not choose any dummy candidate.

The critical candidates are present in ℓ “layers” as shown in the red spiral in Fig. 1, where ℓ is sufficiently large. This figure shows the ranks of the critical candidates in the voters’ profiles. We parametrize this spiral by θ , which maps to the voter at the corresponding angular position $2\pi\theta$. We place critical candidates in order, where each candidate appears a number of times consecutively on the spiral. Therefore, each voter has one critical candidate from each layer $t = 0, 1, \dots, \ell$ in the spiral part of its ranking.

In the t^{th} layer, the parameter θ lies in $[t - 1, t)$. The critical candidate when the parameter is θ has rank $g(\theta) = a\varphi^\theta$ for the voter at angular position $2\pi\theta$. Here, φ denotes the golden ratio $\frac{\sqrt{5}-1}{2} \approx 0.618$, and a is a sufficiently small constant so that rounding to the nearest integer does not change the analysis. This critical candidate is placed for a certain length $h(\theta)$ on the spiral,

which means this candidate appears at rank $g(\theta)$ for voters in the range $[2\pi\theta, 2\pi(\theta + h(\theta))]$. In our construction, $h(\theta)$ will be very small, so that we will say this candidate appears $h(\theta)$ times at rank $g(\theta)$ for parameter θ . The greater θ is, the smaller $h(\theta)$ has to be, and we will calculate its expression later.

For the convenience of analysis, at the layer $t = 0$, that is, for $\theta \in [-1, 0)$, there is a special candidate appearing on the spiral throughout the layer. This special candidate is picked first by GREEDY. Other than its appearance on the spiral, any critical candidate is placed at the very bottom, i.e., rank 1, for the other voters. Denote the total number of critical candidates by k . Then we have $m - k$ dummy candidates. These dummy candidates are symmetrically placed at other ranks. We copy each voter $(m - k)!$ times, once for each possible permutation of the dummy candidates to place in the remaining ranks.

The idea of this construction is to trick GREEDY into picking every critical candidate on the spiral in order, while in fact, lower-layer critical candidates have no contribution to the objective once higher-layer ones have been selected. The following analysis computes the optimal parameters to realize this plan.

Not Choosing a Dummy Candidate. We first ensure GREEDY does not choose a dummy candidate in this instance by setting $h(\theta)$ properly. We assume that GREEDY chooses critical candidates in increasing order of θ , and we will justify this assumption later.

To simplify notation, denote $X = \int_0^1 a\varphi^\theta d\theta$ and $Y = \int_0^1 a^2\varphi^{2\theta} d\theta$. Computing these explicitly:

$$X = \frac{a}{\ln \varphi}(\varphi - 1), \quad Y = \frac{a^2}{2 \ln \varphi}(\varphi^2 - 1) = X^2 \frac{(\varphi + 1) \ln \varphi}{2(\varphi - 1)}.$$

Using this notation, consider the critical candidate at the beginning of the first layer, that is, at $\theta = 0$. Since GREEDY chooses the candidate at layer $t = 0$, the decrease in score due to this critical candidate is:

$$h(0) \cdot (g(-1) - g(0)) = h(0) \cdot a \cdot \left(\frac{1}{\varphi} - 1 \right) = h(0) \cdot a \cdot \varphi. \quad (3)$$

where we have used that since φ is the golden ratio, $\varphi + \varphi^2 = 1$.

Now consider the dummy candidates. Just after GREEDY has chosen the special candidate at layer $t = 0$, each such candidate improves the rank of $g(\theta - 1)$ fraction of voters at $\theta \in [0, 1)$. This is because we placed all permutations of the dummy candidates with each voter θ , and GREEDY has already chosen the special candidate. By the same reasoning, conditioned on improvement, the average improvement is $g(\theta - 1)/2$. Therefore, the decrease in score due to a dummy candidate is:

$$\int_0^1 \frac{g^2(\theta - 1)}{2} d\theta = \frac{a^2}{2\varphi^2} \int_0^1 \varphi^{2\theta} d\theta = \frac{1}{2\varphi^2} \cdot Y. \quad (4)$$

Since we want GREEDY to choose the critical candidate, we need to set

$$h(0) = \frac{Y}{2\varphi^3 a}.$$

By the symmetry of the spiral, an identical calculation now holds for all $\theta > 0$. To make GREEDY choose the critical candidate at this location (assuming it has chosen critical candidates for smaller values of θ), we need:

$$h(\theta) = \frac{Y}{2\varphi^3 a} \varphi^\theta.$$

Note that $h(\theta)$ depends linearly on a , so that for very small a , we can pretend this set of voters lies exactly at θ . Further, $h(\theta)$ is decreasing with θ .

Choosing Critical Candidates in Order. We now show that GREEDY chooses the critical candidates following the order on the spiral.

LEMMA 3.4. GREEDY chooses the critical candidates in increasing order of θ .

PROOF. The calculation is identical at any step of GREEDY, so we focus on the step where GREEDY is at the beginning of the first layer, that is, considering the critical candidate at $\theta = 0$. Recall that GREEDY has chosen the special candidate at layer $t = 0$. The previous analysis showed that the critical candidate at $\theta = 0$ yields decrease of $\frac{Y}{2\varphi^2}$. For critical candidates in the same layer $t = 1$ (that is, for $\theta \in [0, 1)$), the contribution of the candidate at θ is

$$h(\theta) \cdot (g(\theta - 1) - g(\theta)) = \frac{Y}{2\varphi^3} \varphi^\theta (\varphi^{\theta-1} - \varphi^\theta) = \frac{Y}{2\varphi^2} \cdot \varphi^{2\theta},$$

which decreases with θ , so that the current candidate, $\theta = 0$, offers the best decrease. Here, we have used that since φ is the golden ratio, $\varphi^2 + \varphi = 1$.

For $t \geq 1$, suppose we instead considered a candidate $t + \theta$ for $\theta \in [0, 1)$ located in layer $t + 1$. Conditioned on having chosen layer $t = 0$, this candidate gives a contribution of

$$\begin{aligned} h(t + \theta) \cdot (g(\theta - 1) - g(t + \theta)) &\leq h(t) \cdot (g(-1) - g(t)) \\ &\leq \max(h(2) \cdot g(-1), h(1) \cdot (g(-1) - g(1))) \\ &= \max\left(\frac{Y}{2\varphi^2}, \frac{Y}{2\varphi^2 a} \cdot a \left(\frac{1}{\varphi} - \varphi\right)\right) = \frac{Y}{2\varphi^2}, \end{aligned}$$

where the first inequality uses that $h(\theta)$ is decreasing in θ , and that $\varphi < 1$.

Therefore, GREEDY will pick the critical candidate at $\theta = 0$ instead of another candidate at the same or a higher layer. Since the argument is identical at each θ , GREEDY picks critical candidates in order on the spiral. $\square$

The Lower Bound. So far we have shown that GREEDY chooses critical candidates in increasing order of layers and does not choose dummy candidates. We finally put it all together and show the following bound, which completes the proof of Theorem 3.3.

PROOF OF THEOREM 3.3. The number of critical candidates on the t^{th} layer ($\theta \in [t - 1, t)$) is

$$\int_{t-1}^t \frac{1}{h(\theta)} d\theta = \frac{2\varphi^3 a}{Y} \int_{t-1}^t \varphi^{-\theta} d\theta = \frac{2a}{\varphi^{t-3} Y} \int_{-1}^0 \varphi^{-\theta} d\theta = \frac{2a}{\varphi^{t-3} Y} \int_0^1 \varphi^\theta d\theta = \frac{2X}{\varphi^{t-3} Y}.$$

Therefore, when it is done with the ℓ^{th} layer, the number of candidates GREEDY has picked is

$$k = \frac{2X}{\varphi^{\ell-3} Y} (1 + \varphi + \varphi^2 + \dots + \varphi^{\ell-1}) \rightarrow \frac{2X}{(1 - \varphi) \varphi^{\ell-3} Y}$$

when ℓ is large. Meanwhile, the 1-Borda score of GREEDY is

$$r_V(T_k) = \int_0^1 g(\ell - 1 + \theta) d\theta = \varphi^{\ell-1} X.$$

Therefore, the approximation ratio is

$$(k+1)r_V(T_k) \geq \frac{2X}{(1 - \varphi) \varphi^{\ell-3} Y} \cdot \varphi^{\ell-1} X = \frac{2\varphi^2 X^2}{(1 - \varphi) Y} = \frac{2\varphi^2}{(1 - \varphi)} \cdot \frac{2(\varphi - 1)}{(\varphi + 1) \ln \varphi} = -\frac{4\varphi^2}{(\varphi + 1) \ln \varphi} > 1.962. \quad \square$$

4 HARDNESS OF 1-BORDA AND OPTIMAL DETERMINISTIC ALGORITHM

Throughout this section, we focus on 1-Borda score. We justify our choice of benchmark $\text{RAND} = \frac{m+1}{k+1}$, and show that a deterministic algorithm, BANZHAF, achieves this benchmark optimally. First, notice that if the input consists of one voter for each possible preference of m candidates (thus $n = m!$), picking any committee has the same 1-Borda score by symmetry, so OPT is just RAND . Thus, we have the following proposition:

PROPOSITION 4.1. *For any m and k , there exist instances where $\text{OPT} = \text{RAND}$.*

4.1 Hardness Result for 1-Borda: Theorem 1.2

We now show Theorem 1.2: Even if OPT is very small, it is computationally hard to significantly beat RAND . To prove this hardness result, we show a reduction from the decision version of the REGULAR MAX K-COVER problem.

Definition 4.2. In REGULAR MAX K-COVER, there is a universe U of n elements $\{a_1, a_2, \dots, a_n\}$, and a family $\mathcal{F} = \{S_1, S_2, \dots, S_z\}$ of subsets of U . Each S_i has the same size $\frac{n}{k}$. The *value* of an instance is the maximum size of the union of k sets from $\mathcal{F}$. For any constant $\varepsilon > 0$, we consider the following decision version:

- “YES” instances are those with value n . Therefore, there exist k disjoint sets each of size n/k that cover all the elements.
- “NO” instances are those with value at most $\frac{3}{4}n$.

The above problem known to be NP-hard to approximate via the following lemma that is implicit in the proof of Theorem 5.3 in [14].

LEMMA 4.3 ([14]). *The decision version of REGULAR MAX K-COVER from Definition 4.2 is NP-hard, that is, unless $P = NP$, there is no polynomial time algorithm that can decide always answers “YES” for “YES” instances and answers “NO” for “NO” instances.*

Note that if the instance has value n , there exist k disjoint sets each of size n/k that cover all the elements. This aspect will be crucial in our reduction. Also needed in our reduction, we state the following lemma for constructing a profile with polynomial number of voters, where the best solution with score OPT has similar performance as RAND .

LEMMA 4.4. *Fix any $\varepsilon > 0$ and let $n \geq \left\lceil \frac{m(k+1)^2}{\varepsilon^2} \right\rceil$. Consider the instance where the preference of each voter is an independent and uniformly random permutation. Let OPT' denote the expected value of the optimum score, and $\text{RAND}' = \frac{m+1}{k+1}$, then $\Pr[\text{OPT} \leq (1 - \varepsilon) \cdot \text{RAND}] < \frac{1}{2}$, where the probability is over the randomness in the permutations.*

PROOF. Fix any committee T of size k . Notice that $E[r_V(T)] = \text{RAND}'$ since the preferences are uniformly random. We have

$$\begin{aligned} \Pr[r_V(T) - \text{RAND}' \leq -\varepsilon \cdot \text{RAND}'] &= \Pr\left[\frac{1}{n} \sum_{v \in V} r_v(T) - \text{RAND}' > \varepsilon \cdot \text{RAND}'\right] \\ &\leq \exp\left(\frac{-2n(\varepsilon \cdot \text{RAND}')^2}{m^2}\right) \leq \exp\left(\frac{-2n\varepsilon^2}{(k+1)^2}\right), \end{aligned}$$

where the second step comes from Hoeffding's inequality. By union bound,

$$\begin{aligned} \Pr[\text{OPT}' \leq (1 - \varepsilon) \cdot \text{RAND}'] &\leq \binom{m}{k} \cdot \Pr[r_V(T) - \text{RAND}' \leq -\varepsilon \cdot \text{RAND}'] \\ &\leq \exp\left(\frac{-2n\varepsilon^2}{(k+1)^2} + m\right) \leq e^{-m} < \frac{1}{2}. \end{aligned} \quad \square$$

Now we are ready to prove Theorem 1.2.

PROOF OF THEOREM 1.2. Fix a $\varepsilon > 0$ and let $\varepsilon' = 10\varepsilon$. We will choose ε appropriately later. Given any instance of REGULAR MAX K-COVER with n elements and z sets each of size n/k (as in Definition 4.2), we construct the following instance for our problem:

- There are $N = nR$ voters v_{ij} where $i \in [n]$ and $j \in [R]$. We have $m = \frac{2}{\varepsilon}kz$ candidates. The first z candidates $\{c_1, c_2, \dots, c_z\}$ are “critical” candidates, and the other $m - z$ candidates are “dummy” candidates. Each voter corresponds to an element in the universe and each critical candidate corresponds to a set in REGULAR MAX K-COVER.
- If a set S_i covers a_j , then voters v_{ij} for $j \in [R]$ rank c_j within top ε' fraction. Otherwise, v_{ij} 's rank c_j within bottom ε' fraction.
- Independently for each voter, fill the rest of her preferences with the $m - z$ dummy candidates uniformly randomly.
- The copies of a voter only differ in the ranking of the dummy candidates. We set the number of copies to be $R = \left\lceil \frac{10mk^2}{n\varepsilon^2} \right\rceil$. These copies are there to ensure Lemma 4.4 applies to the dummy candidates.

Clearly, the above construction has size $\text{poly}(1/\varepsilon, n, z, k)$. Let OPT denote the optimal score on this instance. Recall that $\text{RAND} = \frac{m+1}{k+1}$. First suppose the instance of REGULAR MAX K-COVER has value n (“YES” instance) so that there are k sets that cover all n elements, then it is easy to check that choosing the corresponding critical candidates as the committee yields $\text{OPT} \leq z < \varepsilon' \cdot \text{RAND}$.

On the other hand, suppose the instance of REGULAR MAX K-COVER is such that any collection of sets of size k only covers at most $(1 - 1/e + \varepsilon)n \leq \frac{3}{4}n$ elements (“NO” instance). Consider any committee T and suppose $T = R \cup D$ where R is a subset of critical candidates and D is a subset of dummy candidates. Let $r = |R|$ and $d = |D| = k - r$. Let n' be the number of elements R covers in the REGULAR MAX K-COVER instance. By assumption, $n - n' \geq \frac{n}{4}$ since any collection R covers at most $\frac{3}{4}n$ elements. Further, since the instance is regular, $n - n' \leq \frac{n}{k}r$, so that $n - n' \geq d\frac{n}{k}$.

Using Lemma 4.4, with probability $> \frac{1}{2}$ over the choice of the ranking of the dummy candidates, the optimal score of D on the $(n - n')R$ uncovered voters using the $m - z$ dummy candidates is greater than $(1 - \varepsilon)\frac{m-z}{d+1}$. Inserting the critical candidates cannot decrease this score for these voters, since the candidates in R appear last in their ordering. Further, we have assumed $m = \frac{2}{\varepsilon}kz$. Therefore, with probability $> \frac{1}{2}$, we have:

$$\text{OPT} > \frac{n - n'}{n}(1 - \varepsilon)\frac{m - z}{d + 1} \geq \frac{n - n'}{n} \cdot \left(1 - \frac{\varepsilon'}{4}\right) \cdot \frac{m + 1}{d + 1}$$

We now split the analysis into two cases:

- (1) Suppose $d + 1 \leq \frac{k+1}{4}$. Since $n - n' \geq \frac{n}{4}$, we have

$$\text{OPT} > \frac{1}{4} \cdot \left(1 - \frac{\varepsilon'}{4}\right) \cdot \frac{m + 1}{d + 1} \geq \left(1 - \frac{\varepsilon'}{4}\right) \cdot \frac{m + 1}{k + 1} \geq (1 - \varepsilon')\text{RAND}.$$

(2) Suppose $d + 1 \geq \frac{k+1}{4}$. Since $n - n' \geq \frac{d}{k}n$, and since $d, k = \omega(1)$, we have:

$$\text{OPT} > \frac{d}{k} \cdot \left(1 - \frac{\varepsilon'}{4}\right) \cdot \frac{m+1}{d+1} = \frac{d}{d+1} \cdot \frac{k+1}{k} \cdot \left(1 - \frac{\varepsilon'}{4}\right) \cdot \frac{m+1}{k+1} \geq (1 - \varepsilon')\text{RAND}.$$

Therefore, our construction ensures that with probability $> \frac{1}{2}$, we have $\text{OPT} \geq (1 - \varepsilon')\text{RAND}$ if the original REGULAR MAX K-COVER instance has value at most $\frac{3}{4}n$.

Now suppose there is a polynomial time algorithm that can distinguish between instances with $\text{OPT} \leq \varepsilon'\text{RAND}$ and $\text{OPT} \geq (1 - \varepsilon')\text{RAND}$. Then, feeding the output of the above construction to this algorithm implies a coRP algorithm for the decision version of REGULAR MAX K-COVER, which by Theorem 4.3 implies $\text{NP} \subseteq \text{coRP}$. Since $\text{RP} \subseteq \text{NP}$, this implies $\text{RP} \subseteq \text{coRP}$, so that $\text{ZPP} = \text{RP} \cap \text{coRP} = \text{RP}$. Since ZPP is symmetric with respect to “YES” and “NO” instances, this implies $\text{ZPP} = \text{coRP}$, so that $\text{ZPP} = \text{NP}$.

We now show how to set ε . For the first part of the theorem, we set $\varepsilon = \frac{1}{10} \left(\frac{k}{m}\right)^{1-\delta}$. This can be achieved by choosing m such that $\left(\frac{m}{k}\right)^\delta = 20kz$. Note that this ensures $m = \text{poly}(k, z)$ when δ is a constant, so that the construction runs in polynomial time. For this setting, we have $\varepsilon' \cdot \frac{m+1}{k+1} \leq \left(\frac{m+1}{k+1}\right)^\delta$, while $\varepsilon' = \left(\frac{k}{m}\right)^{1-\delta} = \left(\frac{1}{20kz}\right)^{\frac{1-\delta}{\delta}} = o(1)$, completing the proof.

For the second part of the theorem, we set $m = 20k^{1+\alpha}z$, and $\varepsilon = \frac{1}{10} \frac{1}{k^\alpha}$. Again, we have $m = \text{poly}(k, z)$, and $\varepsilon' = o(1)$, completing the proof. $\square$

Theorem 1.2 now implies the following easy corollaries.

COROLLARY 4.5. *Unless $\text{NP} = \text{ZPP}$, there is no k^α -approximation to the 1-Borda score for any constant $\alpha > 0$. Similarly, there is no $\left(\frac{m+1}{k+1}\right)^{1-\delta}$ -approximation for any constant $\delta \in (0, 1)$.*

The next corollary adapts the hardness proof to the maximization version of the problem.

COROLLARY 4.6. *For the maximization version of 1-Borda, there is no polynomial time $\left(1 - \frac{1-\varepsilon}{k+1}\right)$ -approximation for constant $\varepsilon \in (0, 1/2)$ unless $\text{NP} = \text{ZPP}$.*

PROOF. Set $\varepsilon > 0$ to be a small constant in the proof of Theorem 1.2. Then, in the “NO” instance, the maximization score is at most $(m+1) \left(1 - \frac{1-\varepsilon}{k+1}\right)$, while for the “YES” instance, the score is at least $(m+1) \left(1 - \frac{\varepsilon}{k+1}\right)$. For $\varepsilon \in (0, 1/2)$, the approximation factor achievable is therefore at most $\left(1 - \frac{1-2\varepsilon}{k+1}\right)$, completing the proof. $\square$

4.2 An Optimal Deterministic Algorithm

Given the lower bound and the hardness result, an immediate question is whether there is a deterministic rule to achieve the benchmark RAND. We answer in the affirmative: The BANZHAF algorithm [3, 11, 13] can be viewed as a derandomization of RAND: Instead of randomly picking a candidate at each iteration, it picks the candidate that gives the best expected performance if the rest of the committee is randomly constructed. It is shown in [13] that this algorithm runs in polynomial time. The following theorem implies Theorem 1.3.

THEOREM 4.7. $\text{BANZHAF} \leq \text{RAND}$.

PROOF. Recall that BANZHAF builds sets $\emptyset = T_0 \subsetneq T_1 \subsetneq \dots \subsetneq T_k$, where at step j , BANZHAF picks $c_j \in C \setminus T_{j-1}$ such that:

$$c_j = \underset{c \in C \setminus T_{j-1}}{\text{argmin}} \sum_{\substack{S \subseteq C: |S|=k \\ S \supseteq T_{j-1} \cup \{c\}}} r_V(S). \quad (5)$$

We now use induction to show that for any $j \in \{0, 1, \dots, k\}$,

$$\frac{1}{\binom{m-j}{k-j}} \sum_{\substack{S \subseteq C: |S|=k \\ S \supseteq T_j}} r_V(S) \leq \text{RAND},$$

which is clearly true when $j = 0$, and gives the desired result $\text{BANZHAF} \leq \text{RAND}$ when $j = k$.

For the inductive step, assume it holds for some $j - 1$. Now in the j^{th} iteration, we have the following inequalities that complete the proof. Here, the first step follows since BANZHAF picks c_j in step j . The second step follows since BANZHAF solves Eq (5), so that the score from adding c_j beats the average score of adding one of the $m - j + 1$ candidates in $C \setminus T_{j-1}$. The final equality follows since $\binom{m-j+1}{k-j+1} = \frac{m-j+1}{k-j+1} \binom{m-j}{k-j}$, and by observing that for any $S \supseteq T_{j-1}$, there are $k - j + 1$ choices of $c \in S \setminus T_{j-1}$.

$$\begin{aligned} \frac{1}{\binom{m-j}{k-j}} \sum_{\substack{S \subseteq C: |S|=k \\ S \supseteq T_j}} r_V(S) &= \frac{1}{\binom{m-j}{k-j}} \sum_{\substack{S \subseteq C: |S|=k \\ S \supseteq T_{j-1} \cup \{c_j\}}} r_V(S) \\ &\leq \frac{1}{\binom{m-j}{k-j}} \cdot \left(\frac{1}{m-j+1} \sum_{c \in C \setminus T_{j-1}} \sum_{\substack{S \subseteq C: |S|=k \\ S \supseteq T_{j-1} \cup \{c\}}} r_V(S) \right) \\ &= \frac{1}{\binom{m-j}{k-j}} \cdot \frac{k-j+1}{m-j+1} \cdot \left(\frac{1}{k-j+1} \sum_{c \in C \setminus T_{j-1}} \sum_{\substack{S \subseteq C: |S|=k \\ S \supseteq T_{j-1} \cup \{c\}}} r_V(S) \right) \\ &= \frac{1}{\binom{m-j+1}{k-j+1}} \sum_{\substack{S \subseteq C: |S|=k \\ S \supseteq T_{j-1}}} r_V(S) \leq \text{RAND}. \quad \square \end{aligned}$$

To complete the proof of Theorem 1.3, for the maximization objective, BANZHAF achieves a value at least $(m+1) \left(1 - \frac{1}{k+1}\right)$. Since the maximum possible value is m , this implies a $\left(1 - \frac{1}{k+1}\right)$ -approximation.

5 LOWER BOUND ON COMMITTEE-MONOTONE ALGORITHMS FOR 1-BORDA

Consider the 1-Borda score. A nice property of GREEDY is that it is committee-monotone: In each iteration, the candidate chosen by GREEDY only depends on which candidates have been chosen in previous iterations and not on k , and thus when k increases, the committee selected by GREEDY includes all the candidates GREEDY used to select. On the other hand, BANZHAF does not satisfy committee monotonicity, as the candidates chosen by BANZHAF does depend on k .

This naturally brings up the question: Is there a committee-monotone algorithm which is optimal with respect to the benchmark RAND? We answer this question in the negative, by presenting a lower bound of 1.015 for all committee-monotone algorithms.

THEOREM 5.1. *For any large enough m , there exist instances with m candidates where any committee-monotone algorithm ALG satisfies $r_V(T_k) > 1.015 \cdot \text{RAND}$ for some value $k \in \{1, 2\}$. Here, T_k is the set of candidates ALG chooses when the size of this set is k .*

PROOF. The construction goes as follows: There are two types of candidates, X and Y . Candidates of type X are ranked between $[am, bm]$ by every voter and candidates of type Y are ranked between $[1, am] \cup [bm, m]$ by every voter, where $0 < a < b < 1$ are two parameters. We construct sufficiently

many voters so that all candidates of the same type are symmetric (by having all permutations of candidates of type X and those of type Y). We want to find proper a and b , so that when $k = 1$, the optimal choice is to choose a candidate of type X , while when $k = 2$, the optimal choice is to choose two candidates both of type Y . This means that no committee-monotone algorithm can produce optimal choice for both the first iteration and the second iteration. We optimize over a and b to find the maximum lower bound.

In particular, the search procedure goes as follows. Let $r_V(Y)$ denote the 1-Borda score of choosing a candidate of type Y ; $r_V(XX)$ denote the score of choosing two candidates both of type X ; and $r_V(XY)$ denote the score of choosing a candidate of type X and a candidate of type Y . We can easily see that, when m goes to infinity, up to an $o(1)$ additive error:

$$\begin{cases} \frac{1}{m+1}r_V(Y) = \frac{a}{2} \cdot \Pr[Y \text{ is at } [1, am]] + \frac{1+b}{2} \cdot \Pr[Y \text{ is at } [bm, m]] = \frac{a}{2} \cdot \frac{a}{1-(b-a)} + \frac{1+b}{2} \cdot \frac{1-b}{1-(b-a)} \\ \frac{1}{m+1}r_V(XX) = \frac{2a+b}{3} \\ \frac{1}{m+1}r_V(XY) = \frac{a}{2} \cdot \Pr[Y \text{ is at } [1, am]] + \frac{a+b}{2} \cdot \Pr[Y \text{ is at } [bm, m]] = \frac{a}{2} \cdot \frac{a}{1-(b-a)} + \frac{a+b}{2} \cdot \frac{1-b}{1-(b-a)} \end{cases}.$$

A committee-monotone algorithm either chooses Y in the first iteration, or chooses XX or XY in the first two iterations. Thus, we maximize $\min(\frac{2}{m+1}r_V(Y), \frac{3}{m+1}r_V(XX), \frac{3}{m+1}r_V(XY))$ (note that the value on the numerator corresponds to the value of $k+1$) over $0 < a < b < 1$, and find that, for $a = 0.377$ and $b = 0.552$, it achieves a lower bound greater than 1.015. $\square$

6 THE s -BORDA SCORE

In this section, we extend our analysis of the greedy algorithms to s -Borda score, and show how to significantly improve on the GREEDY and BANZHAF rules via linear programming.

Recall that $\text{RAND} = \frac{m+1}{k+1}$ and choosing a random committee of size k yields expected score $\frac{s(s+1)}{2} \cdot \text{RAND}$. As a derandomization, BANZHAF has score at most this value similar to Theorem 4.7. Let OPT denote the best possible s -Borda score. Considering the instance with one voter for each permutation of candidates as its preference ordering, we have the following proposition:

PROPOSITION 6.1. *For any s , m and k , there exists instances where $\text{OPT} = \frac{s(s+1)}{2} \cdot \text{RAND}$.*

We first consider a natural extension of GREEDY in the 1-Borda case. In the full version, we show that it achieves an s -Borda score at most $2s^2 \cdot \text{RAND}$ (Theorem 6.2), which is within a factor of $\frac{4s}{s+1}$ of the BANZHAF rule. We then show that this bound cannot be improved even when OPT is small. However, unlike the 1-Borda case, there is no fundamental barrier to an improved algorithm when OPT is small, and we present such an algorithm in Section 6.2.

6.1 The GREEDY Algorithm

The GREEDY algorithm follows exactly the same procedure as for 1-Borda, except that we now compute the score based on s -Borda. We present an upper bound for GREEDY in the following theorem. Since the proof is very similar to the $s = 1$ case, we present it in the full version.

THEOREM 6.2 (PROVED IN THE FULL VERSION). $\text{GREEDY} \leq 2s^2 \cdot \text{RAND}$.

Lower Bound for Small OPT. In general, $\text{OPT} = \Omega(s^2) \cdot \text{RAND}$, in which case the analysis of greedy is tight to within a constant factor. The question we now ask is: Does GREEDY always perform better when OPT is small? We answer this in the negative; the proof is in the full version:

THEOREM 6.3 (PROVED IN THE FULL VERSION). *There exists an instance where $\text{OPT} = O(s^2) = o(1) \cdot \text{RAND}$, while the score of GREEDY is $\Omega(s^2) \cdot \text{RAND}$.*

This shows that GREEDY can perform as bad as random even when OPT is small and thus motivates the improved guarantee in Section 6.2.

6.2 An Improved Algorithm via LP Rounding

As mentioned above, GREEDY can hit its worst-case bound of $\Omega(s^2) \cdot \text{RAND}$ even when OPT is actually small. We know that for the case of 1-Borda, no polynomial-time algorithm can do better. Now the question is, can a different algorithm do better in the case of s -Borda for $s = \omega(1)$? We answer this question in the affirmative by presenting an algorithm that is based on dependent rounding of an LP relaxation combined with uniform random sampling, which provides nontrivial improvement when OPT is small. In particular, it achieves expected score at most $3 \cdot \text{OPT} + O(s^{3/2} \log s) \cdot \text{RAND}$.

6.2.1 LP-Rounding-Based Algorithm. The following linear program (based on [6, 8, 10, 17, 20]) is a natural relaxation for the s -Borda problem.

$$\begin{aligned}
 & \text{minimize} && \sum_{i=1}^n \sum_{\ell=1}^s \sum_{j=1}^m x_{ij}^{\ell} \cdot r_{v_i}(c_j), \\
 & \text{subject to} && \sum_{j=1}^m y_j = k, \\
 & && \sum_{\ell=1}^k x_{ij}^{\ell} \leq y_j, \quad i \in \{1, \dots, n\}, \quad j \in \{1, \dots, m\}, \\
 & && \sum_{j=1}^m x_{ij}^{\ell} \geq 1, \quad i \in \{1, \dots, n\}, \quad \ell \in \{1, \dots, k\}, \\
 & && y_j, x_{ij}^{\ell} \in [0, 1], \quad i \in \{1, \dots, n\}, \quad j \in \{1, \dots, m\}, \quad \ell \in \{1, \dots, k\}.
 \end{aligned}$$

Variable y_j denotes how much candidate c_j is chosen; integral values 1 and 0 mean choosing and not choosing candidate c_j , respectively. The first constraint encodes choosing exactly k candidates. We copy each voter k times, and the ℓ^{th} copy of the voter v_i is assigned to the ℓ^{th} -preferred chosen candidate. Variable x_{ij}^{ℓ} denotes how much the ℓ^{th} copy of voter v_i is assigned to candidate c_j . The second constraint prevents a voter from being assigned to a candidate that is not chosen. The third constraint ensures that each copy of the voter is assigned to some candidate. The objective function computes the s -Borda score.

We will use dependent rounding [15] to round this LP solution. There is a catch though: Dependent rounding can cause a deficit in around $\tilde{O}(\sqrt{s})$ candidates from the top s that are fractionally chosen by the LP. Since any solution must account for the top s scores, we need to ensure these “deficit” candidates do not increase the score too much. Towards this end, we scale down the LP solution, and choose enough candidates uniformly at random so that these candidates can absorb the deficit. However, such scaling creates a further deficit that will have to be absorbed by random sampling. We find that the right trade-off is achieved by scaling down by a factor of $(1 - \frac{1}{\sqrt{s}})$.

Without further ado, the overall algorithm works as follows:

- (1) Solve the above linear program and let $\tilde{y}$ denote the optimal solution.
- (2) For $j = 1, 2, \dots, m$, let $y_j = (1 - \frac{1}{\sqrt{s}})\tilde{y}_j$. Note that $\sum_{j=1}^m y_j = k(1 - \frac{1}{\sqrt{s}})$.
- (3) Apply dependent rounding [15] on the variables $\{y_j\}$ so that exactly $k(1 - \frac{1}{\sqrt{s}})$ candidates are chosen. Let T_1 denote the set of chosen candidates.
- (4) Finally choose a set T_2 of $\frac{k}{\sqrt{s}}$ candidates uniformly at random from $C \setminus S$ and output $T = T_1 \cup T_2$.

We will show the following theorem:

THEOREM 6.4. *When $m = \omega(k)$, $k = \omega(s^{3/2})$, and $s = \omega(1)$, we have:*

$$\mathbb{E}[r_{\mathcal{V}}(T)] \leq 3\text{OPT} + O(s^{3/2} \log s) \cdot \text{RAND}.$$

6.2.2 Analysis: Proof of Theorem 6.4. First consider dependent rounding on $\{y_j\}$. Let Y_j denote the random variable which returns 1 if y_j is rounded to 1 and 0 if y_j is rounded to 0. Note that $\mathbb{E}[Y_j] = y_j$ for all candidates $j \in C$. The following lemma is an easy consequence of Chernoff bounds:

LEMMA 6.5. *For any subset of candidates $\{c_{j_1}, \dots, c_{j_\ell}\}$, let $W = \sum_{t=1}^\ell y_{j_t}$. If $W = \omega(1)$, then*

$$\Pr \left[\sum_{t=1}^\ell Y_{j_t} \in \left(W - 9\sqrt{W \log W}, W + 9\sqrt{W \log W} \right) \right] \geq 1 - \frac{2}{W^3}.$$

We now compute the expected s -Borda score for each voter. Towards this end, we partition the candidates into buckets with geometrically decreasing sum of y_i values, and account for the expected score generated by dependent rounding in each bucket against the LP value of the subsequent bucket. Lemma 6.5 will ensure the number of candidates chosen from each bucket is close to the LP value, and the deficit gets taken care of by the uniformly randomly chosen candidates.

For simplicity of notation, let $\eta = \log_2 \frac{\sqrt{s}}{2}$. Fix a voter v_i , and suppose its preference order is $c_{i_1} > c_{i_2} > \dots > c_{i_m}$. Recall that $\{\tilde{y}, x\}$ is the LP solution. The values x_{ij} in the LP are set as follows: Consider the prefix of the ordering such that $\sum_{t=1}^\ell \tilde{y}_{i_t} \leq s$ and $\sum_{t=1}^{\ell+1} \tilde{y}_{i_t} > s$. The LP sets $x_{ii_t} = \tilde{y}_{i_t}$ for $t \leq \ell$, and sets $x_{ii_{\ell+1}} = s - \sum_{t=1}^\ell \tilde{y}_{i_t}$. The contribution of v_i to the LP objective is therefore

$$\text{OPT}_i = \sum_{t=1}^\ell r_{v_i}(c_{i_t}) \tilde{y}_{i_t} + r_{v_i}(c_{i_{\ell+1}}) \left(s - \sum_{t=1}^\ell \tilde{y}_{i_t} \right) \geq \sum_{t=1}^\ell r_{v_i}(c_{i_t}) \tilde{y}_{i_t}. \quad (6)$$

Consider the first ℓ candidates in the above ordering. We have $\sum_{j=1}^\ell \tilde{y}_j \geq s - 1$, so that

$$\sum_{j=1}^\ell y_j \geq (s - 1) \left(1 - \frac{1}{\sqrt{s}} \right) \geq s - (\sqrt{s} + 1) \geq s - 2\sqrt{s}. \quad (7)$$

We split these ℓ candidates into sets $C_1, \dots, C_\eta$ as follows: We walk down the preference order of v_i . We take C_1 as the set of candidates whose y -values sum to $\frac{s}{2}$; C_2 as the next set of candidates whose y -values sum to $\frac{s}{4}$, and so on until C_η , whose sum of y -values is $\frac{s}{2^\eta} = 2\sqrt{s}$. Now the sum of y -values of all candidates in $\{C_1, \dots, C_\eta\}$ is exactly $s - 2\sqrt{s}$. Formally, we define

$$\theta_0 = 0, \quad \theta_h = \min \left\{ q \mid \sum_{t=1}^q y_{i_t} \geq \left(1 - \frac{1}{2^h} \right) s \right\}, \quad \forall h \in \{1, \dots, \eta\},$$

and correspondingly define the sets $\{C_1, \dots, C_\eta\}$ as:

$$C_h = \{c_{i_q} \mid \theta_{h-1} < q \leq \theta_h\}, \quad \forall h \in \{1, \dots, \eta\}.$$

For all $h \in \{1, \dots, \eta\}$, let $y_{C_h} = \sum_{c_j \in C_h} y_j$ and $Y_{C_h} = \sum_{c_j \in C_h} Y_j$. Note that y_{C_h} decreases by a factor of 2 as h increases. Now consider the outcome of the dependent rounding procedure for each of the sets $C_1, \dots, C_\eta$. We say the rounding fails for v_i if there exists $h \in \{1, \dots, \eta - 1\}$ such that the number Y_{C_h} of chosen candidates in C_h is not in range $y_{C_h} \pm 9\sqrt{y_{C_h} \log y_{C_h}}$. We will not consider C_η when defining failure, and will deal with this set separately.

Let FAIL denote the failure event. We now bound the probability of the event FAIL for voter v_i .

LEMMA 6.6. $\Pr[\text{FAIL}] \leq s^{-3/2}$.

PROOF. By union bound applied to Lemma 6.5, we have:

$$\Pr[\text{FAIL}] \leq \sum_{h=1}^{\eta-1} \frac{2}{y_{C_h}^3} \leq \frac{3}{y_{C_{\eta-1}}^3} \leq s^{-3/2},$$

where we have used that $\{y_{C_h}\}$ is a geometrically decreasing sequence, and that $y_{C_{\eta-1}} = 4\sqrt{s}$. $\square$

We are now ready to compute the expected score for v_i in our algorithm. Recall that T denotes the set of chosen candidates and OPT_i denotes the s -Borda score for v_i in the LP solution. Let BAD denote the expected s -Borda score for v_i in the event FAIL , and GOOD denote the expected score otherwise. We will bound these separately below.

LEMMA 6.7. $\text{BAD} \leq O(s^{5/2}) \cdot \text{RAND}$.

PROOF. If FAIL happens, the final solution is still at least as good as choosing the $\frac{k}{\sqrt{s}}$ random candidates in Step (4) of the algorithm. Note that since we assumed $k = \omega(s^{3/2})$, we have $\frac{k}{\sqrt{s}} \geq s$, so that we will have chosen enough random candidates to fill up at least s positions for computing s -Borda score. Further, since we assume that $m = \omega(k)$, the score of the solution will at most double had we assumed these $\frac{k}{\sqrt{s}}$ candidates are chosen randomly from the entire set of m candidates instead of from the remaining $m - k(1 - \frac{1}{\sqrt{s}})$ candidates after dependent rounding. Thus, we have:

$$\text{BAD} \leq 2 \mathbb{E}_{T \subseteq C, |T|=k/\sqrt{s}}[r_{v_i}(T)] \leq 2 \frac{s(s+1)}{2} \frac{m+1}{\frac{k}{\sqrt{s}}+1} \leq 4s^{3/2}(s+1) \cdot \text{RAND},$$

which yields that $\text{BAD} \leq O(s^{5/2}) \cdot \text{RAND}$. $\square$

LEMMA 6.8. $\text{GOOD} \leq 3\text{OPT}_i + O(s^{3/2} \log s) \cdot \text{RAND}$.

PROOF. Suppose FAIL does not happen. Denote the set of candidates chosen by the algorithm from $\{C_1, \dots, C_{\eta-1}\}$ as T_1 , and the randomly chosen $\frac{k}{\sqrt{s}}$ candidates as T_2 . Therefore $T = T_1 \cup T_2$. From Eq (7), we have

$$\sum_{h=1}^{\eta-1} y_{C_h} = s - 2\sqrt{s} - y_{C_\eta} = s - 4\sqrt{s}.$$

Since FAIL does not happen, we have:

$$|T_1| \geq s - 4\sqrt{s} - \sum_{h=1}^{\eta-1} \sqrt{y_{C_h} \log y_{C_h}} \geq s - 4\sqrt{s} - \sum_{h=1}^{\eta-1} \sqrt{\frac{s}{2^h} \log s} \geq s - 4\sqrt{s} - (\sqrt{2} + 1)\sqrt{s \log s}.$$

Denote $u = 4\sqrt{s} + (\sqrt{2} + 1)\sqrt{s \log s} = O(\sqrt{s \log s})$, so that $|T_1| \geq s - u$. The quantity u is the total “deficit” in candidates from the top s that is caused by scaling the LP and dependent rounding. We make up this deficit using the set T_2 . Specifically, consider the subsets,

$$T_1^* = \arg \min_{Q \subseteq T_1, |Q|=s-u} \sum_{c \in Q} r_{v_i}(c) \quad \text{and} \quad T_2^* = \arg \min_{Q \subseteq T_2, |Q|=u} \sum_{c \in Q} r_{v_i}(c).$$

Note that $T_1^* \subseteq T_1$, and $T_2^* \subseteq T_2$. We will evaluate the score of these subsets of candidates, which will be an upper bound on the score of the algorithm. Towards this end, we define μ_h as the scaled LP score of C_h , that is:

$$\mu_h = \frac{\sum_{j=\theta_{h-1}+1}^{\theta_h} r_{v_i}(c_{i_j}) y_{i_j}}{y_{C_h}}, \quad \forall h \in \{1, \dots, \eta\}.$$

Since $y_i \leq \tilde{y}_i$, combining the previous inequality with Eq (6), we have:

$$\sum_{h=1}^{\eta} \mu_h y_{C_h} \leq \text{OPT}_i.$$

Since $y_{C_h} \geq 2\sqrt{s} = \omega(1)$ for all $h \in \{1, 2, \dots, \eta\}$, we have:

$$|T_1 \cap C_h| \leq y_{C_h} + 9\sqrt{y_{C_h} \log y_{C_h}} \leq \frac{3}{2}y_{C_h}.$$

Since $\mu_h > r_{v_i}(c)$, $\forall c \in C_{h-1}$ and since $y_{C_h} \leq 2y_{C_{h+1}}$, we can bound the expected score of T_1^* as:

$$\sum_{c \in T_1^*} r_{v_i}(c) \leq \sum_{h=1}^{\eta-1} \frac{3}{2} \cdot y_{C_h} \mu_{h+1} \leq \sum_{h=1}^{\eta-1} 3 \cdot y_{C_{h+1}} \mu_{h+1} \leq 3\text{OPT}_i.$$

We can again assume that the $\frac{k}{\sqrt{s}}$ random candidates are chosen randomly from the entire set of m candidates. This yields a bound on the score of T_2^* as:

$$\sum_{c \in T_2^*} r_{v_i}(c) \leq \sum_{t=1}^u t \cdot (\sqrt{s} \cdot \text{RAND}) = O(s^{3/2} \log s) \cdot \text{RAND}.$$

where we used $u = O(\sqrt{s \log s})$ to derive $\sum_{t=1}^u t = O(s \log s)$.

Therefore, we can bound GOOD as follows:

$$\text{GOOD} \leq \sum_{c \in T_1^*} r_{v_i}(c) + \sum_{c \in T_2^*} r_{v_i}(c) \leq 3\text{OPT}_i + O(s^{3/2} \log s) \cdot \text{RAND}. \quad \square$$

Synthesizing the bounds from Lemmas 6.7 and 6.8, we can conclude that:

$$\begin{aligned} \mathbb{E}[r_{v_i}(T)] &= \text{BAD} \cdot \Pr[\text{FAIL}] + \text{GOOD} \cdot (1 - \Pr[\text{FAIL}]) \\ &\leq s^{-3/2} \cdot O(s^{5/2}) \cdot \text{RAND} + 3\text{OPT}_i + O(s^{3/2} \log s) \cdot \text{RAND} \\ &= 3\text{OPT}_i + O(s^{3/2} \log s) \cdot \text{RAND}. \end{aligned}$$

Taking expectation over all voters, this yields Theorem 6.4.

7 CONCLUSION

Our work opens some interesting directions for further research. One open question is to extend our results to the stronger notion of approximate core stability under the CC rule for which the best known result is a 16-approximation [9, 19]. It would be interesting to explore if our techniques can help improve the approximation factor via a simple-to-implement procedure.

We conjecture that there is a lower bound of $\Omega(s^{3/2}) \cdot \text{RAND}$ on the score achievable by poly-time algorithms for s -Borda, i.e., that the algorithm in Section 6.2 is almost optimal. This will require a non-trivial strengthening of known hardness results for maximum multicover [4]. It would also be interesting to explore if there are greedy rules that can match these bounds.

In the same vein, another interesting question is to map the landscape of approximation ratios for generalizations such as committee scoring rules. The work of [6] shows strong positive results when voters assign a smooth set of weights to all candidates in the committee, while our work considers the case where the weights are concentrated on higher-ranked candidates. There is a large middle ground where the approximability of this problem is poorly understood.

ACKNOWLEDGMENTS

We thank Brandon Fain for several discussions, and the anonymous reviewers for their suggestions. This work is supported by NSF grant CCF-1637397, ONR award N00014-19-1-2268, and DARPA award FA8650-18-C-7880.

REFERENCES

- [1] Vijay Arya, Naveen Garg, Rohit Khandekar, Adam Meyerson, Kamesh Munagala, and Vinayaka Pandit. 2004. Local Search Heuristics for k-Median and Facility Location Problems. *SIAM J. Comput.* 33, 3 (2004), 544–562.
- [2] Haris Aziz and Barton E. Lee. 2020. The expanding approvals rule: improving proportional representation and monotonicity. *Social Choice and Welfare* 54, 1 (2020), 1–45.
- [3] J. Banzhaf. 1965. Weighted Voting Doesn't Work: A Mathematical Analysis. *Rutgers Law Review* 19, 2 (1965), 317–343.
- [4] Siddharth Barman, Omar Fawzi, Suprovat Ghoshal, and Emirhan Gürpınar. 2020. Tight approximation bounds for maximum multi-coverage. In *International Conference on Integer Programming and Combinatorial Optimization, IPCO*. 66–77.
- [5] Markus Brill, Piotr Faliszewski, Frank Sommer, and Nimrod Talmon. 2019. Approximation Algorithms for BalancedCC Multiwinner Rules. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '19, Montreal, QC, Canada, May 13-17, 2019*. International Foundation for Autonomous Agents and Multiagent Systems, 494–502.
- [6] Jarosław Byrka, Piotr Skowron, and Krzysztof Sornat. 2018. Proportional Approval Voting, Harmonic k-median, and Negative Association. In *45th International Colloquium on Automata, Languages, and Programming, ICALP (LIPIcs, Vol. 107)*. 26:1–26:14.
- [7] John R. Chamberlin and Paul N. Courant. 1983. Representative Deliberations and Representative Decisions: Proportional Representation and the Borda Rule. *The American Political Science Review* 77, 3 (1983), 718–733.
- [8] Moses Charikar, Sudipto Guha, Éva Tardos, and David B. Shmoys. 2002. A Constant-Factor Approximation Algorithm for the k-Median Problem. *J. Comput. System Sci.* 65, 1 (2002), 129–149.
- [9] Yu Cheng, Zhihao Jiang, Kamesh Munagala, and Kangning Wang. 2020. Group Fairness in Committee Selection. *ACM Trans. Economics and Comput.* 8, 4 (2020), 23:1–23:18.
- [10] Gérard Cornuéjols, George Nemhauser, and Laurence Wolsey. 1983. *The uncapacitated facility location problem*. Technical Report. Cornell University Operations Research and Industrial Engineering.
- [11] Pradeep Dubey and Lloyd S. Shapley. 1979. Mathematical Properties of the Banzhaf Power Index. *Mathematics of Operations Research* 4, 2 (1979), 99–131.
- [12] Edith Elkind, Piotr Faliszewski, Piotr Skowron, and Arkadii Slinko. 2017. Properties of multiwinner voting rules. *Social Choice and Welfare* 48, 3 (2017), 599–632.
- [13] Piotr Faliszewski, Martin Lackner, Dominik Peters, and Nimrod Talmon. 2018. Effective Heuristics for Committee Scoring Rules. *Proceedings of the AAAI Conference on Artificial Intelligence, AAAI* 32, 1 (Apr. 2018).
- [14] Uriel Feige. 1998. A Threshold of $\ln n$ for Approximating Set Cover. *J. ACM* 45, 4 (July 1998), 634–652.
- [15] R. Gandhi, S. Khuller, S. Parthasarathy, and A. Srinivasan. 2002. Dependent rounding in bipartite graphs. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science, FOCS*. 323–332.
- [16] Ashish Goel, Anilesh K. Krishnaswamy, Sukolsak Sakshuwong, and Tanja Aitamurto. 2019. Knapsack Voting for Participatory Budgeting. *ACM Trans. Econ. Comput.* 7, 2, Article 8 (Jul 2019).
- [17] Mohammadtaghi Hajiaghayi, Wei Hu, Jian Li, Shi Li, and Barna Saha. 2016. A Constant Factor Approximation Algorithm for Fault-Tolerant k-Median. *ACM Trans. Algorithms* 12, 3, Article 36 (April 2016), 19 pages.
- [18] Kamal Jain and Vijay V. Vazirani. 2001. Approximation Algorithms for Metric Facility Location and k-Median Problems Using the Primal-Dual Schema and Lagrangian Relaxation. *J. ACM* 48, 2 (2001), 274–296.
- [19] Zhihao Jiang, Kamesh Munagala, and Kangning Wang. 2020. Approximately stable committee selection. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC*. 463–472.
- [20] Jyh-Han Lin and Jeffrey Scott Vitter. 1992. ϵ -approximations with Minimum Packing Constraint Violation (Extended Abstract). In *Proceedings of the 24th Annual ACM Symposium on Theory of Computing, STOC*. 771–782.
- [21] Tyler Lu and Craig Boutilier. 2011. Budgeted Social Choice: From Consensus to Personalized Decision Making. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence, IJCAI*. 280–286.
- [22] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. 1978. An analysis of approximations for maximizing submodular set functions—I. *Mathematical Programming* 14, 1 (1978), 265–294.
- [23] PBP. 2016. Where has it worked? - The Participatory Budgeting Project. <http://www.participatorybudgeting.org/about-participatory-budgeting/where-has-it-worked/>.
- [24] Piotr Skowron, Piotr Faliszewski, and Jerome Lang. 2016. Finding a collective set of items: From proportional multirepresentation to group recommendation. *Artificial Intelligence* 241 (2016), 191 – 216.
- [25] Piotr Skowron, Piotr Faliszewski, and Arkadii Slinko. 2015. Achieving fully proportional representation: Approximability results. *Artificial Intelligence* 222 (2015), 67 – 103.
- [26] T. N. Thiele. 1895. Om flerfoldevalg. *Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandlinger* (1895), 415–441.
- [27] Nicolaus Tideman and Daniel Richardson. 2000. Better Voting Methods Through Technology: The Refinement-Manageability Trade-Off in the Single Transferable Vote. *Public Choice* 103, 1 (2000), 13–34.