

Promoting Fairness and Priority in k -Winners Selection Using IRV

Anonymous Author(s)

ABSTRACT

We investigate the problem of finding winner(s) given a large number of users (voters) preferences casted as *ballots*, one from each of the m users, where each ballot is a ranked order of preference of up to ℓ out of n items (candidates). Given a group protected attribute with k different values and a priority that imposes a selection order among these groups, the goal is to satisfy the priority order and select a winner per group that is most representative. It is imperative that at times the original users' preferences may require further manipulation to meet these fairness and priority requirement. We consider manipulation by modifications and formalize the *margin finding problem under modification* problem. We study the suitability of Instant Run-off Voting (IRV) as a preference aggregation method and demonstrate its advantages over positional methods. We present a suite of technical results on the hardness of the problem, design algorithms with theoretical guarantees and further investigate efficiency opportunities. We present exhaustive experimental evaluations using multiple applications and large-scale datasets to demonstrate the effectiveness of IRV, and efficacy of our designed solutions qualitatively and scalability-wise.

ACM Reference Format:

Anonymous Author(s). 2024. Promoting Fairness and Priority in k -Winners Selection Using IRV. In *Proceedings of 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'24)*. ACM, New York, NY, USA, 15 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 INTRODUCTION

The task of finding the winner, i.e., the most favorable item or candidate from a given set of m users' (voters') preferences over n items (candidates), has found a wide variety of applications such as in hiring candidate(s) for a job, selecting member(s) of a committee, finding winning candidate(s) in a competition, in electoral voting, or even in recommender systems. IRV (Instant Run-off Voting) is a ranked choice voting mechanism that has been gaining popularity lately as an electoral system in the US [11, 13, 22, 26, 34]. In this paper, we study the applicability and computational implications of adapting IRV to preference data to enable group fairness while satisfying a priority order.

Preference data considering faculty hiring. Table 1 represents ranked order of up to top-5 preferences over 7 candidates who have applied to a faculty position. Preferences are provided by 10 committee members (voters). Each of these ranked orders of preferences constitutes a ballot.

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 the author(s) 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.
KDD'24, August 25–29, 2024, Barcelona, Spain

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-XXXX-X/18/06
<https://doi.org/XXXXXXX.XXXXXXX>

Fairness and priority order. Group fairness is studied considering the protected attribute of the candidates to ensure equal representation of each group [19, 29]. The example assumes that research area is one such protected attribute with $k = 3$ different values (it is not hard to extend this to race, gender, ethnicity, or any other protected attribute). The value of this attribute for each of the candidates and the priority among these values is given in Table 2.

Goal. The goal is to return one candidate per protected attribute group that is most representative of the committee members' preferences while obeying the priority order. In our example, we need to first select a Data Mining (DM) candidate, then a Machine Learning (ML) candidate, and finally an Artificial Intelligence (AI) candidate.

The IRV process. The IRV process [21, 27] is a multi-stage process [31] that simulates $n - 1$ run-off rounds, where in each such round one item is eliminated. The single item that survived the eliminations after all rounds is the winner. More specifically, given the original preferences of the users (voters), an initial tally of the first choice votes of every candidate is performed in round 1. The item that has the lowest number of first choice votes is eliminated. Ties are broken arbitrarily. After the elimination, all the ranked orders that include the eliminated item are updated, and the items following this eliminated item in the ranked order are advanced one place up. This concludes round 1. This is iterated $n - 1$ times, namely, the tally is recomputed, and the item that has the lowest number of first choice votes is eliminated, where ties are broken arbitrarily.

Using the running example, as shown in the left of Table 4, the IRV process eliminates *Molly* in round 1, *Mira* in round 2 (and the respective vote gets transferred to *Sara*), and *Gina* in round 3. This process continues further making *Sara* the winner after 6 rounds.

Motivation. The resurgence of IRV is motivated by a range of expected benefits, including, ensuring majority support for the winner, reducing conflict within the electorate, reducing strategic voting, and increasing diversity of the winners [27]. IRV is amenable to incomplete ranked order, making the process further suitable for

Committee member	1st choice	2nd choice	3rd choice	4th choice	5th choice
Jack	Zoey	Mira			
Emma	Laura	Gina	Molly	Kim	Zoey
Monica	Zoey	Molly	Kim	Gina	Sara
Daniel	Zoey	Molly	Sara	Gina	
Max	Mira	Molly	Sara	Kim	Zoey
John	Sara	Gina	Kim	Zoey	
Amy	Gina	Sara	Kim	Mira	Zoey
Alice	Sara	Gina	Kim	Molly	Zoey
Bob	Kim	Gina	Sara	Molly	Zoey
Steve	Kim	Gina	Sara		

Table 1: Preferences over 7(n) candidates by 10 committee members(m) upto 5-th position (ℓ)

Priority Order	Protected attribute (area)	Candidates
First	DM	Molly, Laura
Second	ML	Gina, Kim, Sara
Third	AI	Zoey, Mira

Table 2: Fairness and Priority Orders

applications where users are not obligated to provide full order. Multiple recent works [14], [8] have demonstrated the superiority of IRV over plurality voting [23], as well as positional voting mechanisms (such as Borda [17]) to promote proportional representation of solid coalition and anti-plurality. Table 3 summarizes some of the advantages of IRV compared to other selection methods. Refer to Appendices A.1, A.2 for further details.

Method	Anti-plurality	Proportional representation	Suitable to incomplete order
Scoring based	×	×	×
Plurality	×	×	×
Positional	×	×	✓
IRV	✓	✓	✓

Table 3: Comparison of aggregation methods

IRV Margin computation. Recall that in our example the IRV process chooses *Sara* as the winner of the ballots. Clearly, *Sara* does not satisfy the priority order of selecting a DM candidate first. Hence, some ballot modifications are needed. If Jack’s ballot in Table 1 is changed by replacing *Zoey* with *Molly*, a series of 6 run-off rounds are simulated after that, as listed in the right of Table 4, which makes *Molly* the winner. If instead *Laura* is to be made the winner, this will require at least 3 ballot modifications, for example, by replacing the top choice of Alice, Monica, and John with *Laura*. Intuitively too, *Molly* is a better choice because it is liked as the second choice for 3 out of 7 original committee members. A similar process must also be carried out for ML and AI independently. Our goal is to find the minimum ballot modification that results in an outcome that satisfies the k priority orders. We refer to this problem as IRV margin computation [23] to satisfy k priority orders (denoted by **MqKIRV** for $k \geq 1$ and **MqIRV** for $k = 1$). To the best of our knowledge, we are the first to initiate a principled study on this.

Candidate	Tally	R1	R2	R3	R4	R5	R6	New Tally	R1	R2	R3	R4	R5	R6
Zoey	3	3	3	3	3	4		2	2	2	2			
Sara	2	2	3	4	4	6	8	2	2	3	3	3	5	
Kim	2	2	2	2	3			2	2	2	2	2		
Laura	1	1	1	1				1	1	1				
Gina	1	1	1					1	1					
Mira	1	1						1						
Molly	0							1	2	2	3	5	5	7

Table 4: IRV rounds after ballot modification (left): *Sara* winner
IRV rounds after ballot modification (right): *Molly* winner

Why ballot modification? An alternative to ballot modification could be the following - filter out candidates that do not satisfy the priority order (e.g., delete all ML and AI candidates for finding the top DM candidate) and run preference aggregation over the remaining ones. We note that such a filtering process that does not consider the order of all candidates on the ballot could heavily skew the results. Appendix A.3 has more details.

Technical Contributions (Sections 3 and 4). We make multiple technical contributions in terms of analyzing the studied problems as well as designing solutions for them. We prove that **MqIRV** is NP-hard, even when the ballot size is at most $\ell = 2$ by reducing an instance of the known NP-complete problem Exact Cover by 3-Sets (X3C) to an instance of **MqIRV**. Inspired by [11, 26] we then design

an algorithmic framework **ALGEXACT** that gives an exact solution and considers all possible permutations of the candidates that end in a candidate that satisfies the priority order. Solving **ALGEXACT** thus requires repeatedly solving a subproblem **DISTTo**, which, given a permutation, finds the smallest number of ballot modifications needed to ensure that the order of the candidates eliminated in $n - 1$ run-off rounds of IRV follow this order. Unfortunately, we prove that even the decision version of **DISTTo** is NP-hard by reducing an instance of X3C to **DISTTo**, even when $\ell = 3$.

We further study efficiency opportunities of **ALGEXACT** by enabling early terminations via branch and bound. The idea is to avoid making expensive **DISTTo** calls by computing a lower bound on the margin for every possible suffix of every permutation, and eliminating a permutation in its entirety if the lower bound on its margin is not smaller than the current upper bound on the margin of the **MqIRV** instance. To that end, we design a lower bound computation algorithm **DISTToLB** and an upper bound computation algorithm **MqIRVUB** that are highly effective and computationally lightweight. We also study the **DISTTo** problem under different preference manipulation models – for example, we study *how to only add the smallest number of ballots* to the existing set of ballots, such that the priority orders are satisfied. We refer to this as the **DISTToADD** problem. We present an efficient exact solution for the **DISTToADD** problem. We also present an integer programming formulation for **MqIRV** which is non-trivial. We finally design a highly scalable heuristics that is shown to work well in practice.

Experimental Evaluations (Section 5). Our final contribution is experimental - we use four real world large scale datasets motivated by different electoral voting and recommender systems applications, as well as one synthetically generated very large datasets. Our experimental evaluations have the following findings: **(a)** We empirically show that **MqIRV** results in a significantly smaller anti-plurality index [14] (i.e., it does not select candidates that are disliked by the majority of voters) compared to alternative approaches such as plurality voting [23] or Borda [17]. **(b)** We present an in-depth case study demonstrating that ballot modification results in a lower anti-plurality index compared to alternative approaches such as filtering. **(c)** We demonstrate that **ALGEXACT** is optimal, yet more scalable than existing solutions that could be adapted to our problem [26], [11]. **(d)** We empirically demonstrate the optimality of **DISTToADD** and its scalability, as well as the quality and scalability of our designed approximate solution by varying several pertinent parameters and comparing with appropriate additional baseline algorithms.

We present the discussion of related work in Section 6 (and in Appendix A.10) and conclude in Section 7.

2 DATA MODEL & PROBLEM

In this section, we describe the data model, following which we formally define the problem, and prove its hardness.

2.1 Data model

Ballot/preference. Preference of a user is elicited using a ballot b containing a ranking up to position at most ℓ , where c_i is the i -th preferred candidate. Using the running example, c_1 and c_5 are *Gina*, and *Zoey*, respectively of user *Amy*’s ballot.

Ballot profiles. The data contains the preferences/ballots $\mathcal{B}$ of m users/voters over a set C of n items/candidates. Using the running example, $m = 10$, $n = 7$. The columns in Table 1 show $\mathcal{B}$.

Preference aggregation. A preference aggregation method $\mathcal{F}$ takes $\mathcal{B}$ as input and selects a winner from the candidates/items. Given fairness criteria and priority order, the goal is to make use of $\mathcal{B}$ and $\mathcal{F}$ multiple (k) times to select k different winners. Table 2 shows $k = 3$ such constraints for recommending top DM, ML, and AI candidates. We use IRV as $\mathcal{F}$, as discussed in more detail in Section A.2.

Preference manipulation models. We consider two different preference manipulation models, where only the first one satisfies the number of ballot invariance property (i.e., the total number of votes remains unchanged) and is our primary focus in this work.

- (1) **Manipulation by modification.** Given a ballot b with ranking up to position j ($j \leq \ell$) positions, update any number of entries in b considering candidates from C . As an example, Jack's ballot (see Table 1) is changed from *Zoey*, *Mira* to *Molly*, *Mira*. Note that changing Daniel's ballot from *Zoey*, *Molly*, *Sara*, *Gina* to *Mira*, *Kim* also constitutes to a single ballot modification.
- (2) **Manipulation by addition.** Add a new ballot b with ranking of up to ℓ candidates from C .

Handling ties in IRV Recall that according to our definition ties during the IRV process are broken arbitrarily. It is not difficult to see that the way these ties are broken may impact the value of the margin. Indeed, in our example of ballot modification candidate *Molly* is the winner after just a single modification only in case the ties are broken in a very specific way. We postulate that any consistent choice would be effective in our case, since we use the margin to distinguish among choices and are not interested in the actual value of the margin.

2.2 Problem Definitions & Hardness

Problem Definition 1. *MqIRV (IRV Margin satisfying a single query constraint).* Given a set of ballots $\mathcal{B}$ eliciting m voters ranked preferences of up to ℓ positions over a given set C of n candidates, and a query constraint that specifies a subset of the candidates, find the minimum number of ballots that need to be modified in order to ensure that the winner of the IRV election belongs to the subset specified in the query constraint.

Running Example. Referring to Table 2, if the query constraint specifies selecting a DB candidate, then the minimum number of ballot modifications required to ensure that is 1, where *Zoey* in Jack's ballot is swapped by *Molly*. If instead *Laura* is to be made the winner, this will require 3 ballot modifications. Hence, the margin to satisfy the query constraint is 1.

THEOREM 2.1. *MqIRV is NP-Complete, even when $\ell = 2$.*

Problem Definition 2. *MqKIRV (IRV Margin satisfying k query constraints.)* Given a set of ballots $\mathcal{B}$ eliciting m voters ranked preferences of up to ℓ positions over a given set C of n candidates, and a query with k constraints, each specifies a subset of the candidates, find the minimum number of ballots that need to be modified in order to ensure that the winners of k independent invocations of the IRV election (each starting from the original ballots) belong to the respective subsets specified k query constraints.

THEOREM 2.2. *MqKIRV is NP-Complete, even when $\ell = 2$.*

PROOF. Follows trivially from Theorem 2.1. $\square$

Running Example. Considering the running example again (Table 2), $k = 3$ and the ballots are shown in Table 1. The winner for DB is *Molly* (margin = 1), for ML it is *Sara* (margin = 0), for horror it is *Zoey* (margin = 1). The minimum number of ballot modifications (margin) required to ensure all three constraints is $1+0+1 = 2$.

3 ALGORITHMS FOR MqIRV AND MqKIRV

In this section, we focus on designing exact solutions for **MqIRV** and **MqKIRV**. In Section 3.2 we discuss **ALGEXACT**, a branch-and-bound algorithm for **MqIRV** that is capable of effective pruning of the search space. In Section 3.3 we present a non-trivial integer programming formulation of **MqIRV**. These exact algorithms apply also exact algorithms for **MqKIRV** as follows from the following simple theorem.

THEOREM 3.1. *An optimal solution for MqKIRV corresponds to solving MqIRV optimally k times.*

3.1 Required Definitions

We first give some definitions that will be useful when discussing our algorithms.

Signature. Let $\mathcal{S}$ be the set of all possible partial or total rankings over C (including those that do not appear in $\mathcal{B}$). A signature $s \in \mathcal{S}$ is one such partial or total ranking. The total number of possible signatures is $|\mathcal{S}| = \sum_{x=1}^{\ell} \binom{n}{x} \cdot x!$. For example, both $\{Molly, Sara\}$ and $\{Zoey, Molly, Sara, Gina\}$ are valid signatures even though the former is not present in Table 1.

Tally $t_r(c)$ or first choice votes. The tally or first choice votes of a candidate c at round r , denoted as $t_r(c)$, is defined as the number of ballots in round r in which c is the first choice candidate. Using the running example, tally of *Sara*, *Zoey*, and *Kim* at the beginning of round 5 are: $t_5(Sara) = 4$, $t_5(Zoey) = 3$, and $t_5(Kim) = 3$.

3.2 ALGEXACT for MqIRV

We propose an algorithmic framework **ALGEXACT** that is an exact solution to the **MqIRV** problem. The algorithmic solution is developed by creating a branch and bound tree, akin to two prior works [11, 26].

For a given winner w , the solution considers all possible permutations of candidates that need to be eliminated (i.e., $(n-1)!$), where each permutation represents an elimination order simulating $n-1$ run-off rounds of IRV. The height of the tree is at most n . Each node of the tree contains two values: (a) an elimination order π , (b) a score that represents the number of ballot modifications to realize π (we formalize that as **DISTTo** below). Each edge of the tree represents the next candidate to be eliminated. An artificial root node connects the branches of the subtree, where each subtree represents a $w \in W$ as the winner, where W is the constrained winner set specified by the query. Except for the fake root node, the relationship between any parent and child nodes in the tree is as follows: (i) At any parent node with elimination order π , the child node has elimination order $\pi' = c + \pi$, for some $c \in C - \pi$, and (ii) **DISTTo**(π) $\leq$ **DISTTo**(π') [26]. The leaf nodes end with a full

permutation, where the last candidate is from W . The maximum number of possible leaf nodes is $|W| \times (n-1)!$. ALGEXACT solves the sub-problem DISTTO formalized below, repeatedly, at each node of the branch and bound tree.

Problem Definition 3. DISTTO. *Given an elimination order over the candidates π (complete or partial order, $|\pi| \leq |C|$) and a database of ballot profiles $\mathcal{B}$, find the minimum number of ballots that must be modified to achieve π .*

THEOREM 3.2. DISTTO is NP-hard, even when $\ell = 3$.

ALGEXACT explores the tree level by level (refer to Figure 9 in Appendix A.11) and makes an attempt to prune part of the tree based on the lower bound of a branch (which corresponds to an elimination order) and an upper bound of the value the MqIRV instance.

Definition 3.3. Upper bound of an instance MqIRVUB. Given an MqIRV instance, MqIRVUB is defined as an upper bound of the number of ballots that must be modified to satisfy the query constraint.

Definition 3.4. Lower bound (DISTTOLB) of DISTTO(π). Given an MqIRV instance and an elimination order π , DISTTOLB is a lower bound on the number of ballots that must be modified to achieve π , namely, $\text{DISTTOLB}(\pi) \leq \text{DISTTO}(\pi)$.

Running Example. Figure 9 shows one such partially constructed tree for our running example. The candidates are represented by their unique ids, and any red node and the sub-tree under them are fully pruned. Each such red node has $\text{DISTTOLB}(\pi)$ that is not smaller than the MqIRVUB of the MqIRV instance (e.g., $\text{DISTTOLB}([1, 3, 5]) = 4$ is larger than $\text{MqIRVUB} = 2$). Compared to prior works [11, 26], we propose both effective as well as computationally efficient MqIRVUB and DISTTOLB solutions, as we discuss in Section 4.

3.3 IP for MqIRV

MqIRV can be formulated as an integer linear program (IP). The objective of the IP is to minimize the number of ballot modifications required to ensure that the winner is the preferred candidate. Next, we describe how to formulate this IP.

For each ballot signature $s \in \mathcal{S}$, let m_s denote the number of ballots with signature s in the original ballot profile. Define $m = \sum_{s \in \mathcal{S}} m_s$, so that m counts the total number of ballots in the original election profile. Note that the values of m_s and m are determined by the original election profile. Let a_s denote the number of ballots that are modified to s from a different ballot signature, d_s denote the number of ballots that are modified from s to another ballot signature, and y_s denote the total number of ballots with signature s after the modifications. The following equations must be satisfied.

$$m_s + a_s - d_s = y_s \quad (1)$$

$$m \geq y_s \geq 0 \quad (2)$$

$$m_s \geq d_s \geq 0 \quad (3)$$

$$m - m_s \geq a_s \geq 0 \quad (4)$$

Equation 1 requires that the number of ballots with a new signature s be equal to the number of ballots that originally had the

Algorithm 1 ALGEXACT

Input: Ballot profile $\mathcal{B}$, set of Candidates C , set of preferred candidates W .

Output: MqIRV

```

1:  $ub = \infty$ 
2:  $lb = 0$ 
3: initialize priority queue with tuples  $(w, 0)$  where  $w \in W$ 
4: while queue.notEmpty() do
5:    $\pi, lb = \text{queue.pop}()$ 
6:   for  $c \in C \setminus \pi$  do
7:      $\pi' = c + \pi$ 
8:      $lb = \text{DISTTOLB}(\mathcal{B}, C, \pi')$ 
9:     if  $lb > ub$  then
10:      prune  $\pi'$ 
11:   else
12:     queue.add( $\pi', lb$ )
13:   end if
14:   if  $|\pi'| == n$  then
15:      $ub = \min(ub, \text{DISTTO}(\mathcal{B}, C, \pi'))$ 
16:   end if
17: end for
18: end while
19: MqIRV =  $ub$ 
20: Return MqIRV

```

signature s , plus the number that changed from something else to s , minus the number that changed from s to something else. Equation 2 constrains that the number of ballots that end with signature s cannot be more than the total number of ballots that were cast in the election. The next two equation requires that the number of ballots that are modified to have signature s must be nonnegative and no more than the number of ballots that had a signature different than s originally, and one cannot change more ballots of signature s than the number of ballots that originally had the signature s .

The next constraint is that the total number of ballots changed from any signature is equal to the total number of ballots changed to any signature.

$$\sum_{s \in \mathcal{S}} a_s = \sum_{s \in \mathcal{S}} d_s \quad (5)$$

The next two constraints correspond to the elimination order. Assume C is the set of all candidates. For every pair $\{c_i, c_j\} \subseteq C$, define u_{c_i, c_j} as a binary variable that is 1 iff candidate c_j is eliminated before candidate c_i . For completeness also define $u_{c_i, c_i} = 1$, for every $c_i \in C$. The following constraints guarantee that the variables u_{c_i, c_j} define an order. Equation 6 constrains it to be antisymmetric and Equation 7 constrains it to satisfy the triangle inequality.

$$u_{c_i, c_j} + u_{c_j, c_i} = 1 \quad \forall \{c_i, c_j\} \subseteq C \quad (6)$$

$$u_{c_i, c_j} + u_{c_j, c_r} + u_{c_r, c_i} \geq 1 \quad \forall \{c_i, c_j, c_r\} \subseteq C \quad (7)$$

For a signature s of an original ballot and candidates $\hat{c}$ and $\tilde{c}$ (which may be equal), define the binary variable $v_{s, \hat{c}, \tilde{c}}$ to be 1 iff when candidate $\tilde{c}$ is eliminated $\hat{c}$ is the top candidate in the signature that had originally been signature s . Bit $v_{s, \hat{c}, \tilde{c}}$ is trivially 0 if $\hat{c}$ does not appear in s . Let signature $s = c_1, c_2, \dots, c_\ell$, where c_x is the x -th candidate on the ballot, c_1 is the top choice while c_ℓ is the bottom.

Assume from now on that $\hat{c} = c_i$. For candidate c_j in s , where $j < i$, bit v_{s,c_i,c_j} is 0 since c_j is ranked higher than c_i in s . Assume from now on that $\tilde{c} \neq c_j$, for $j \in \{1, \dots, i-1\}$. We have the following constraint on $v_{s,c_i,\tilde{c}}$.

$$v_{s,c_i,\tilde{c}} = u_{c_i,\tilde{c}} \cdot \prod_{x=1}^{i-1} u_{\tilde{c},c_x} \quad (8)$$

This constraint ensures that all the candidates $c_1, c_2, \dots, c_{i-1}$ are eliminated before $\tilde{c}$ is eliminated, and in case $c_i \neq \tilde{c}$, candidate c_i is eliminated after $\tilde{c}$ is eliminated. Thus signature s contributes to c_i 's tally when $\tilde{c}$ is eliminated. Note that since by definition $u_{c_1,c_1} = 1$, we get that $v_{s,c_1,c_1} = 1$, which holds trivially. The constraint in its current format is not linear since it is a product of bits. Later, we show how to convert it to linear constraints.

The next constraint is for every ordered pair of candidates $\hat{c} \neq \tilde{c}$. It guarantees that if $u_{\hat{c},\tilde{c}} = 1$, namely $\hat{c}$ is eliminated after $\tilde{c}$, then in the round in which $\tilde{c}$ is eliminated the number of ballots in which $\hat{c}$ is the top candidate is at least the number of ballots in which $\tilde{c}$ is the top candidate. The constraint is written as a product of bits and an integer (later, we show how to convert it to linear constraints).

$$\sum_s (y_s \cdot v_{s,\hat{c},\tilde{c}}) \geq u_{\hat{c},\tilde{c}} \cdot \sum_s (y_s \cdot v_{s,\tilde{c},\hat{c}}) \quad (9)$$

If we want to force candidate $\hat{c}$ to be the winner we need to add the constraints $u_{\hat{c},\tilde{c}} = 1$, for every $\tilde{c} \neq \hat{c}$. Alternatively, if we want to force candidate $\tilde{c}$ not to be the winner we need to add the constraint $\sum_{\hat{c} \neq \tilde{c}} u_{\hat{c},\tilde{c}} \geq 1$. In addition, we can change the objective function to count only additions or only deletions or any linear combination of additions, deletions, and modifications. For our case we set the objective function to be: minimize $\sum a_s$, which is the number of ballots modifications.

In the last two constraints, we used (i) product of bits, and more generally (ii) product of a nonnegative number and bits. We show how to linearize a product of a nonnegative number and bits as long as we have an upper bound on the number. Let $u_1, \dots, u_x$ be x bits, and A be a non-negative number. Assume that m is an upper bound on A . (As in our case, since m is the total number of signatures.) The constraints that replace $Z = A \cdot \prod_{i=1}^x u_i$ are as follows.

$$\begin{aligned} z &\leq u_i \cdot m && \text{for } i \in \{1, \dots, x\} \\ z &\leq A \\ z &\geq A + \left(\sum_{i=1}^x u_i - x \right) \cdot m \\ z &\geq 0 \end{aligned}$$

4 EFFICIENT ALGORITHMS

This section is dedicated to further investigation of computational efficiency. In Section 4.1, we describe an improved algorithm for computing DistToLB . In Section A.6 (in the Appendix), we discuss an improved **MqIRVUB** algorithm that is computationally efficient. Interestingly, this algorithm can be applied as an efficient heuristic for the **MqIRV** problem. In Section A.7 (in the Appendix), we describe an efficient (polynomial time) algorithm for DistTo in case only ballot additions are allowed. Thus, demonstrating that DistTo becomes a computationally tractable in this special case.

4.1 An Improved DistToLB Algorithm

In this section, we discuss an improved lower bound calculation algorithm for $\text{DistTo}(\pi)$. The intuition is the following: given π and two candidates c and c' , if c needs to be eliminated before c' in round i , where $t_i(c)$ and $t_i(c')$ are the number of first choice votes of c and c' in round i , respectively, then at least $\left\lceil \frac{t_i(c) - t_i(c')}{2} \right\rceil$ number of first choice votes from c needs to go to c' . That is, lb , the lower bound of round i is calculated as the half of the difference of tally between c and c' . Finally, the maximum over all of these is returned as the output of the algorithm. Algorithm 2 has the pseudocode.

Algorithm	Efficient ALGEXACT	Blom
Number of IP calls	AI: 1 ML: 1 DM: 2	AI: 143 ML: 108 DM: 107
Runtime (s)	0.057	0.626

Table 5: Efficiency improvement using **MqIRVUB and DistToLB for the running example**

Running example. Assume, $\pi = [\text{Gina}, \text{Molly}, \text{Zoey}] = [4, 6, 0]$ where 4 is eliminated first. Initially, $t_1(\text{Gina}) = 6$, $t_1(\text{Zoey}) = 3$, $t_1(\text{Molly}) = 1$. To ensure *Gina* is eliminated, at least $\max\{\left\lceil \frac{6-1}{2} \right\rceil, \left\lceil \frac{6-3}{2} \right\rceil\} = 3$ ballot modifications are required. After *Gina* is eliminated, $t_2(\text{Zoey}) = 5$, $t_2(\text{Molly}) = 4$. Required modifications of the ballot to ensure that *Zoey* wins = $\left\lceil \frac{5-4}{2} \right\rceil = 1$. Therefore, $lb = \max(3, 1) = 3$.

Using the running example, Algorithm 2 reduces a significant number of DistTo (which is solved using IP) calls. For example, $lb = \text{DistToLB}([4, 6, 0]) = 3 \leq \text{DistTo}([4, 6, 0])$. Hence **ALGEXACT** prunes the branch $[4, 6, 0]$ without having to make an expensive DistTo call (this is because lb for this branch $> ub$). Table 5 shows efficiency improvement using DistToLB and **MqIRVUB** inside **ALGEXACT** over prior work [11].

Algorithm 2 Algorithm for DistToLB

Input: Set of ballots $\mathcal{B}$, an elimination order π

Output: $\text{DistToLB}(\text{DistTo}(\pi))$

```

1:  $lb = 0$ 
2: while  $|\pi| > 1$  do
3:    $c = \pi.pop\_front()$ 
4:   for  $c' \in \pi \setminus c$  do
5:      $lb = \max(lb, \left\lceil \frac{t_i(c) - t_i(c')}{2} \right\rceil)$ 
6:   end for
7: end while
8: Return  $lb$ 
```

THEOREM 4.1. *Algorithm 2 returns a valid lower bound on $\text{DistTo}(\pi)$.*

Lemma 1. *The running time of Algorithm 2 is $O(n^2 + m\ell)$.*

5 EXPERIMENTAL EVALUATIONS

We conducted experiments to analyze our algorithms, implemented in Python 3.8 on a Windows 11, i7, 16GB RAM setup. Results are averages from 10 runs. The code and data could be found in the [github \[3\]](#).

5.1 Experiment Design

We have three goals. (a) Assess the effectiveness of **MqKIRV** (Section 5.2). (b) Evaluate the quality of our designed algorithms for **MqIRV** and **MqKIRV** problems (Section 5.3). (c) Evaluate their scalability (Section 5.4). We analyzed four real-world and one synthetic dataset, with comprehensive details provided in Table 6.

Dataset Name	m	n	Description
NSW Senate Election data	533	1,520k	Candidates from five parties.
San Francisco Election data	18	193k	Board of supervisors, district attorney, and mayoral results.
MovieLens	100k	100k	User movie ratings.
Adressa News	100k	100k	News articles with user ratings.
Synthetic	1,000k	1,000k	Random preference rankings.

Table 6: Real world and synthetic datasets
(m denotes number of candidates and n number of voters)

5.1.1 Baseline Algorithms. The following algorithms are implemented.

1. Filtering-Borda [28]. We implement a baseline where candidates who do not satisfy the query constraints are first filtered out. Then, considering the remaining candidates, the preferences of the voters are aggregated using the "positional" scoring mechanism Borda [28] that assigns a score to each candidate corresponding to the positions in which a candidate appears within each voter's ballot. This baseline is implemented to evaluate two aspects: 1. Why a ballot modification is necessary and 2. effectiveness of a different positional aggregation mechanism and its effectiveness over IRV.

2. Plurality voting [18, 20]. The winner is the candidate who represents a plurality of voters' first choice or, in other words, receives more first choice votes than any other candidate. That makes plurality voting among the simplest of all electoral systems. This baseline is implemented to evaluate effectiveness of a non positional aggregation mechanism and its effectiveness over IRV.

3. Blom et al. [11]. Magrino et al. [26] propose a simple lower bound based on the DistTo of any π of length n . Given two elimination orders, if one is the suffix of another, then, the DistTo of the suffix could be used as the DistToLB of DistTo for the longer elimination order. Blom et al. [11] propose an improved lower bound over [26] based on the last round margin $l(c', c)$ between any pair of candidates c and c' (to ensure c' is eliminated before c), where $l(c', c)$ is half of the difference in the tallies of c' and c (first choice votes). This idea is generalized to generate lower bound of margin to ensure an elimination order ending in π , which is $\max\{l(c', c)\}$, where $c' \in C - \pi$, $c \in \pi$.

4. Random. We implement an algorithm that runs iteratively. In the first iteration, it randomly selects a ballot and modifies it. In the next iteration, it doubles the number of selected ballots to be modified (and selects those ballots randomly) and repeats the process until the query constraints are satisfied.

5. IP for DistToAdd. We implement an integer programming based solution for the DistToAdd problem.

These algorithms are compared against our proposed DistToLB and **MqIRVUB** solutions inside **ALGEXACT**. We also compare **ALGAPPRX** against these solutions and the implemented IP for **MqIRV**. Finally, we compare our designed solution DistToAddAlg with its corresponding IP implementation.




Figure 1: Anti-plurality index using NSW election dataset

5.1.2 Measures. To evaluate anti-plurality, we measure the anti-plurality index that is proposed in a related work [14]. Anti-plurality index of a preference aggregation method is computed by looking at each winner candidate i that the method produces and then calculating the percentage of voters who prefer i the least (i.e., it is the last choice on their ballots). The average anti-plurality index is then calculated by taking the average over multiple queries. To evaluate the quality of our designed algorithms, we compare approximation factors of margins produced by different algorithms (margin produced by the proposed algorithm/ exact margin), as well as compare the exact margin values. Finally, we compare the effectiveness of the proposed algorithms based on the number of expensive DistTo calls they make (smaller is better). To evaluate scalability, we evaluate the pruning effectiveness of the algorithms and the overall running time.

5.1.3 Query and Parameters. Query constraints are generated randomly but by using party affiliation for NSW datasets, race of the candidates from the San Francisco Election dataset, and movie genre, and news type of the last two datasets, respectively. For evaluating **MqIRV**, we vary the size of the ballot (ℓ), number of users (m), and the number of candidates (n). We consider various combinations over these parameters to cover a wide range of recommendation settings. The default values are $n = 10$, $\ell = 4$ and $m = 1000$.

5.2 Goal 1: Analyzing Anti-plurality

For these experiments, NSW dataset is used. For each query, the set W is selected arbitrarily based on the 5 different party affiliation of the candidates – Labor Party or LAB, Christian Democratic Party or CDP, National Party or NLT, Liberal Party or LIB, The Greens or GRN. We compare average anti-plurality index of **MqIRV** margin computation based on plurality voting [18, 20] and margin computation based on Filtering-Borda in Figure 1 after running 133 queries. These results clearly indicate that **MqKIRV** results in significantly anti-plurality compared to the other baselines.

5.2.1 A case study. We present a case study to demonstrate efficacy of **MqIRV** to overcome anti-plurality. A smaller subset of NSW election data is used that contains 12 candidates and 33, 553 voters. A query is generated to select candidates that are either LIB or LAB. This makes $W = \{2, 5, 8, 10\}$ (these numbers are the unique ids of the candidates). **MqIRV** selects candidate 8 as the winner, while, Plurality voting and Filtering-Borda both select candidate 5. Upon further analysis, it appears that a total of 9884 voters like candidate 5 as their first choice, while a total of 5411 voters dislike candidate 5 (these voters place candidate 5 as their last choice on their ballots). For candidate 8, these two numbers are 9483 and 1863, respectively. In fact, about 25% of the voters put candidate 5 as one

of their 3 least preferred candidates compared to only 2% voters that do so for candidate 8. This case study anecdotally demonstrates the efficacy of **MqIRV** to overcome anti-plurality. This case study also demonstrates why filtering based approach could skew the results, which **MqIRV** avoids by looking at the entire ballot and the order of all candidates.

5.3 Goal 2: Analyzing Quality

Approximation factor. In Table 7, we present the approximation factors of the **MqIRV** problems solved using different algorithms. The results are shown for 4 real datasets. Two of the exact solutions are compared against the IP formulation of **MqIRV** and exhibit approximation ratio of 1, as expected. **ALGAPPRX** has an approximation ratio between 1.91 and 3.15. On the other hand, *Random* has an approximation ratio between 3.61 and 4.21. As analyzed analytically, **DISTToADDALG** is an exact solution of **DISTToADD** and has an approximation ratio of 1.

Dataset	ALGEXACT	DISTToADD	ALGAPPRX	Random
NSW dataset	1	1	1.97	3.41
San Francisco Election	1	1	1.98	3.96
MovieLens	1	1	1.99	3.42
Adressa News	1	1	3.15	4.21

Table 7: Approximation factor of the algorithms

Margin. Figure 2 shows the box plot of difference in margin for **ALGAPPRX** and **ALGEXACT** varying n for all 4 real datasets over 10 different queries. These results corroborate that **ALGAPPRX** is an effective solution across all 4 datasets.

We also analyze the margin difference between **ALGAPPRX** and *Random* using one synthetic dataset and 3 real datasets varying n up to 1 million. For each run, we keep the number of ballots $m = n$. Figure 5 shows **ALGAPPRX** always returns smaller margin than *Random*. Using MovieLens data, *Random* margin is about 20 times larger than **ALGAPPRX**.

Number of DISTTo IP calls. Finally, we show that **ALGEXACT** requires significantly less number of IP calls compared to *Blom* (Figure 6). On Adressa News dataset on $n = 10$, **ALGEXACT** invokes about 17 times less number of IP calls than what *Blom* does. These results demonstrate the effectiveness of our proposed **DISTToLB** and **MqIRVUB** solutions, compared to the adapted version of [11].

5.4 Goal 3: Analyzing Scalability

Running time. In these experiments (Figure 3), we compare running time in seconds for **ALGEXACT**, **ALGAPPRX**, and *Blom* on 4 real world datasets by varying n , while keeping ℓ and m fixed. The exact algorithms show that the running time increases exponentially with increasing n . **ALGAPPRX** is almost 24333 times faster than *Blom* for $n = 12$ using MovieLens dataset. While **ALGEXACT** is 7.6 times faster than *Blom* for $n = 12$ using MovieLens dataset.

Figure 7 presents effect of varying ℓ and m on running time of **ALGEXACT**, **ALGAPPRX**, and *Blom* on 2 real world datasets. As expected, running time **ALGEXACT** does not significantly vary with increasing m and ℓ , as it is mostly driven by exponential 2^n cost of branch & bound tree exploration.

Running time in very large scale data. For these experiments, we compare running time of our efficient solution **ALGAPPRX** and compare that with *Random*. Figure 8 shows that the running time of

ALGAPPRX is significantly smaller than *Random*. Using the Adressa News dataset with $n = 100k$, $m = 100k$ and $\ell = 4$, the runtime for *Random* is about 100 times higher than **ALGAPPRX**.

Running time of DISTToADDALG. Figure 4 compares the running time between our exact solution **DISTToADDALG** for **DISTToADD** with IP based implementation (**DISTToIPADD**). **DISTToIPADD** runtime increases exponentially with n as expected, whereas, **DISTToADDALG** runs in $O(n^2)$ time. For MovieLens dataset with $n = 10$ **DISTToADDALG** is 53 times faster than **DISTToIPADD**.

5.5 Summary of Results

Our first observation is that, **MqKIRV** significantly promotes lower anti-plurality, whereas, the other baselines do not. The case study demonstrates that ballot modification selects winner with lower anti-plurality index than plurality voting and a filtering based approach (Filtering-Borda) that could be myopic at times. Our second major observation is that our designed **ALGEXACT** enabled by effective lower bound **DISTToLB** and upper bound **MqIRVUB** algorithm is highly effective as well as computationally efficient compared to their counterparts *Blom*. Third, **ALGAPPRX** exhibits empirical approximation factor around 2 (for 3 of the datasets) and runs significantly faster than the exact solutions (order of magnitude faster) and the *Random* baseline. Finally, consistent with our theoretical analysis, **DISTToADDALG** returns an exact solution for **DISTToADD**, runs in polynomial time, and is significantly faster (about 53 times for some datasets) than the IP based solution.

6 PRIOR WORK

We present related work covering three areas: (a) preference aggregation methods, (b) how to minimally update preferences so that the produced outputs satisfy additional criteria, and (c) multi-stage preference aggregation methods and their margin of victory computation.

Due to space constraints, the details are given in Appendix A.10. We remark that it is evident from this prior work that we are the first to study an IRV based multi-stage preference aggregation procedures [31]. Also, our margin finding problem **MqKIRV** is different from previously known MoV problems, and our hardness results and algorithmic solutions to this problem extend the state of the art in this area.

7 CONCLUSION

We study the suitability of Instant Run-off Voting (IRV) as a preference aggregation method to select k different winners to promote group fairness and priority. We formalize an optimization problem that aims at finding the margin, i.e., the smallest number of modifications of original users' preferences (ballots) so that the selected k winners satisfy all these query constraints. We present principled models and several non-trivial algorithmic and theoretical results. Our experimental analyses demonstrate suitability of IRV as a preference aggregation method over plurality voting and a filtering based approach, as well as corroborate our theoretical analysis.

This work opens up many interesting directions – as an ongoing work, we are investigating how to design approximation algorithms with theoretical guarantees for **MqIRV**. We are also studying how our proposed solution **ALGEXACT** could be adapted to compute the margin for single transferable voting (STV) schemes.




Figure 2: Margin difference between ALGAPPRX and ALGEXACT varying n Figure 3: Runtime for ALGAPPRX, ALGEXACT, and *Blom* varying n 

Figure 4: Runtime for DISTTOADDALG and DISTTOIPADD




Figure 5: Margin for ALGAPPRX and Random




Figure 6: #IP calls for ALGAPPRX and ALGEXACT varying n Figure 7: Runtime for ALGAPPRX, ALGEXACT, and *Blom* varying l, m 

Figure 8: Runtime for ALGAPPRX & Random

REFERENCES

- [1] Sihem Amer-Yahia, Behrooz Omidvar-Tehrani, Senjuti Basu, and Nafiseh Shabib. 2015. Group recommendation with temporal affinities. In *International Conference on Extending Database Technology (EDBT)*.
- [2] Sihem Amer-Yahia, Senjuti Basu Roy, Ashish Chawlat, Gautam Das, and Cong Yu. 2009. Group recommendation: Semantics and efficiency. *Proceedings of the VLDB Endowment* 2, 1 (2009), 754–765.
- [3] Anonymous. 2023. Git link. https://anonymous.4open.science/r/selection_queries_using_irv-5AD0/README.md.
- [4] Manel Ayadi, Nahla Ben Amor, Jérôme Lang, and Dominik Peters. 2019. Single Transferable Vote: Incomplete Knowledge and Communication Issues. In *18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS 19)*. International Foundation for Autonomous Agents and Multiagent Systems, Montreal QC, Canada, 1288–1296. <https://hal.science/hal-02307486>
- [5] Linas Baltrunas, Tadas Makcinskas, and Francesco Ricci. 2010. Group recommendations with rank aggregation and collaborative filtering. In *Proceedings of the fourth ACM conference on Recommender systems*. 119–126.
- [6] John J. Bartholdi and James B. Orlin. 1991. Single transferable vote resists strategic voting. *Social Choice and Welfare* 8, 4 (1991), 341–354. <http://www.jstor.org/stable/41105995>
- [7] Senjuti Basu Roy, Laks VS Lakshmanan, and Rui Liu. 2015. From group recommendations to group formation. In *Proceedings of the 2015 ACM SIGMOD international conference on management of data*. 1603–1616.
- [8] Rachel Behar and Sara Cohen. 2022. Representative Query Results by Voting. In *Proceedings of the 2022 International Conference on Management of Data*. 1741–1754.
- [9] Arnab Bhattacharyya and Palash Dey. 2021. Predicting winner and estimating margin of victory in elections using sampling. *Artificial Intelligence* 296 (2021), 103476. <https://doi.org/10.1016/j.artint.2021.103476>
- [10] Michelle Blom, Peter J. Stuckey, and Vanessa J. Teague. 2017. Towards Computing Victory Margins in STV Elections. [arXiv:1703.03511](https://arxiv.org/abs/1703.03511) [cs.GT]
- [11] Michelle Blom, Peter J. Stuckey, Vanessa J. Teague, and Ron Tidhar. 2015. Efficient Computation of Exact IRV Margins. [arXiv:1508.04885](https://arxiv.org/abs/1508.04885) [cs.AI]
- [12] Da Cao, Xiangnan He, Lianhai Miao, Yahui An, Chao Yang, and Richang Hong. 2018. Attentive group recommendation. In *The 41st International ACM SIGIR conference on research & development in information retrieval*. 645–654.
- [13] David Cary. 2011. Estimating the Margin of Victory for Instant-Runoff Voting. In *Conference on Electronic voting technology/workshop on trustworthy elections*. USENIX Association, San Francisco, CA.
- [14] Abhijnan Chakraborty, Gourab K Patro, Niloy Ganguly, Krishna P. Gummadri, and Patrick Loiseau. 2018. Equality of Voice: Towards Fair Representation in Crowdsourced Top-K Recommendations. [arXiv:1811.08690](https://arxiv.org/abs/1811.08690) [cs.SI]
- [15] Vincent Conitzer, Tuomas Sandholm, and Jérôme Lang. 2007. When Are Elections with Few Candidates Hard to Manipulate? *J. ACM* 54, 3 (Jun 2007), 14:1–14:33. <https://doi.org/10.1145/1236457.1236461>
- [16] Palash Dey and Y. Narahari. 2015. Estimating the Margin of Victory of an Election using Sampling. [arXiv:1505.00566](https://arxiv.org/abs/1505.00566) [cs.AI]
- [17] Peter Emerson. 2013. The original Borda count and partial voting. *Social Choice and Welfare* 40 (2013), 353–358.
- [18] Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. 2017. Multiwinner voting: A new challenge for social choice theory. *Trends in computational social choice* 74, 2017 (2017), 27–47.
- [19] David García-Soriano and Francesco Bonchi. 2021. Maxmin-fair ranking: individual fairness under group-fairness constraints. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 436–446.
- [20] Bernard Grofman, Guillermo Owen, and Scott L Feld. 1983. Thirteen theorems in search of the truth. *Theory and decision* 15, 3 (1983), 261–278.
- [21] Wm. H. Hare. 1871. Application of Mr. Hare's System of Voting to the Nomination of Overseers of harvard College. *Journal of Social Science: Containing the Transactions of the American Social Science Association* 3-4 (1871), 192–198. <https://books.google.com/books?id=W7QRAAAAYAAJ>
- [22] Steven Hill and Robert Richie. 2005. Success for instant runoff voting in San Francisco. *National Civic Review* 94, 1 (2005), 65–69.
- [23] Md Mouinul Islam, Dong Wei, Baruch Schieber, and Senjuti Basu Roy. 2022. Satisfying complex top-k fairness constraints by preference substitutions. *Proceedings of the VLDB Endowment* 16, 2 (2022), 317–329.
- [24] Alborz Jelvani and Amelie Marian. 2022. Identifying Possible Winners in Ranked Choice Voting Elections with Outstanding Ballots. *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing* 10, 1 (Oct. 2022), 114–123. <https://doi.org/10.1609/hcomp.v10i1.21992>
- [25] Jae Kyeong Kim, Hyea Kyeong Kim, Hee Young Oh, and Young U Ryu. 2010. A group recommendation system for online communities. *International journal of information management* 30, 3 (2010), 212–219.
- [26] Thomas Magrino, Ronald Rivest, Emily Shen, and David Wagner. 2011. Computing the margin of victory in IRV elections. In *2011 Electronic Voting Technology Workshop/Workshop on Trustworthy Elections (EVT/WOTE 11)*. USENIX Association, San Francisco, CA, 4–4. <https://www.usenix.org/conference/evtwote-11/computing-margin-victory-irv-elections>
- [27] Eamon McGinn. 2020. Rating Rankings: Effect of Instant Run-off Voting on participation and civility. http://eamonmcginn.com.s3.amazonaws.com/papers/IRV_in_Minneapolis.pdf
- [28] Shmuel Nitzan and Ariel Rubinstein. 1981. A further characterization of Borda ranking method. *Public choice* (1981), 153–158.
- [29] Evaggelia Pitoura, Kostas Stefanidis, and Georgia Koutrika. 2022. Fairness in rankings and recommendations: an overview. *The VLDB Journal* (2022), 1–28.
- [30] Abinash Pujahari and Dilip Singh Sisodia. 2020. Aggregation of preference relations to enhance the ranking quality of collaborative filtering based group recommender system. *Expert Systems with Applications* 156 (2020), 113476. <https://doi.org/10.1016/j.eswa.2020.113476>
- [31] Senjuti Basu Roy. 2022. Returning Top-K: Preference Aggregation or Sortition, or is there a Better Middle Ground? *SIGMOD Blog* (2022).
- [32] Senjuti Basu Roy, Saravanan Thirumuruganathan, Sihem Amer-Yahia, Gautam Das, and Cong Yu. 2014. Exploiting group recommendation functions for flexible preferences. In *2014 IEEE 30th international conference on data engineering*. IEEE, Chicago, IL, USA, 412–423.
- [33] Senjuti Basu Roy, Saravanan Thirumuruganathan, Sihem Amer-Yahia, Gautam Das, and Cong Yu. 2014. Exploiting group recommendation functions for flexible preferences. In *2014 IEEE 30th international conference on data engineering*. IEEE, 412–423.
- [34] A. D. Sarwate, S. Checkoway, and H. Shacham. 2012. Risk-Limiting Audits and the Margin of Victory in Nonplurality Elections. *Statistics, Politics and Policy* 3, 3 (December 2012), 29–64. <https://doi.org/10.1515/spp-2012-0003>
- [35] Dong Wei, Md Mouinul Islam, Baruch Schieber, and Senjuti Basu Roy. 2022. Rank aggregation with proportionate fairness. In *Proceedings of the 2022 International Conference on Management of Data*. 262–275.
- [36] Lirong Xia. 2012. Computing the Margin of Victory for Various Voting Rules. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC '12)* (Valencia, Spain). Association for Computing Machinery, New York, NY, USA, 982–999. <https://doi.org/10.1145/2229012.2229086>
- [37] Meike Zehlke, Ke Yang, and Julia Stoyanovich. 2021. Fairness in ranking: A survey. *arXiv preprint arXiv:2103.14000* (2021).

A APPENDIX

A.1 Different Aggregation Mechanisms

Plurality voting. In a plurality voting system, each voter is allowed to vote for one candidate, and the candidate who receives the most votes wins, regardless of whether they secure a majority of the votes. This system is straightforward and easy to understand but can result in a "winner-takes-all" outcome where the elected representative may not reflect the preference of the majority of voters. For instance, in our current example, *Zoey* wins the plurality vote with just 3 ballots in her favor. However, *Zoey* is also the least favored choice of 5 voters, underscoring the system's limitation in capturing the majority's true preference.

Scoring based. In scoring-based voting systems, voters score each candidate independently on a scale (e.g., 0 to 5 or 1 to 10). The scores for each candidate are then aggregated to determine the winner. This system allows voters to express not just a preference order but also the intensity of their preferences. Examples of scoring-based voting systems include Range Voting and Approval Voting. However, when the users provide preferences in a ranked order, there is no standard way to convert those preferences to scores.

Positional voting. Positional voting systems allow voters to rank candidates in order of preference. The most common form of positional voting is the *Borda Count*, where points are assigned to positions on the voters' preference lists. Using the running example, *Gina* emerged as the clear winner of *Borda Count* with a total of 17 points (3 points from Emma, 1 point from Monica, 0 point from Daniel, 2 points from John, 4 points from Amy, 3 points from Alice, 3 points from Bob, 1 points from Steve.). *Gina*'s consistent ranking as a second top choices of many voters secured her victory.

A.2 IRV Properties

IRV is known to satisfy properties [14] that other preference aggregation measures are unable to accommodate.

IRV promotes proportional representation for solid coalitions. **Definition.** In social choice theory, a solid coalition for a set of candidates is defined as a set of voters who all rank every candidate in that set higher than any candidates outside that set. This criterion requires that if the number of such voters is at least half of the total number of voters, then one of those candidates from that set must win.

Consider a scenario in which two candidates with similar ideologies compete over the same pool of voters, resulting in divided votes and potentially allowing a third candidate with a different ideology that has fewer overall votes to win. IRV fulfills this criterion, whereas plurality voting [20] fails to do so. To demonstrate this property, notice that in our running example, there exists a solid coalition of voters who like ML (refer to Table 1 which shows 5 of the 10 users, John, Amy, Alice, Bob, and Steve rank the three ML candidates *Gina*, *Kim*, *Sara* higher than any other candidate). Clearly, if user preferences are aggregated using plurality voting, none of the ML candidates will be returned as the winner since *Zoey* has the highest number of first place votes, and will be selected as the winner. On the contrary, IRV will select *Sara* as the winner, and hence it is resistant to the ballot splitting problem.

IRV promotes anti-plurality. In social choice theory, the *majority loser criterion* was proposed to evaluate single-winner elections. It states that if a majority of voters prefer every other candidate over a given candidate, then that candidate must not win. IRV fulfills this criterion [14] (as there is a solid coalition for the rest of the candidates). Indeed, the candidate *Zoey* is the last choice of 6 out of the 10 users (Table 1), and thus IRV will not select it. Contrarily, plurality voting will select *Zoey* as the winner. In [14] this criterion is extended to define *anti-plurality* which requires that no candidate among the bottom $x\%$ of the ranked choices for the majority of the voters should be selected. Although not guaranteed, it is empirically shown in [14] that IRV fulfills this extended criterion anti-plurality frequently.

IRV vs. Plurality Voting A popular voting mechanism is plurality voting, that selects that the winner that receives the highest number of top ranked votes. Using Table 1, note that plurality voting will choose *Laura* as the winner among the candidates in DM area, even though it is clear that between *Laura* and *Molly*, the latter is more preferred by the users. As we will demonstrate later our IRV based process will indeed choose *Molly*. Finally, it is known that finding the *margin* (the number of ballots that must be substituted in order to change the original winner [16, 23, 34, 36]) for IRV is NP-hard [11], making IRV less susceptible to manipulation.

IRV vs. Scoring based Voting Scoring-based voting systems face challenges in terms of the proportionality of solid coalitions, anti-plurality, and having the complete preferences of voters in a way that accurately reflects voter intent. In scoring-based voting systems, the proportionality of solid coalitions as a strong preference for a particular candidate can be undermined if voters give nearly as high scores to other candidates, thus not providing a clear advantage to the coalition's preferred candidate. This aspect can also impact anti-plurality; since voters might give high scores to a candidate who is the last choice by the majority of the voters, this can lead to a winner who is disliked by most. In contrast, IRV upholds all three of these properties.

IRV vs. Positional Voting Positional Voting, like the Borda Count, focuses on selecting the most preferred candidate rather than reflecting the depth of support among multiple choices. This may lead to a winner-takes-all outcome, often favoring the candidate with the highest first-choice or second-choice support, potentially disregarding the proportional strength of coalitions. For the same reason, Positional Voting can select a winner who is the last choice by the majority of voters. Let's consider an example where voters rank five candidates A, B, C, D, and E in order of preference. In this election, there are 20 votes of preference order {A, D, E, C}, indicating that A is ranked first, followed by D, E, and C; 10 votes of preference order {B, C, E, A}; and 11 votes of preference order {C, B, D, A} are cast. Applying the Borda Count, Candidate A accumulates 60 points (20 votes * 3 points), Candidate B, C, D, E accumulates 52, 53, 51 and 30 points respectively. Therefore, under the Borda Count, Candidate A emerges as the winner with the highest total points of 60. However, notice there is a solid coalition of B and C formed by 21 voters (they are preferred by 21 voters, more than half of voters, than any other candidates). IRV will select one of the candidates of group B or C. Again, A is the last preference by the majority of voters (21 voters), yet still wins the election, showing an anti-plurality scenario.

A.3 Why Ballot Modification?

Without loss of generality, imagine that there are n voters that have provided their top-3 (ℓ) preferences over m candidates, out of which only 3 candidates satisfy priority orders (these candidates are A, B, C). Voters preferences are as follows: $n/4$ voters have provided the top-3 choices as $A > B >$ someone other than C . Another $n/4$ voters have provided the top-3 choices as $B > A >$ and someone other than C . The remaining $n/2$ voters prefer other candidates in the top-2 positions, but like C in the third position. Any aggregation mechanism based on filtering will choose C as the winner, but clearly A or B are better choices (as they are in the higher ranked order for half of the voters). Ballot modification, on the other hand, will select either A or B . Section 5.2.1 contains a case study validating this aspect using an election dataset.

A.4 Hardness Results: MqIRV is NP-Complete

THEOREM A.1. **MqIRV** is NP-Complete, even when $\ell = 2$.

PROOF. (sketch). The proof is done by taking an instance of the known NP-hard problem 3-Exact Cover problem (3XC) to an instance of **MqIRV**. $\square$

Consider an election in which m voters need to elect $k = 1$ candidate out of n candidates. In the election, each voter casts his/her ballot for two candidate in ranked order. The final candidate is determined using the IRV process. For a given instance of the election, we define the margin as the number of ballot changes required to ensure that a specific candidate wins.

We prove that computing the margin is NP-Complete. Our proof is inspired by the NP-Hardness proof of [6]. It is straightforward that the problem is in NP since the outcome of an IRV election can be computed in polynomial time. The NP-hardness is proved by reduction from the 3-Exact Cover problem (3XC). In this problem, we are given a universal set $\{e_1, \dots, e_{3n}\}$, and $m \geq n$ subsets $S_1, \dots, S_m$ of size 3 each. We need to determine whether there are n sets whose union is the universal set.

Suppose that we are given an instance of the 3XC problem. We show how to define a related IRV margin problem and then prove that the IRV has a margin n if and only if the answer to the respective 3XC problem is Yes.

The IRV problem has $2m+3n+2$ candidates $b_1, \dots, b_m, c_1, \dots, c_m, e_1, \dots, e_{3n}$ and u, v . We must ensure that u wins the election. There are $6m + m^2 + m(m+5) + 3n(2m+10) + 2m + 11 + 2m + 13 = 2m^2 + 6mn + 15m + 30n + 24$ ballots as follows:

- For every $i \in [1..m]$, let $S_i = \{e_x, e_y, e_z\}$. There are 6 ballots that we call “cover ballots”. These ballots are two of each $[b_i, e_x]$, $[b_i, e_y]$, and $[b_i, e_z]$
- For every $i \in [1..m]$ there are m ballots $[b_i, c_i]$
- For every $i \in [1..m]$ there are $m+5$ ballots $[c_i, b_i]$
- For every $i \in [1..3n]$ there are $2m+10$ ballots $[e_i, v]$
- There are $2m+11$ ballots $[v, u]$
- There are $2m+13$ ballots $[u, v]$

Suppose that the 3XC instance has an exact cover. Let the indices of the sets in the cover be $j_1, \dots, j_n$. We change n ballots as follows. For every $i \in [1..n]$ change a ballot $[b_{j_i}, c_{j_i}]$ to $[c_{j_i}, b_{j_i}]$.

We successively eliminated all candidates who got the least number of votes, which is initially $m+5$. There are m candidates with this

number of votes: $m - n$ candidates c_x , for $x \in [1..m] \setminus \{j_1, \dots, j_n\}$, and n candidates b_x , for $x \in \{j_1, \dots, j_n\}$. As a result of eliminating the $m - n$ candidates c_x , the number of votes of the candidates b_x , for $x \in [1..m] \setminus \{j_1, \dots, j_n\}$ jumps to $2m+11$. As a result of eliminating the n candidates b_x , the number of votes of the candidates c_x , for $x \in \{j_1, \dots, j_n\}$, jumps to $2m+5$. Also, since the union of the n sets S_x , $x \in \{j_1, \dots, j_n\}$, is the universal set, the elimination of b_x in the $6n$ “cover ballots” causes the number of votes of every e_i to jump to $2m+12$.

Next, the n remaining candidates c_x , for $x \in \{j_1, \dots, j_n\}$, with $2m+5$ votes are eliminated. This does not change the vote of any other candidate. Lastly, the $m - n$ candidates b_x , for $x \in [1..m] \setminus \{j_1, \dots, j_n\}$, and v each with $2m+11$ votes are eliminated. None of the e_i is eliminated at this point because all of them have $2m+12$ votes. Then, all e_i s will be deleted, each with $2m+12$ votes, and, finally, u wins with $2m+11+2m+13 = 4m+24$ votes.

We need to prove the other direction. Namely, if the margin is n then there is an exact cover. Suppose that the outcome of the elections can be changed to be u by at most n ballot changes. Since candidate v has one more vote than each of the $3n$ candidates $e_1, \dots, e_{3n}$, we need to increase the votes of all the candidates $e_1, \dots, e_{3n}$ by at least 2 so that none of the e_i is eliminated before v is eliminated. Because if any of e_i s is eliminated before v is eliminated, then the second choice of e_i 's ballot goes to v and the votes of v increase to $4m+21$. Then all e_i and u will be eliminated, and v wins the election, and u loses. The only way to ensure that none of e_i s is eliminated before v is by eliminating some of the candidates b_j . This can be done by ballot changes that reduce the number of votes of some of the candidates b_j by 1 and increase the number of votes of the respective candidates c_j . This will cause some candidates b_j to be eliminated and thus increase the votes of the resulting elements e_i in the “cover ballots” corresponding to these candidates b_j . Since we can make only n ballot changes and since the cover ballots of any candidate b_j change the votes of only the 3 candidates from $\{e_1, \dots, e_{3n}\}$ that correspond to the set S_j , the n candidates b_j eliminated first must correspond to an exact set.

A.5 Hardness Results: DISTTo is NP-hard

THEOREM A.2. **DISTTo** is NP-hard, even when $\ell = 3$.

PROOF. First, we prove that **DISTTo** is NP-hard when instead of ballot modifications we consider ballot deletions. The proof is by reduction from the 3-Exact Cover problem (3XC) described earlier. In the 3XC problem we are given a universal set $\{v_1, \dots, v_{3n}\}$, and $m > n$ subsets $S_1, \dots, S_m$ of size 3 each. We need to determine whether there are n subsets whose union is the universal set. Given an instance of the 3XC problem, we show how to reduce it to an instance of **DISTTo**. The instance of **DISTTo** consists of $3n+1$ candidates $v_1, v_2, \dots, v_{3n+1}$, and the elimination order $\pi = v_1, v_2, \dots, v_{3n+1}$ ($\pi[1] = v_1$ is eliminated first, and $\pi[3n+1] = v_{3n+1}$ is the winner). We show that this elimination order can be achieved with n ballot deletions iff the 3XC instance has a positive answer. The polynomial number of ballots in the instance varies in size from 3 to 1 and is described below.

Ballots of size 3: There are m ballots of size 3, one per every subset S_i , $1 \leq i \leq m$. Consider a subset $S_i = \{v_x, v_y, v_z\}$. From now on, we

assume that the subset is “ordered”, that is, $1 \leq x < y < z \leq 3n$. For every such subset S_i , the ballot (v_x, v_y, v_z) is added, namely v_x is the top candidate in the ballot, v_y is the second candidate, and v_z is the bottom candidate.

Ballots of size 2: For $1 \leq x < y \leq 3n$, let c_{xy} be the sum of the following 2 numbers: (1) number of ballots of size 3 in which v_x is the top candidate and v_y is the second candidate and (2) the number of ballots of size 3 in which v_x is the second candidate and v_y is the bottom candidate (note that the index of the top candidate in this case is lower than x). Let $maxc_x = \max_{y=x+1}^{3n} \{c_{xy}\}$. For every $x < y \leq 3n$, there are $maxc_x - c_{xy}$ ballots of size 2 consisting of candidate v_x as the top candidate and v_y as the second candidate. There are also $maxc_i$ ballots consisting of candidate v_i as the top candidate and candidate v_{3n+1} as the second candidate.

The total number of size 2 ballots is bounded by $6nm - 2m$ since there are at most $(3n - 1) \cdot maxc_x$ size 2 ballots with v_x as the top candidate for $1 \leq i \leq 3n$, and $\sum_{x=1}^{3n} maxc_x \leq 2m$.

Ballots of size 1: For $1 \leq x \leq 3n$, let d_x be the total number of ballots of size 3 and size 2 in which v_x is the top candidate. Let $maxd = \max_{y=1}^{3n} \{d_y\}$. There are $maxd - d_x$ ballots of size 1 consisting only of candidate v_x as the top candidate and the only candidate. There are also $maxd - 1$ ballots consisting of only candidate v_{3n+1} as the top and only candidate. The number of ballots of size 1 is bounded by $18n^2m - 3nm$ since at most $3n$ candidates have single ballots, and for each of these candidates, there are at most $m + 6nm - 2m$ ballots of size 1, since this is an upper bound on the number of ballots of size 2 and 3 per candidate.

We prove that if there is an exact cover, then the margin is n . Suppose that the 3XC instance has an exact cover consisting of n sets. Each such set corresponds to a ballot of size 3. We call these ballots the “cover ballots”. For $1 \leq x \leq 3n$, let $b(x)$ be the unique cover ballot containing x . We prove below that after deleting the n cover ballots the IRV process will result in the elimination order $v_1, v_2, \dots, v_{3n+1}$.

By our construction, before the deletion of the cover ballots, each of the candidates $v_1, \dots, v_{3n}$ is the top candidate on the $maxd$ ballots and v_{3n+1} is the top candidate on the $maxd - 1$ ballots. Since the candidates on every ballot are ordered, v_1 must be the top candidate in ballot $b(1)$ and thus after the removal of this ballot, v_1 is the top candidate in $maxd - 1$ ballots. Also, since no candidate appears more than once in the cover ballots, after their removal, each of the candidates $v_2, \dots, v_{3n}$ is the top candidate on either $maxd - 1$ or $maxd$ ballots. Recall that ties are broken arbitrarily, and thus we can eliminate v_1 . As a result of the elimination of v_1 the top candidate in all ballots that included v_1 (and are not of size 1) is updated. By our construction, there are exactly $maxc_1$ such ballots for each of the candidates $v_2, \dots, v_{3n+1}$. After the elimination of v_1 , v_2 must be the top candidate in ballot $b(2)$ and therefore after the removal of this ballot v_2 is the top candidate in $maxc_1 + maxd - 1$ ballots. Again, no candidate can be the top candidate in less than $maxc_1 + maxd - 1$ ballots and thus v_2 can be eliminated. Continuing in the same manner, after the elimination of $v_1, \dots, v_{x-1}$, candidate v_x must be the top candidate in ballot $b(x)$ and thus after the removal of this ballot v_x is the top candidate in $\sum_{y=1}^{x-1} maxc_y + maxd - 1$ ballots and can be eliminated as dictated by the required elimination order.

In the other direction, we prove that if the margin is n then there is an exact cover. To achieve this goal, we show that any set of

ballots whose removal results in the elimination order $v_1, v_2, \dots, v_{3n+1}$ must include the candidates $v_1, v_2, \dots, v_{3n}$. We prove this by contradiction. Assume that this is not the case and that there exists a set of ballots that do not include a candidate v_x whose removal results in the required elimination order. Let v_x be the candidate with the minimum index that is not included in the deleted ballots. In this case, by our construction, when v_x is about to be eliminated, it is the top candidate of $\sum_{y=1}^{x-1} maxc_y + maxd$ ballots, while v_{3n+1} is the top candidate of $\sum_{y=1}^{x-1} maxc_y + maxd - 1$ ballots. A contradiction. Clearly, the only way to delete n ballots that include all $3n$ candidates $v_1, v_2, \dots, v_{3n}$ is by choosing ballots of size 3 that correspond to an exact cover.

Next, we extend this proof to the case of ballot modifications. We use the same ballot profile as before with only one difference: candidate v_{3n+1} has $maxd - n - 1$ ballots, that is, $n + 1$ fewer ballots than any other candidate (instead of having 1 ballot less than the others). By a similar reduction, it can be shown that in this scenario, the 3XC problem instance has an exact cover iff the optimal solution to the DISTOinstance consists of n ballot modifications where the ballots removed in these modifications include candidates $v_1, v_2, \dots, v_{3n}$ and each of the added n ballots includes candidate v_{3n+1} as the top and only candidate. $\square$

A.6 Algorithm ALGAPPRX

In this section, we discuss a highly scalable Algorithm ALGAPPRX which could be used as a subroutine inside ALGEXACT to calculate **MqIRVUB**, as well as, could serve as a standalone algorithm to solve **MqIRV**.

The basic idea of ALGAPPRX simply leverages the fact that for every possible winner $w \in W$, w must have more first choice votes (tally) than the rest of the candidates ($e \in C \setminus w$). An upper bound of ballot modification to ensure the winning of candidate w is thus the maximum difference in the first choice votes (tally) between w and each e . Finally, given W , **MqIRVUB** is the smallest (minimum) over these bounds considering $w \in W$.

Algorithm 3 Algorithm ALGAPPRX: An Improved **MqIRVUB**

```

Input:  $\mathcal{B}$ , candidate set  $C$ , winners  $W$ .
Output: MqIRVUB or margin
1: MqIRVUB = 0
2: for  $w \in W$  do
3:    $ub = 0, C' = C$ 
4:    $i = 1$ 
5:   while  $i \leq n - 1$  do
6:      $e = \arg \min_{c \in C \setminus w} t_i(c)$ 
7:      $C.remove(e)$ 
8:     Distribute  $e$ 's vote following IRV rules and update tally
       of the remaining candidates
9:      $ub = \max(ub, [t_i(e) - t_i(w)])$ 
10:  end while
11:   $C = C'$ 
12: end for
13: MqIRVUB =  $\min(\mathbf{MqIRVUB}, ub)$ 
14: Return MqIRVUB or margin

```

Algorithm 4 DISTToADDALG

Input: $\mathcal{B}, C, l, \pi = \{c_1, \dots, c_n\}$
Output: DISTToADD
1: $addone = \text{Add}(\mathcal{B}, C, \pi)$
2: DISTToADD = **Merge**($\mathcal{B}, C, \pi, addone, l$)
3: Return DISTToADD

Algorithm 3 has the pseudocode, which simulates $n - 1$ rounds of IRV run-offs for each $w \in W$. In round i , the candidate e with the smallest tally is removed from C . After that, the remaining first-choice votes of e are redistributed and the tally of the remaining candidates is updated. The current upper bound ub is updated by the difference $t_i[e] - t_i[w]$ of tally between the eliminated candidate e and w (indeed, if $t_i[e] - t_i[w]$ number of extra votes could be added to w , it will never be eliminated before e). Finally, if $|W| > 1$, Algorithm 3 runs for all $w \in W$, and the minimum of the ub 's is returned as the output of **MqIRVUB** problem.

Running example. Consider that *Mira* is the preferred winner ($w = \text{Mira}$). Initially, it has 1 ballot in its tally. The candidates *Zoey*, *Gina*, *Laura* (having ballot 0, 1 and 1 respectively) are eliminated in the first three rounds. To ensure that *Kim* with ballot 2 is eliminated, 1 ballot needs to be added to *Mira*. Similarly, to ensure that *Zoey* is eliminated next, 2 ballots must be added to *Mira*. In the last round, *Sara* will have 8 ballots in its tally and *Mira* will have 2 ballots. As a result, 6 more ballots are required for *American Psycho* to avoid elimination. Therefore, the **MqIRVUB**(*American Psycho*) = $\max(0, 0, 0, 1, 2, 6) = 6$. Using the running example in figure 9, for $W = \{\text{Mira}, \text{Zoey}\}$, **MqIRVUB**(*Mira*, *Zoey*) = 2, which is a tighter upper bound than ∞ and saves expensive DISTTo calls.

THEOREM A.3. *Algorithm 3 returns a valid upper bound on MqIRV.*

PROOF SKETCH. Each round of the algorithm calculates the difference of tally between the eliminated candidate in that round and w . Let us assume that ub is the maximum of those differences after $n - 1$ rounds. Indeed, if the tally of w increases by ub , w will be the surviving candidate after $n - 1$ rounds of elimination. Modifying a single ballot amounts to adding a new ballot and removing an existing one. This could be facilitated starting from the candidate who is eliminated first, then repeat the process for the next eliminated candidate, and so on, until ub number of ballot additions has been accounted for. Similarly, the **MqIRVUB** will be the smallest of ub 's for each candidate $w \in W$. $\square$

THEOREM A.4. *Algorithm ALGAPPRX is an approximate solution for MqIRV.*

Lemma 2. *The running time of Algorithm 3 is $O(m\ell + \min\{n^2, n + m(\log n)^2\})$.*

A.7 DISTToADDALG for DISTToADD

Algorithm DISTToADDALG (Pseudocode in Algorithm 4) takes $\mathcal{B}, C, l$ as inputs, and returns the minimum number of ballot additions to ensure π . The algorithm has two main procedures: **Add** and **Merge**. **Add** finds the number of size 1 ballots needed to ensure π . **Merge** merges multiple size 1 ballots and produces ballots up to size n . Algorithm 4 first calls Subroutine **Add** which returns $addone$. Then,

it passes $addone$ to Subroutine **Merge**, which returns the output of DISTToADDALG.

Subroutine **Add** (Algorithm 5) returns a two dimensional array $addone$. Each element $addone[c][r]$ represents the number of ballots of size one added to candidate c 's tally at round r . It repeats in $|\pi|$ rounds. In round r , it computes the tally $t_r(c)$ of candidates $c \in \pi$, as well as keeping track of the sum of ballot additions up to round $r - 1$ in $t'_{r-1}[c] = \sum_x addone[c][x]$ ($x \in 1, \dots, r - 1$). To ensure c is not eliminated in round r , $\max(0, t_r(e) + t'_{r-1}(e) - t_r(c) - t'_{r-1}(c))$ number of ballots of size one ballot additions is required for c . $addone[c][r]$ is updated based on that. Finally, when all the candidates in π are popped, $addone$ is returned.

Subroutine **Merge** (Algorithm 6) reduces the number of ballots by merging the ballots of size 1 into ballots of size at most n . The intuition behind this subroutine is as follows. A ballot of signature (c_{x_1}) corresponding to $addone[c_{x_1}][r_{y_1}]$ can be merged with a ballot (c_{x_2}) corresponding to $addone[c_{x_2}][r_{y_2}]$ into a new ballot of signature (c_{x_1}, c_{x_2}) if $\pi^{-1}[c_{x_1}] < \pi^{-1}[c_{x_2}]$ and $\pi^{-1}[c_{x_1}] \leq r_{y_2}$. Since, first (c_{x_1}, c_{x_2}) will contribute to c_{x_1} in round r_{y_1} , and then after c_{x_1} is eliminated, this ballot will contribute to c_{x_2} at round r_{y_2} . After the merge, we can reduce $addone[c_{x_1}][r_{y_1}]$ by one, keeping the value of $addone[c_{x_1}][r_{y_1}]$ the same. We can keep merging ballots this way as long as it is feasible. The size of a merged ballot $(c_{x_1}, c_{x_2}, \dots, c_{x_n})$ is at most n since $\pi^{-1}[x_1] < \pi^{-1}[x_2] < \dots < \pi^{-1}[x_n]$.

Subroutine **Merge** runs in n rounds. We maintain two variable $mergeFrom$ and $mergeTo$, initially they are 0 (line 1). In each round r , the sum of the $addone$ entries in the row corresponding to the candidate $\pi[r]$ is added to $mergeFrom$, and $mergeTo$ is set to the sum of the column $r + 1$ (line 2-4). If we merge the ballots from $mergeFrom$ with the ballots counted in $mergeTo$ then the resulting ballot will always satisfy the conditions specified above. As we are merging in n rounds, the merged ballot length will never be more than n . After merging, we reduce $mergeFrom$ by $mergeTo$, making sure $mergeFrom$ is not negative (line 5). Finally, we return $mergeFrom$.

Running example. Consider an elimination order $\pi = \{\text{Molly}, \text{Mira}, \text{Gina}, \text{Kim}, \text{Laura}, \text{Zoey}, \text{Sara}\}$. To ensure that *Laura* is not eliminated before *Kim* in round 4, we need to add 1 ballot of signature (*Laura*). Similarly, to make sure *Zoey* is the winner, 4 ballots of signature (*Zoey*) have to be added at round 6. Total ballots of size one equals $4 + 1 = 5$. We can merge (*Laura*) and (*Zoey*) to (*Laura, Zoey*). When *Laura* is eliminated this ballot counts toward *Zoey*. Hence, required ballot additions = 4.

Lemma 3. *The minimum number of ballots of size one required to be added to ensure elimination order π is $\sum_{c \in C} \sum_{r=1}^{|\pi|} addone[c][r]$.*

PROOF SKETCH. Consider a round r where e is the eliminated candidate and c is a standing candidate. To ensure c is not eliminated in round c , it must satisfy: $t_r(e) + t'_{r-1}(e) \leq t_r(c) - t'_{r-1}(c)$. For a candidate c and round r , $addone[c][r]$ is the number of ballots of size one that are required to ensure that c is not eliminated before e . Hence, $\sum_{c \in C} \sum_{r=1}^{|\pi|} addone[c][r]$ is the minimum number of ballots of size one required to ensure π . $\square$

THEOREM A.5. DISTToADDALG returns an exact solution.

PROOF SKETCH. Using Lemma 3, $addone$ (returned by Subroutine **Add**) represents all the ballots of size one required to be added

Algorithm 5 Subroutine: **Add**

Input: $\mathcal{B}, C, \pi$
Output: *addone*

```

1: addone[ $c$ ][ $x$ ] = 0,  $\forall c \in C, x \in \{1, \dots, |\pi|\}$ 
2:  $r = 1$ ,
3: while  $|\pi|.notEmpty()$  do
4:    $t_r(c)$  = determine tally of  $c$  at round  $r$ ,  $\forall c \in \pi$ 
5:    $t'_{r-1}(c) = \sum_{x=1}^{r-1} \text{addone}[c][x]$ ,  $\forall c \in C$ 
6:    $e = \pi.pop\_front()$ 
7:    $e$ 's first choice votes are redistributed according to IRV
8:   for  $c \in \{\pi - e\}$  do
9:     addone[ $c$ ][ $r$ ] =  $\max(0, t_r(e) + t'_{r-1}(e) - t_r(c) - t'_{r-1}(c))$ 
10:  end for
11:   $r = r + 1$ 
12: end while
13: Return addone

```

Algorithm 6 Subroutine: **Merge**

Input: $\mathcal{B}, C, \pi, \text{addone}$
Output: *margin*

```

1: mergeFrom = 0, mergeTo = 0
2: for  $r = 1$  to  $n$  do
3:   mergeFrom = mergeFrom +  $\sum_{i=1}^{r-1} \text{addone}[\pi[r]][i]$ 
4:   mergeTo =  $\sum_{j=r+1}^n \text{addone}[\pi[j]][r]$ 
5:   mergeFrom =  $\min(0, \text{mergeFrom} - \text{mergeTo})$ 
6: end for
7: Return mergeFrom

```

to ensure π . Next, we show that Subroutine **Merge** merges maximum number of size one ballots of *addone*. Subroutine **Merge** always produces a merged ballot such that after replacing the original ballots with the merged ballot, the resulting elimination order of the election does not change. In each round, the algorithm 6 merges the maximum number of ballots possible. Repeating this process n times produces a minimum number for *mergeFrom*. Hence, **DIST-TOADDALG** returns the optimum value of **DISTTOADD**. $\square$

Lemma 4. *The running time for **DISTTOADDALG** is $O(m\ell + n^2)$.*

Extension to ballots of bounded size. We remark that Subroutine **Merge** can be generalized also to the case of ballots of bounded size $\ell < n$. In this case, we need to optimize the way we merge ballots as it may not be beneficial to merge a ballot (c_x) corresponding to *addone*[c_x][r_y] where $\pi[x] = n$ and $r_y \ll n$ to a ballot of length $\ll \ell$ as this will block us from using this ballot in future rounds (after round r_y). One way to compute the best way to merge the ballots is by modeling this problem as a min cost flow problem where the (negative) cost rewards the merged ballots and the flow value is the total number of ballots of size 1.

A.8 Proofs

THEOREM A.6. *An optimal solution for **MqKIRV** corresponds to solving **MqIRV** optimally k times.*

PROOF. It follows from the definition of the problem that an optimal solution of **MqKIRV** can be expressed as a sum of k independent **MqIRV** instances. Therefore, solving **MqKIRV** optimally is equivalent to solving each of the k **MqIRV** instances optimally. $\square$

THEOREM A.7. *Algorithm 2 returns a valid lower bound on **DISTTO**(π).*

PROOF SKETCH. Each round of the algorithm calculates half of the difference of the first choice votes between the eliminated candidate and other standing candidates based on π . Notice that the eliminated candidate must have fewer or equal votes in its tally than any of the standing candidates. For any pair of candidates, the minimum number of ballot modifications required to ensure that the eliminated candidate has less or equal votes than the standing candidate could be achieved by reducing $lb \left\lceil \frac{t_i(c) - t_i(c')}{2} \right\rceil$ number of votes from the eliminated one and adding that to the standing one. This is true for all pairs of eliminated and standing candidates in all rounds. Therefore, the maximum of all *lb*'s serves is indeed **DISTTO**(π). $\square$

THEOREM A.8. *Algorithm **ALGAPPRX** is an approximate solution for **MqIRV**.*

PROOF SKETCH. According to Theorem A.3, **ALGAPPRX** is an upper bound of **MqIRV**. Therefore, **ALGAPPRX** also solves an instance of **MqIRV** approximately. $\square$

A.9 Running Time Proofs

Lemma 5. *The running time of Algorithm 2 is $O(n^2 + m\ell)$.*

PROOF. The running time of Algorithm 2 has two components: (i) time to calculate the tally and (ii) time to find the maximum *lb* (lines 4-6). Tally can be calculated efficiently as follows: for each candidate, maintain the number of ballots in which this candidate is the top choice, as well as a linked list of all these ballots. In every elimination round, pick a candidate that appears as a top candidate in the minimum number of ballots, and eliminate this candidate by going over its linked list and adding each ballot in the linked list to the next surviving candidate (and update this candidate's number of ballots). While finding the next surviving candidate, delete the ones that have already been eliminated from the ballot. In this way, the number of operations performed on a single ballot during the tally calculation is $O(\ell)$. Hence, the running time to calculate the tally is $O(m\ell)$. To find the maximum of *lb* in each of the n rounds (lines 4-6), the algorithm iterates over the remaining $O(n)$ candidates. This totals to $O(n^2)$ time. Therefore, the running time of Algorithm 2 is $O(n^2 + m\ell)$. $\square$

Lemma 6. *The running time of Algorithm 3 is $O(m\ell + \min\{n^2, n + m(\log n)^2\})$.*

PROOF. The running time of Algorithm 3 has two components: (i) time for calculating the tally (ii) time for finding the candidate with minimum tally. Tally can be calculated efficiently in $O(m\ell)$ time as explained in the analysis of Algorithm 2. Finding the candidate with a minimum tally can be done using two methods depending on the value of n and m . *Method 1:* Perform a linear search on all remaining candidates to find the one with the minimum tally in each round. The linear search requires $O(n)$ time per round, and thus a total $O(n^2)$ time in n rounds. *Method 2:* The candidate with the minimum tally can be found using a min heap to store the tally of the remaining candidates. The creation of the heap takes $O(n)$ time.

Finding the initial candidate with the smallest tally takes constant time. A single update of the heap takes $O(\log n)$ time. The number of times that heap needs to be updated is bounded by the number of ballots that need to be redistributed when a candidate is eliminated. Since we eliminate the candidate with the minimum tally, if a round has x surviving candidates, then the minimum tally is no more than m/x . So, summing over all elimination rounds, we get that the number of heap updates is upper bounded by $m(1/n + 1/(n-1) + \dots + 1/2)$ which is $O(m \log n)$ (Harmonic number). Therefore, the total time to update the heap is $O(m(\log n)^2)$. and the running time for Algorithm 3 is $O(m\ell + \min\{n^2, n + m(\log n)^2\})$. $\square$

Lemma 7. *The running time for DISTTOADDALG is $O(m\ell + n^2)$.*

PROOF. (a) **Add:** the runtime for counting tally is $O(m\ell)$, and for calculating *addone* is $O(n^2)$. (b) **Merge:** each cell of *addone* is visited a constant number of times, and hence takes $O(n^2)$ time. It follows that the total running time for DISTTOADDALG is $O(m\ell + n^2)$. $\square$

A.10 Prior Work

We present three types of related work in this section.

Preference aggregation. Preference aggregation is closely studied in the context of group recommendation [1, 2, 5, 7, 12, 25, 30, 32, 33], with the goal of selecting one or top- k items that are most suitable to the preference of all users in the group. These are also studied while promoting fairness in ranking and recommendation [29, 37]. In [14], the authors empirically demonstrate that multi-stage voting methods, such as STV and IRV offer benefits over positional preference aggregation methods (e.g., plurality voting, approval voting) in the recommendation contexts by handling hyperactive users in a more equitable and fair way.

Changing original preferences. The second line of related work exists in how to minimally update the original preferences of the users so that the produced outputs satisfy additional criteria. Some leading criteria include maximizing the satisfaction of some specific users considering rating based preference aggregation methods in top- k recommendation [33], changing the original winner, that is, computing margin, or producing Margin of victory (MoV), or satisfying fairness criteria, [23, 35], to name a few. Among these, the most relevant to this work is the previous work on computing MoV. There are two types of MOV: constructive and destructive. In the constructive (destructive) version, the goal is to find the minimum number of changes to the ballots that is needed so that a special candidate is (not) elected. [36] has investigated the computational

complexity and (in)approximability of computing MoV for various voting rules, including approval voting, all positional scoring rules, etc. [9] has introduced a sampling based probabilistic algorithm for finding the margin of victory, which can be used for many voting rules.

Multi-stage preference aggregation methods and their margin of victory computation. Multi-stage methods, such as STV and IRV, were introduced in the 19th century in electoral voting systems. [6] demonstrated that determining whether the MoV in an IRV election is at most 1 is NP-hard for both constructive and destructive versions. Moreover, there is no 2-approximation algorithm for it unless $P = NP$. In [15], the coalitional weighted manipulation is investigated. In [24], the authors have shown a branch and bound algorithm that calculates possible winners when only some part of the ballots are accessible, not all. The usage of [24] is to generate information on the result of an election and to announce it on election night, when there are still some ballots that have not arrived at the specified place to count the votes. MoV of IRV [27] and STV [21] is studied in many related works [26], [11], [4], [10].

A.11 Figures




Figure 9: Partially explored tree for ALGEXACT, the candidates are represented with their ids, where red nodes and their subtrees are pruned