



HAL
open science

Rank Aggregation Using Scoring Rules

Niclas Boehmer, Robert Bredereck, Dominik Peters

► **To cite this version:**

Niclas Boehmer, Robert Bredereck, Dominik Peters. Rank Aggregation Using Scoring Rules. Proceedings of the 37th AAAI Conference on Artificial Intelligence, Feb 2023, Washington, D. C., United States. pp.5515-5523, 10.1609/aaai.v37i5.25685 . hal-04546555

HAL Id: hal-04546555

<https://hal.science/hal-04546555>

Submitted on 15 Apr 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Rank Aggregation Using Scoring Rules

Niclas Boehmer¹, Robert Bredereck², and Dominik Peters³

¹TU Berlin

¹TU Clausthal

³CNRS, LAMSADE, Université Paris Dauphine–PSL

Contents

1. Introduction	2
1.1. Our Contributions	3
2. Preliminaries	4
3. Scoring-Based Rank Aggregation	4
4. Axiomatic Properties	5
5. Simulations	6
5.1. Setup	6
5.2. Comparison of Scoring-Based Ranking Methods	7
5.3. Comparison to Kemeny Ranking	8
5.4. Further Simulations	8
6. Complexity	9
6.1. Parameter Number of Candidates	10
6.2. Sequential Loser	10
6.2.1. Plurality	10
6.2.2. Veto	12
6.2.3. Borda	13
6.3. Sequential Winner	13
7. Future Directions	13
8. Bibliography	14
A. Additional Material for Section 4	17
B. Additional Material for Section 5	20
B.1. Setup—Euclidean Model	20
B.2. Comparison of Ranking Methods—Euclidean Model	20
B.3. Similarity in Different Vote Parts	21
B.4. Number of Ties in Executions of the Rules	22
B.5. Influence of Profile Size	23
B.6. Further Voting Rules	25
C. Additional Material for Section 6.2	26
C.1. Plurality	26
C.2. Veto	28
C.3. Borda	35
D. Additional Material for Section 6.3	39
D.1. Plurality	39
D.2. Veto	42
D.3. Borda	43

Abstract

To aggregate rankings into a social ranking, one can use scoring systems such as Plurality, Veto, and Borda. We distinguish three types of methods: ranking by score, ranking by repeatedly choosing a winner that we delete and rank at the top, and ranking by repeatedly choosing a loser that we delete and rank at the bottom. The latter method captures the frequently studied voting rules Single Transferable Vote (aka Instant Runoff Voting), Coombs, and Baldwin. In an experimental analysis, we show that the three types of methods produce different rankings in practice. We also provide evidence that sequentially selecting winners is most suitable to detect the “true” ranking of candidates. For different rules in our classes, we then study the (parameterized) computational complexity of deciding in which positions a given candidate can appear in the chosen ranking. As part of our analysis, we also consider the WINNER DETERMINATION problem for STV, Coombs, and Baldwin and determine their complexity when there are few voters or candidates.

1. Introduction

Rank aggregation, the task of aggregating several rankings into a single ranking, sits at the foundation of social choice as introduced by Arrow [1951]. Besides preference aggregation, it has numerous important applications, for example in the context of meta-search engines [Dwork et al., 2001], of juries ranking competitors in sports tournaments [Truchon, 1998], and multi-criteria decision analysis.

One of the best-known methods for aggregating rankings is Kemeny’s [1959] method: A Kemeny ranking is a ranking that minimizes the average swap distance (Kendall-tau distance) to the input rankings. It is axiomatically attractive [Young and Levenglick, 1978, Can and Storcken, 2013, Bossert and Sprumont, 2014] and has an interpretation as a maximum likelihood estimator [Young, 1995] making it well-suited to epistemic social choice that assumes a ground truth.

However, Kemeny’s method is hard to compute [Bartholdi et al., 1989, Hemaspaandra et al., 2005] which makes the method problematic to use, especially when there are many candidates to rank (for example, when ranking all applicants to a university). Even if computing the ranking is possible, it is coNP-hard to verify if a ranking is indeed a Kemeny ranking [Fitzsimmons and Hemaspaandra, 2021]. Thus, third parties cannot easily audit, interpret, or understand the outcome, making systems based on Kemeny’s method potentially unaccountable. This limits its applicability in democratic contexts.

These two drawbacks motivate the search for computationally simpler and more transparent methods for aggregating rankings. There is a significant literature on polynomial-time approximation algorithms for Kemeny’s method [Coppersmith et al., 2006, Kenyon-Mathieu and Schudy, 2007, Ailon et al., 2008, van Zuylen and Williamson, 2009], but these algorithms are typically not attractive beyond their approximation guarantee. In particular, they would typically not fare well in an axiomatic analysis, and are unlikely to be understood by and appealing to the general public (many are based on derandomization).

Instead, we turn to one of the fundamental tools of social choice: positional scoring rules. These rules transform voter rankings into scores for the candidates. For example, under the *Plurality* scoring rule, every voter gives 1 point to their top-ranked candidate. Under the *Veto* (or anti-plurality) scoring rule, voters give -1 point to their last-ranked candidate and zero points to all others. Under the *Borda* scoring rule, every voter gives m points to their top-ranked candidate, $m - 1$ points to their second-ranked candidate, and so on, giving 1 point to their last-ranked candidate. We study three ways of using scoring rules to aggregate rankings:

- *Score*: We rank the candidates in order of their score, higher-scoring candidates being

ranked higher.

- *Sequential-Winner*: We take the candidate c with the highest score and rank it top in the aggregate ranking. We then delete c from the input profile, re-calculate the scores, and put the new candidate with the highest score in the second position, and so on.
- *Sequential-Loser*: We take the candidate c with the lowest score and rank it last. We then delete c , re-calculate the scores, and put the new candidate with the lowest score in the second-to-last position, and so on.

Ranking by score is the obvious way of using scoring rules for rankings, and so it has been studied in the social choice literature [Smith, 1973, Levenshlick, 1977]. Sequential-Loser captures as special cases the previously studied rules Single Transferable Vote (also known as Instant Runoff Voting, among other names), Coomb’s method, and Baldwin’s method. These are typically used as voting rules that elect a single candidate, but they can also be understood as rank aggregation methods. On the other hand, despite being quite natural, Sequential-Winner methods appear not to have been formally studied in the literature (to our knowledge).

1.1. Our Contributions

Axiomatic Properties (Section 4) Based on the existing literature, we begin by describing some axiomatic properties of the methods in our three families. For example, we check which of the methods are Condorcet or majority consistent, and which are resistant to cloning. We also consider independence properties and state some characterization results.

Simulations (Section 5) To understand how and whether the three families of methods practically differ from each other, and how they relate to Kemeny’s method, we perform extensive simulations based on synthetic data (sampled using the Mallows and Euclidean models). We find that, for Plurality and Borda, ranking by score and Sequential-Loser usually produce very similar results, whereas Sequential-Winner offers a new perspective (which is typically closer to Kemeny’s method). Moreover, we observe that Sequential-Loser rules seem to be particularly well suited to identify the best candidates (justifying their usage as single-winner voting rules), while Sequential-Winner rules are best at avoiding low quality candidates.

Computational Complexity (Section 6) The rules in all three of our families are easy to compute in the sense that their description implies a straightforward algorithm for obtaining an output ranking. However, for the sequential rules there is a subtlety: During the execution of the rule, ties can occur. It matters how these are broken, because candidates could end up in significantly different positions. For high-stakes decisions and in democratic contexts, it would be important to know which output rankings are possible.

Thus, we study the computational problem of deciding whether a given candidate can end up in a given position. This and related problems have been studied in the literature under the name of *parallel universe tie-breaking*, including theoretical and experimental studies for some of the rules in our families [Conitzer et al., 2009, Brill and Fischer, 2012, Mattei et al., 2014, Freeman et al., 2015, Wang et al., 2019]. We extend the results of that literature and find NP-hardness for all the sequential methods that we study. We show that the problem becomes tractable if the number of candidates is small. In contrast, for several methods we find that the problem remains hard even if the number of input rankings is small. Curiously, for few input rankings, methods based on Plurality, Borda, or Veto each induce a different parameterized complexity class.

2. Preliminaries

For $k \in \mathbb{N}$, write $[k] = \{1, \dots, k\}$.

Let $C = \{c_1, \dots, c_m\}$ be a set of m candidates. A ranking \succ of C is a linear order (irreflexive, total, transitive) of C . We write $\mathcal{L}(C)$ for the set of all rankings of C .

A (ranking) profile $P = (\succ_1, \dots, \succ_n)$ is a list of rankings. We sometimes say that the rankings are voters.

For a subset $C' \subseteq C$ of candidates and ranking $\succ \in \mathcal{L}(C)$, we write $\succ|_{C'}$ for the ranking obtained by restricting \succ to the set C' . For a profile P , we write $P|_{C'}$ for the profile obtained by restricting each of its rankings to C' .

A social preference function¹ f is a function that assigns to every ranking profile P a non-empty set $f(P) \subseteq \mathcal{L}(C)$ of rankings. Here, $f(P)$ may be a singleton but there can be more than one output ranking in case of ties. For a ranking \succ , we say that f selects \succ on P if $\succ \in f(P)$.

For a ranking $\succ \in \mathcal{L}(C)$ and a candidate $c \in C$, let $\text{pos}(\succ, c) = |\{d \in C : d \succ c\}| + 1$ be the position of c in \succ . For example, if $\text{pos}(\succ, c) = 1$ then c is the most-preferred candidate in \succ . We write $\text{cand}(\succ, r) \in C$ for the candidate ranked in position $r \in [m]$ in $\succ \in \mathcal{L}(C)$.

For a ranking $\succ \in \mathcal{L}(C)$, $\text{rev}(\succ)$ denotes the ranking where the candidates are ranked in the opposite order as in \succ , i.e., for each $r \in [m]$, $\text{cand}(\succ, r) = \text{cand}(\text{rev}(\succ), m - r + 1)$. For a profile $P = (\succ_1, \dots, \succ_n)$, we write $\text{rev}(P) = (\text{rev}(\succ_1), \dots, \text{rev}(\succ_n))$.

For an integer $m \in \mathbb{N}$, a scoring vector $\mathbf{s}^{(m)} = (s_1, \dots, s_m) \in \mathbb{R}^m$ is a list of m numbers. A scoring system is a family of scoring vectors $(\mathbf{s}^{(m)})_{m \in \mathbb{N}}$ one for each possible number m of candidates. For the sake of conciseness, we sometimes write \mathbf{s} instead of $(\mathbf{s}^{(m)})_{m \in \mathbb{N}}$. We will mainly focus on three scoring systems:

- *Plurality* with $\mathbf{s}^{(m)} = (1, 0, \dots, 0)$ for each $m \in \mathbb{N}$,
- *Veto* with $\mathbf{s}^{(m)} = (0, \dots, 0, -1)$ for each $m \in \mathbb{N}$,
- *Borda* with $\mathbf{s}^{(m)} = (m, m - 1, \dots, 1)$ for each $m \in \mathbb{N}$.

Given a profile P over m candidates, the \mathbf{s} -score of candidate $c \in C$ is $\text{score}_{\mathbf{s}}(P, c) = \sum_{i \in [n]} \mathbf{s}_{\text{pos}(\succ_i, c)}^{(m)}$. We say that a candidate is an \mathbf{s} -winner if it has maximum \mathbf{s} -score, and an \mathbf{s} -loser if it has minimum \mathbf{s} -score. For a scoring system \mathbf{s} we denote by \mathbf{s}^* the scoring system where we reverse each scoring vector and multiply all its entries by -1 , i.e., for each $m \in \mathbb{N}$ and $i \in [m]$, we have $(\mathbf{s}^*)^{(m)}_i = -\mathbf{s}_{m-i+1}^{(m)}$. Note that $(\mathbf{s}^*)^* = \mathbf{s}$ for every \mathbf{s} , that *Plurality*^{*} = *Veto*, that *Veto*^{*} = *Plurality*, and that *Borda*^{*} is the same as *Borda*, up to a shift.

For two rankings \succ_1 and \succ_2 , their swap distance (or Kendall-tau distance) $\kappa(\succ_1, \succ_2)$ is the number of pairs of candidates on whose ordering the two rankings disagree, i.e., $\kappa(\succ_1, \succ_2) = |\{(c, d) \in C \times C : c \succ_1 d \text{ and } d \succ_2 c\}|$. Note that the maximum swap distance between two rankings is $\binom{m}{2}$. Given a profile P , *Kemeny's rule* selects those rankings which minimize the average swap distance to the rankings in P , so it selects $\arg \min_{\succ \in \mathcal{L}(C)} \sum_{i \in [n]} \kappa(\succ, \succ_i)$. We refer to the selected rankings as *Kemeny rankings*.

3. Scoring-Based Rank Aggregation

We now formally define the three families of scoring-based social preference functions that we study.

Definition 3.1 (s-Score). Let \mathbf{s} be a scoring system. For the social preference function \mathbf{s} -Score on profile P , we have $\succ \in \mathbf{s}\text{-Score}(P)$ if and only if for all $c, d \in C$ with $\text{score}_{\mathbf{s}}(c, P) > \text{score}_{\mathbf{s}}(d, P)$, we have $c \succ d$.

¹This terminology is due to [Young and Levenglick \[1978\]](#). The term *social welfare function* from [Arrow \[1951\]](#) usually refers to resolute functions that may only output a single ranking.

Definition 3.2 (Sequential- \mathbf{s} -Winner; Seq.- \mathbf{s} -Winner). Let \mathbf{s} be a scoring system. The social preference function Seq.- \mathbf{s} -Winner is defined recursively as follows: For a profile P , we have $\succ \in \text{Seq.-}\mathbf{s}\text{-Winner}(P)$ if and only if

- the top choice $c = \text{cand}(\succ, 1)$ is an \mathbf{s} -winner in P ,
- if $|C| > 1$, then $\succ|_{C \setminus \{c\}} \in \text{Seq.-}\mathbf{s}\text{-Winner}(P|_{C \setminus \{c\}})$.

Definition 3.3 (Sequential- \mathbf{s} -Loser; Seq.- \mathbf{s} -Loser). Let \mathbf{s} be a scoring system. The social preference function Seq.- \mathbf{s} -Loser is defined recursively as follows: For a profile P , we have $\succ \in \text{Seq.-}\mathbf{s}\text{-Loser}(P)$ if and only if

- the bottom choice $c = \text{cand}(\succ, |C|)$ is an \mathbf{s} -loser in P ,
- if $|C| > 1$, then $\succ|_{C \setminus \{c\}} \in \text{Seq.-}\mathbf{s}\text{-Loser}(P|_{C \setminus \{c\}})$.

Example 3.4. Let P be the following ranking profile:

$$3 \times a \succ b \succ c, \quad 2 \times b \succ c \succ a, \quad 2 \times c \succ b \succ a$$

Then for the three methods based on Plurality, we have:

- $\text{Plurality-Score}(P) = \{a \succ b \succ c, a \succ c \succ b\}$,
- $\text{Seq.-Plurality-Wi.}(P) = \{a \succ b \succ c\}$, and
- $\text{Seq.-Plurality-Lo.}(P) = \{b \succ a \succ c, c \succ a \succ b\}$.

We sometimes view Seq.- \mathbf{s} -Winner (or Seq.- \mathbf{s} -Loser) rules as round-based voting rules, where in each round an \mathbf{s} -winner (or an \mathbf{s} -loser) is deleted from the profile and added in the highest (or lowest) position of the ranking that has not yet been filled. Notably, if there are multiple \mathbf{s} -winners (or \mathbf{s} -losers) in one round, each selection gives rise to different output rankings. Seq.-Plurality-Loser is also known as *STV*, Seq.-Veto-Loser as *Coombs*, and Seq.-Borda-Loser as *Baldwin*.

Sequential-Winner and Sequential-Loser rules are formally closely related: If a candidate is an \mathbf{s} -winner in some profile P , then it is an \mathbf{s}^* -loser in the reverse profile $\text{rev}(P)$. Hence, we can conclude the following:

Lemma 3.5. Let \mathbf{s} be a scoring system. Then for each ranking profile P and for every ranking $\succ \in \mathcal{L}(C)$, we have:

$$\begin{aligned} & \succ \in \text{Sequential-}\mathbf{s}\text{-Winner}(P) \\ \iff & \text{rev}(\succ) \in \text{Sequential-}\mathbf{s}^*\text{-Loser}(\text{rev}(P)). \end{aligned}$$

For example, this lemma establishes a close connection between Seq.-Veto-Winner and Seq.-Plurality-Loser, as a ranking \succ is selected under Seq.-Veto-Winner on profile P if and only if $\text{rev}(\succ)$ is selected under Seq.-Plurality-Loser on profile $\text{rev}(P)$. This equivalence will prove useful in our axiomatic analysis and in our complexity results.

4. Axiomatic Properties

In this section, we will briefly and informally discuss some axiomatic properties and characterizations of the methods in our three families. A more formal treatment appears in [Appendix A](#). See [Table 1](#) for an overview.

A desirable property of a ranking aggregation rule is that if one candidate is deleted from the profile, then the relative rankings of the other candidates does not change (*independence*

	Kemeny	Score			Sequential-Winner			Sequential-Loser		
		Plurality	Veto	Borda	Plurality	Veto	Borda	Plurality	Veto	Borda
Independence at the top	✓	×	×	×	✓	✓	✓	×	×	×
Independence at the bottom	✓	×	×	×	×	×	×	✓	✓	✓
Reinforcement	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Reinforcement at the top	×	✓	✓	✓	✓	✓	✓	×	×	×
Reinforcement at the bottom	×	✓	✓	✓	×	×	×	✓	✓	✓
Condorcet winner at top	✓	×	×	×	×	×	×	×	×	✓
Copy majority	✓	×	×	×	✓	×	×	×	✓	×
Independence of clones	×	×	×	×	×	×	×	✓	×	×

Table 1: An overview of the axiomatic properties of our studied rules. See [Appendix A](#) for definitions.

of irrelevant alternatives, IIA). [Arrow’s \[1951\]](#) impossibility theorem shows that this property cannot be satisfied by unanimous non-dictatorial rules. [Young \[1988\]](#) proves that Kemeny’s method satisfies a weaker version that he calls *local IIA*: removing the candidate that appears in the first or last position in the Kemeny ranking does not change the ranking of the other candidates. Splitting this property into its two parts, we can easily see from their definitions that Seq.-s-Winner satisfies independence at the top, and Seq.-s-Loser satisfies independence at the bottom.

Another influential axiom is known as *consistency* or *reinforcement*. A rule f satisfies reinforcement if whenever some ranking \succ is chosen in two profiles, $\succ \in f(P) \cap f(P')$, then it is also chosen if we combine the profiles into one, and in fact $f(P + P') = f(P) \cap f(P')$. All the methods in this paper satisfy reinforcement. Notably, [Young \[1988\]](#) shows that Kemeny is the only anonymous, neutral, and unanimous rule satisfying reinforcement and local IIA. Focusing on Seq.-s-Loser, [Freeman et al. \[2014\]](#) define reinforcement at the bottom to mean that if the same candidate c is placed in the last position in the selected ranking in two profiles, then c is also placed in the last position in the selected ranking in the combined profile. They show that independence at the bottom and reinforcement at the bottom characterize Seq.-s-Loser rules (under mild additional assumption). Using [Lemma 3.5](#), a simple adaptation of their proof shows that Seq.-s-Winner rules can be similarly characterized by independence at the top and reinforcement at the top. (s-Score methods do not satisfy similar independence assumptions; they have been characterized by [Levenglick \[1977\]](#) and [Smith \[1973\]](#).)

Refining their characterization of Seq.-s-Loser rules, [Freeman et al. \[2014\]](#) characterize Seq.-Plurality-Loser (aka STV) as the only Seq.-s-Loser rule satisfying independence of clones [[Tideman, 1987](#)], Seq.-Veto-Loser (aka Coombs) as the only one that, in case a strict majority of voters have the same ranking, copies that ranking as the output ranking, and Seq.-Borda-Loser (aka Baldwin) as the only one always placing a Condorcet winner in the first position. Using [Lemma 3.5](#), we can similarly characterize Seq.-Plurality-Winner as the only method in its class that copies a majority ranking.

5. Simulations

We analyze our three families of scoring-based ranking rules for Plurality and Borda on synthetically generated profiles.

5.1. Setup

To deal with ties in the computation of our rules, each time we sample a ranking profile over candidates C , we also sample a ranking $\succ_{\text{tie}} \in \mathcal{L}(C)$ uniformly at random and break ties according

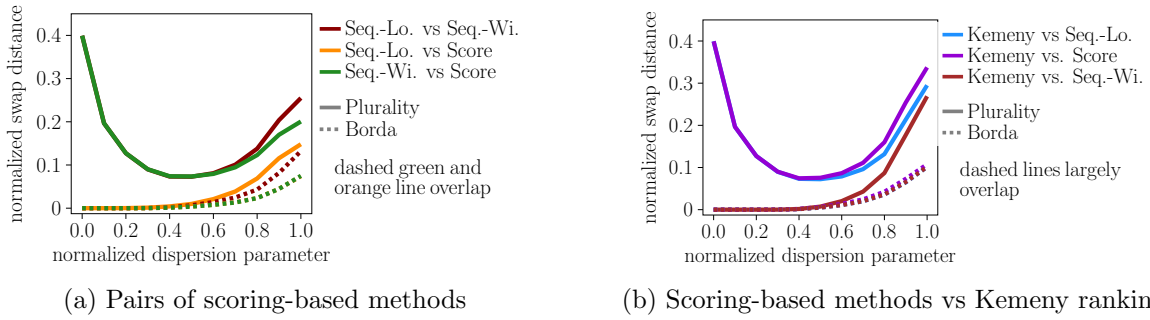


Figure 1: Pairwise average normalized swap distance between rankings produced by different methods for Plurality (solid) and Borda (dashed) on Mallows profiles with 10 candidates and 100 voters.

to \succ_{tie} for all rules. To quantify the difference between two rankings $\succ_1, \succ_2 \in \mathcal{L}(C)$, we use their normalized swap distance, i.e., their swap distance $\kappa(\succ_1, \succ_2)$ divided by the maximum possible swap distance between two rankings $\binom{m}{2}$.

(Normalized) Mallows We conduct simulations on profiles generated using the Mallows model [Mallows \[1957\]](#) (as observed by [Boehmer et al. \[2021\]](#) real-world profiles seem often to be close to some Mallows profiles). This model is parameterized by a dispersion parameter $\phi \in [0, 1]$ and a central ranking $\succ^* \in \mathcal{L}(C)$. Then, a profile is assembled by sampling rankings i.i.d. so that the probability of sampling a ranking $\succ \in \mathcal{L}(C)$ is proportional to $\phi^{\kappa(\succ, \succ^*)}$. We use the normalization of the Mallows model proposed by [Boehmer et al. \[2021\]](#), which is parameterized by a normalized dispersion parameter $\text{norm-}\phi \in [0, 1]$. This parameter is then internally converted to a dispersion parameter ϕ such that the expected swap distance between a sampled vote and the central vote is $\text{norm-}\phi \cdot (m(m-1)/4)$. Then $\text{norm-}\phi = 0$ results in profiles only containing the central vote, and $\text{norm-}\phi = 1$ leads to profiles where all rankings are sampled with the same probability, so that on average rankings disagree with the central ranking \succ^* on half of the pairwise comparisons. Choosing $\text{norm-}\phi = 0.5$ leads to profiles where rankings on average disagree with \succ^* on a quarter of the pairwise comparisons.

5.2. Comparison of Scoring-Based Ranking Methods

We analyze the average normalized swap distance between the rankings selected by our three families of scoring-based ranking methods on profiles containing 100 rankings over 10 candidates. For this, we sampled 10 000 profiles for each $\text{norm-}\phi \in \{0, 0.1, \dots, 0.9, 1\}$ and depict the results in [Figure 1\(a\)](#). Let us first focus on Plurality: We find that the rankings produced by Seq.-Plurality-Loser and Plurality-Score are quite similar, whereas the ranking produced by Seq.-Plurality-Winner is substantially different. This observation is particularly strong for $\text{norm-}\phi \leq 0.3$: In such profiles, all the rankings are similar to each other. Accordingly, many candidates initially have a Plurality score of zero, and thus there are many ties in the execution of Plurality-Score and Seq.-Plurality-Loser (for the latter, ties occur in more than half of the rounds). Thus, the rankings computed by the two rules fundamentally depend on the (shared) random tie-breaking order \succ_{tie} . In contrast, for Seq.-Plurality-Winner, for $\text{norm-}\phi \leq 0.3$, no ties in its execution appear. In particular, Seq.-Plurality-Winner is thereby able to meaningfully distinguish the weaker candidates on these profiles.

Turning to $\text{norm-}\phi \geq 0.3$ (where more candidates have non-zero Plurality score and thus the tie-breaking is no longer as important), Seq.-Plurality-Loser and Plurality-Score are still clearly more similar to each other than to Seq.-Plurality-Winner; this indicates that Seq.-Wi. rules

indeed add a new perspective to existing scoring-based ranking rules.

Switching to Borda, the rankings returned by the three methods are quite similar. This is intuitive given that Borda scores capture the general strength of candidates in a profile much better than Plurality scores. Thus, the Borda score of a candidate also changes less drastically in case some candidate is deleted. Increasing $\text{norm-}\phi$, the selected rankings become more different from each other (as profiles get more chaotic, leading to more similar Borda scores of candidates). Interestingly, for larger values of $\text{norm-}\phi$, Borda-Score has the same (small) distance to the other two rules, whereas Seq.-Borda-Winner and Seq.-Borda-Loser are more different.

5.3. Comparison to Kemeny Ranking

To assess which method produces the “most accurate” rankings, we compare them to Kemeny’s method. For 10 000 profiles for each $\text{norm-}\phi \in \{0, 0.1, \dots, 0.9, 1\}$, in [Figure 1\(b\)](#), we show the average normalized swap distance of the Kemeny ranking to the rankings selected by our rules.

For Plurality, independently of the value of $\text{norm-}\phi$, Seq.-Plurality-Winner produces the ranking most similar to the Kemeny ranking, then Seq.-Plurality-Loser and lastly Plurality-Score, indicating the advantages of sequential rules. What sticks out is that for $\text{norm-}\phi \leq 0.3$, Seq.-Plurality-Loser and Plurality-Score are far away from the Kemeny ranking. As discussed above, the reason is that, for both methods, large parts of the ranking are simply determined by the random tie-breaking order in such profiles. In contrast, Seq.-Plurality-Winner is not affected, and its output ranking is very close to the Kemeny ranking until $\text{norm-}\phi \leq 0.5$ (when their average normalized distance is only 0.004). For a larger dispersion parameter and in particular for $\text{norm-}\phi \geq 0.7$, the distance from the Kemeny ranking become more similar for our three methods. This behavior is intuitive, recalling that for $\text{norm-}\phi = 1$, profiles are “chaotic”, with many different rankings having comparable quality.

For Borda, the rankings produced by the three methods are all around the same (small) distance from the Kemeny ranking. This distance increases steadily from 0 for $\text{norm-}\phi = 0$ to around 0.1 for $\text{norm-}\phi = 1$.

5.4. Further Simulations

In [Appendix B](#), we describe the results of further experiments. For instance, we analyze in which parts of the computed ranking the considered methods agree or disagree most. We find that for both Plurality and Borda, for the top positions the Kemeny ranking agrees frequently with the Seq.-Loser rule. For the bottom positions it agrees with the Seq.-Winner rule. This suggests that one should use Seq.-Loser for identifying the best candidates and Seq.-Winner for avoiding the worst candidates. Moreover, Seq.-Winner and Score agree more commonly on the top half of candidates, whereas Seq.-Loser and Score agree more commonly on the bottom half of candidates.

We repeat all our experiments on profiles sampled from Euclidean models. Obtaining similar results, this confirms that our general observations from above also hold for profiles sampled from other distributions. We also analyze the influence of the number of voters and candidates on the results, observing that increasing the number of voters leads to an increased similarity of the rankings for Plurality and Borda, whereas increasing the number of candidates leads to an increased similarity for Borda but not for Plurality. We also consider additional scoring vectors. For instance, we find that for Veto the roles of Seq.-Loser and Seq.-Winner are reversed, which is to be expected, recalling [Lemma 3.5](#).

		n	m
Sequential-Plurality-Loser (STV)	NP-c. (Thm. 6.2)	FPT (Obs. 6.4)	FPT (Thm. 6.1)
Sequential-Veto-Loser (Coombs)	NP-c. (Thm. 6.6)	W[1]-h. (Thm. 6.6), XP (Thm. 6.7)	FPT (Thm. 6.1)
Sequential-Borda-Loser (Baldwin)	NP-c. (Thm. 6.8)	NP-c. for $n = 8$ (Thm. 6.8)	FPT (Thm. 6.1)

Table 2: Our results for Sequential-Loser rules. All hardness results hold for WINNER DETERMINATION; all algorithmic results also apply to POSITION- k DETERMINATION. The unparameterized NP-hardness results in the first column were already stated or proven by Conitzer et al. [2009] and Mattei et al. [2014]

6. Complexity

We study various computational problems related to Sequential-Winner and Sequential-Loser rules. By breaking ties arbitrarily, it is easy to compute *some* ranking that is selected by such a rule. However, in some (high-stakes) applications, it might not be sufficient to simply output some ranking selected by the rule. For instance, some candidate could claim that there also exist other rankings selected by the same rule where that candidate is ranked higher. To check such claims, and understand which rankings can be selected in the presence of ties, we need an algorithm that for a given candidate d and position k , decides whether d is ranked in position k in some ranking selected by the rule. Accordingly, we introduce the following computational problem:

POSITION- k DETERMINATION for social preference function f	
Given:	A ranking profile P over candidate set C , a designated candidate $d \in C$, and an integer $k \in [C]$.
Question:	Is there a ranking \succ selected by f on P where d is in position k , i.e., $\succ \in f(P)$ with $\text{pos}(\succ, d) = k$?

Where possible, we will design (parameterized) algorithms that solve this problem. We also prove hardness results, which will apply even to restricted versions of this problem that are most relevant in practice. Specifically, we would expect candidates to mainly be interested if they can be ranked highly. Thus, we introduce the TOP- k DETERMINATION problem, where we ask whether a given candidate can be ranked in one of the first k positions.² Lastly, the special case of both problems with $k = 1$ is of particular importance: The WINNER DETERMINATION problem asks whether the designated candidate is ranked in the first position in some ranking selected by the rule.

For the three Sequential-Loser rules, it is known that their WINNER DETERMINATION problem is NP-complete. For STV, this was stated by Conitzer et al. [2009], and for Baldwin and Coombs, this was proven by Mattei et al. [2014]. We will see that the corresponding TOP- k DETERMINATION problems for the Sequential-Winner rules are also NP-complete. Thus, since almost all of our problems turn out to be NP-hard, we take a more fine-grained view. In particular, we will study the influence of the number n of voters and the number m of candidates on the complexity of our problems. This analysis is not only of theoretical interest but also practically relevant, as in many applications one of the two parameters is considerably smaller than the other (e.g., in political elections m is typically much smaller than n , while in applications such as meta-search engines or ranking applicants, n is often much smaller than m). Tables 2 and 3 provide overviews of our results.

²If we have an algorithm for POSITION- k DETERMINATION, we can solve the TOP- k DETERMINATION problem by using the algorithm for positions $i = 1, \dots, k$. (This is a Turing reduction.)

6.1. Parameter Number of Candidates

We start by considering the parameter m , the number of candidates. It is easy to see that POSITION- k DETERMINATION for all Sequential-Winner and Sequential-Loser rules is fixed-parameter tractable with respect to m (by iterating over all $m!$ possible output rankings). However, it is possible to improve the dependence on the parameter in the running time.

Theorem 6.1. *For every scoring system \mathbf{s} , POSITION- k DETERMINATION can be solved in*

[Proof]

- $\mathcal{O}(2^m \cdot nm^2)$ time and $\mathcal{O}(m^k \cdot nm^2)$ time for Sequential- \mathbf{s} -Winner, and
- $\mathcal{O}(2^m \cdot nm^2)$ time and $\mathcal{O}(m^{m-k} \cdot nm^2)$ time for Sequential- \mathbf{s} -Loser.

Proof (algorithm). We present an algorithm for Seq.- \mathbf{s} -Winner (the results for Seq.- \mathbf{s} -Loser directly follow from this by applying Lemma 3.5). We solve the problem via dynamic programming. We call a subset $C' \subseteq C$ of candidates an *elimination set* if there is a selected ranking where the candidates from C' are ranked in the first $|C'|$ positions. We introduce a table T with entry $T[C']$ for each subset $C' \subseteq C$ of candidates. $T[C']$ is set to true if C' is an elimination set. We initialize the table by setting $T[\emptyset]$ to true. Now we compute T for each subset $C' \subseteq C$ in increasing order of the size of the subset using the following recurrence relation: We set $T[C']$ to true if there is a candidate $c \in C'$ such that $T[C' \setminus \{c\}]$ is true and c is an \mathbf{s} -winner in $P|_{C \setminus (C' \setminus \{c\})}$.

After filling the table, we return “true” if and only if there is a subset $C' \subseteq C \setminus \{d\}$ with $|C'| = k - 1$ such that $T[C']$ is true and d is an \mathbf{s} -winner in $P|_{C \setminus C'}$. By filling the complete table we get a running time in $\mathcal{O}(2^m \cdot nm^2)$. However, it is sufficient to only fill the table for all subsets of size at most $k - 1$, resulting in a running time in $\mathcal{O}(m^k \cdot nm^2)$. \square

6.2. Sequential Loser

We study Seq.-Plurality/Veto/Borda-Loser (aka STV, Coombs, and Baldwin). The WINNER DETERMINATION problem is NP-hard for all three rules. Table 2 shows an overview of our results. In particular, we get a clear separation of the rules for the number n of voters:

- Seq.-Plurality-Loser admits a simple FPT algorithm,
- Seq.-Veto-Loser is W[1]-hard but in XP,
- Seq.-Borda-Loser is NP-hard for 8 voters.

6.2.1. Plurality

Conitzer et al. [2009] stated that WINNER DETERMINATION for Seq.-Plurality-Loser (aka STV) is NP-hard. This result has been frequently cited and used. The proof was omitted in the conference paper, and to our knowledge no proof has ever appeared in published work. To aid future research, we include a simple reduction here.

Theorem 6.2. *WINNER DETERMINATION for Sequential-Plurality-Loser (aka STV) is NP-hard.*

Proof. We reduce from the NP-hard variant of SATISFIABILITY where each clause contains at most three literals and each literal appears exactly twice Berman et al. [2003]. Let φ be a formula fulfilling these restrictions with clause set $F = \{c_1, \dots, c_m\}$ and variable set $X = \{x_1, \dots, x_n\}$. Let $L = X \cup \bar{X}$ be the set of literals. We construct a ranking profile with candidate set

	n	k	$n + k$	m
Sequential-Plurality-Winner	NP-c. W[1]-h., XP	W[1]-h., XP	FPT	FPT
Sequential-Veto-Winner	NP-c. FPT	W[2]-h., XP	FPT	FPT
Sequential-Borda-Winner	NP-c. NP-h. for $n = 8$	W[1]-h., XP	?	FPT

Table 3: Our results for Sequential-Winner rules. All hardness results hold for the TOP- k DETERMINATION problem; all algorithmic results also apply to the general POSITION- k DETERMINATION problem.

$C = \{d, w\} \cup F \cup L$, where d is our designated candidate, and the following voters:

- 100 voters $d \succ \dots$
- 99 voters $w \succ d \succ \dots$
- 98 voters $c_j \succ w \succ d \succ \dots \quad \forall j \in [m]$
- 60 voters $\ell \succ \bar{\ell} \succ w \succ d \succ \dots \quad \forall \ell \in L$
- 2 voters $\ell \succ c_j \succ w \succ d \succ \dots \quad \forall \ell \in L, j \in [m]$ where ℓ appears in c_j

For this ranking profile, in every execution of Sequential-Plurality-Loser the first n eliminated candidates must be a subset $L' \subseteq L$ of literals such that for every variable we select either its positive literal or its negative literal (but not both). In other words, L' must satisfy $\ell \in L' \leftrightarrow \bar{\ell} \notin L'$. To see this, note that all literal candidates initially have a Plurality score of 64, which is the lowest Plurality score in the profile, and that all other candidates have a higher Plurality score. Thus, in the first round an arbitrary literal ℓ of some variable x is eliminated. This increases the Plurality score of the opposite literal $\bar{\ell}$ to over 120. In the second round, we have to eliminate again an arbitrary literal (however, this time a literal corresponding to a variable different from x). We repeat this process for n rounds until for each variable exactly one of the corresponding literals has been eliminated. We claim that an execution of Sequential-Plurality-Loser eliminates d last if and only if the assignment that sets all literals from L' to *true* satisfies φ .

Suppose φ is satisfied by some variable assignment α , and consider an execution of Sequential-Plurality-Loser that begins by eliminating the n literals set to true in α . After this, the scores of the remaining candidates are:

- (i) d has 100 points,
- (ii) w has 99 points,
- (iii) c_j for $j \in [m]$ has between 100 and 104 points (as at least one of the literals occurring in c_j has been eliminated), and
- (iv) each literal $\ell \in L$ set to false by α has 124 points.

In the next round, w is eliminated, reallocating its 99 points to d . Then, in the next m rounds, each clause candidate c_j is eliminated, in each round reallocating its points to d . Finally, the remaining literals are eliminated, also each reallocating their points to d . Thus, d is the last remaining candidate and ranked in the first position in the selected ranking.

Let $L' \subseteq L$ be the set of literals eliminated in the first n rounds in some execution of the Sequential-Plurality-Loser rule (recall that $\ell \in L' \leftrightarrow \bar{\ell} \notin L'$). Suppose that the assignment α setting all literals from L' to true does not satisfy φ . After the literals from L' have been eliminated, the scores of the remaining candidate are:

- (i) d has 100 points,
- (ii) w has 99 points,
- (iii) c_j for $j \in [m]$ where α satisfies c_j has between 100 and 104 points,
- (iv) c_j for $j \in [m]$ where α does not satisfy c_j has 98 points, and
- (v) each literal $\ell \in L$ set to false by α has 124 points.

Thus, in the next round, one of the unsatisfied clauses is eliminated, redistributing its 98 points to w bringing the score of w to 197. Because all but 100 voters prefer w to d , the Plurality score of d will never exceed the score of w in consecutive rounds, so d cannot be eliminated last. \square

Motivated by this hardness result, we now turn to the problem’s parameterized complexity. We have already seen in [Theorem 6.1](#) that the problem is solvable in $\mathcal{O}(2^m \cdot nm^2)$ time. Indeed, we show that unless the Exponential Time Hypothesis (ETH)³ is false, we cannot hope to substantially improve the exponential part of this running time.

Theorem 6.3. *If the ETH is true, then WINNER DETERMINATION for Sequential-Plurality-Loser (aka STV) cannot be solved in $2^{o(m)} \cdot \text{poly}(n, m)$ time.* [Proof]

Turning to the number n of voters, we can observe that initially only at most n candidates have a non-zero Plurality score. All other candidates (which are not ranked first in any ranking) will be eliminated immediately, without thereby changing the Plurality scores of other candidates. After these eliminations, we are left with at most n candidates. This makes it easy to see that POSITION- k DETERMINATION is fixed-parameter tractable with respect to n (by using [Theorem 6.1](#)).

Observation 6.4. POSITION- k DETERMINATION for Sequential-Plurality-Loser (aka STV) is solvable in $\mathcal{O}(2^n \cdot nm^2)$ time.

6.2.2. Veto

We now turn to Seq.-Veto-Loser (aka Coombs). [Mattei et al. \[2014\]](#) showed that the WINNER DETERMINATION problem for this rule is NP-hard. We give an alternative NP-hardness proof that also implies an ETH-based lower bound for the parameter m .

Theorem 6.5. WINNER DETERMINATION for Sequential-Veto-Loser (aka. Coombs) is NP-complete. *If the ETH is true, then the problem cannot be solved in $2^{o(m)} \cdot \text{poly}(n, m)$ time.* [Proof]

For the parameter n , we show that the problem is W[1]-hard with respect to the number of voters. This is shown via an involved reduction from MULTICOLORED INDEPENDENT SET. This result suggests that Seq.-Veto-Loser behaves quite differently from Seq.-Plurality-Loser, even if these two rules might seem “symmetric” to each other.

Theorem 6.6. WINNER DETERMINATION for Sequential-Veto-Loser (aka. Coombs) is W[1]-hard with respect to the number n of voters. [Proof]

However, on the positive side, WINNER DETERMINATION and even POSITION- k DETERMINATION are solvable in polynomial-time if the number of voters is a constant. The intuition behind this result is that for Seq.-Veto-Loser, the “status” of an execution is fully captured by the *bottom list* of the ranking profile, i.e., a list containing the bottom-ranked candidate of each voter. Indeed, if we know the current bottom list, we can deduce exactly which candidates have

been eliminated thus far. As there are only m^n many possibilities for the bottom list, dynamic programming yields an XP algorithm for POSITION- k DETERMINATION.

Theorem 6.7. POSITION- k DETERMINATION for Sequential-Veto-Loser is in XP with respect to the number n of voters. [Proof]

6.2.3. Borda

We conclude by studying Seq.-Borda-Loser (aka Baldwin). Mattei et al. [2014] proved that WINNER DETERMINATION for this rule is NP-hard, adapting an earlier reduction about hardness of manipulation due to Davies et al. [2014]. In fact, by giving a construction based on weighted majority graphs and using tools from Bachmeier et al. [2019], we prove that this NP-hardness persists even for only $n = 8$ voters. This result suggests that the Borda scoring system leads to the hardest computational problems.

Theorem 6.8. Let $n \geq 8$ be a fixed even integer. Then WINNER DETERMINATION for Sequential-Borda-Loser (aka Baldwin), restricted to instances with exactly n voters, is NP-complete. In addition, if the ETH is true, then the problem cannot be solved in $2^{o(m)} \cdot \text{poly}(m)$ time. [Proof]

6.3. Sequential Winner

In this subsection, we briefly summarize our results for Seq.-Plurality/Veto/Borda-Winner, which to the best of our knowledge have not been previously studied (for formal statements and proofs see Appendix D). As WINNER DETERMINATION is trivial for these rules, we focus on TOP- k DETERMINATION. Table 3 displays an overview of our results. For all three rules, it turns out that TOP- k DETERMINATION is NP-hard and W[1]-hard with respect to k . In contrast, for the parameter n , the picture is again more diverse: For Borda, we once more get NP-hardness for a constant number of voters ($n = 8$), while Plurality and Veto switch their role (we have a fixed-parameter tractable algorithm for Veto and W[1]-hardness for Plurality). Recalling the equivalence from Lemma 3.5, this switch is unsurprising. Indeed, similar reductions are used here as for the corresponding results for Sequential-Loser for the other scoring system.

7. Future Directions

There are many directions for future work. In our complexity study, we have focused on the analysis of the space of possible outcomes. However, if we are happy to break ties immediately (e.g. by some fixed order), one could focus on finding the fastest algorithms for computing the output ranking. Interestingly, it is known that computing STV is P-complete [Csar et al., 2017], so its computation is unlikely to be parallelizable. An additional challenging open problem will be to determine if the hard problems we have identified become tractable if preferences are structured, for example single-peaked. Note that for single-peaked preferences, it is known that Coombs becomes a Condorcet extension and easy to compute [Grofman and Feld, 2004, Prop. 2]. Further, ranking candidates by Borda score is known to give a 5-approximation of Kemeny’s method [Coppersmith et al., 2006]. This raises the question whether any other of the rules from our families provide an approximation? Other specific questions left open by our work are whether ETH lower bounds can be obtained for additional problems, and whether they can be strengthened to SETH bounds. Finally, one could try to extend our results to other scoring vectors, and potentially prove dichotomy theorems.

³The ETH was introduced by Impagliazzo et al. [2001] and states that 3-SAT cannot be solved in $2^{o(n)} \cdot \text{poly}(n)$ time where n is the number of variables.

8. Bibliography

- Nir Ailon, Moses Charikar, and Alantha Newman. Aggregating inconsistent information: ranking and clustering. *Journal of the ACM*, 55(5):1–27, 2008. [→ p. 2]
- Saeed Akhoondian Amiri. A note on the fine-grained complexity of MIS on regular graphs. *Information Processing Letters*, 170:106123, 2021. [→ p. 26]
- Kenneth J. Arrow. *Social Choice and Individual Values*. New Haven: Cowles Foundation, 1st edition, 1951. [→ p. 2, 4, 6]
- Georg Bachmeier, Felix Brandt, Christian Geist, Paul Harrenstein, Keyvan Kardel, Dominik Peters, and Hans Georg Seedig. k -majority digraphs and the hardness of voting with a constant number of voters. *Journal of Computer and System Sciences*, 105:130–157, 2019. [→ p. 13, 36]
- John Bartholdi, Craig A. Tovey, and Michael A. Trick. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare*, 6(2):157–165, 1989. [→ p. 2]
- Piotr Berman, Marek Karpinski, and Alex D. Scott. Approximation hardness of short symmetric instances of MAX-3SAT. *Electronic Colloquium on Computational Complexity*, 2003. [→ p. 10]
- Niclas Boehmer, Robert Brederick, Piotr Faliszewski, Rolf Niedermeier, and Stanislaw Szufa. Putting a compass on the map of elections. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 59–65, 2021. [→ p. 7]
- Walter Bossert and Yves Sprumont. Strategy-proof preference aggregation: Possibilities and characterizations. *Games and Economic Behavior*, 85:109–126, 2014. [→ p. 2]
- Markus Brill and Felix Fischer. The price of neutrality for the ranked pairs method. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1299–1305, 2012. [→ p. 3]
- Burak Can and Ton Storcken. Update monotone preference rules. *Mathematical Social Sciences*, 65(2):136–149, 2013. [→ p. 2]
- Vincent Conitzer, Matthew Rognlie, and Lirong Xia. Preference functions that score rankings and maximum likelihood estimation. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 109–115, 2009. [→ p. 3, 9, 10]
- Don Coppersmith, Lisa Fleischer, and Atri Rudra. Ordering by weighted number of wins gives a good ranking for weighted tournaments. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 776–782, 2006. [→ p. 2, 13]
- Theresa Csar, Martin Lackner, Reinhard Pichler, and Emanuel Sallinger. Winner determination in huge elections with MapReduce. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 451–458, 2017. [→ p. 13]
- Jessica Davies, George Katsirelos, Nina Narodytska, Toby Walsh, and Lirong Xia. Complexity of and algorithms for the manipulation of Borda, Nanson’s and Baldwin’s voting rules. *Artificial Intelligence*, 217:20–42, 2014. [→ p. 13]
- Bernard Debord. Caractérisation des matrices des préférences nettes et méthodes d’agrégation associées. *Mathématiques et sciences humaines*, 97:5–17, 1987. [→ p. 36, 45]
- Cynthia Dwork, Ravi Kumar, Moni Naor, and Dandapani Sivakumar. Rank aggregation methods for the web. In *Proceedings of the 10th International Conference on the World Wide Web (WWW)*, pages 613–622, 2001. [→ p. 2]

- Paul Erdős and Leo Moser. On the representation of directed graphs as unions of orderings. *Publications of the Mathematical Institute of the Hungarian Academy of Science*, 9:125–132, 1964. [→ p. 36]
- Zack Fitzsimmons and Edith Hemaspaandra. Kemeny consensus complexity. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 196–202, 2021. [→ p. 2]
- Rupert Freeman, Markus Brill, and Vincent Conitzer. On the axiomatic characterization of runoff voting rules. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, pages 675–681, 2014. [→ p. 6, 18, 19, 20]
- Rupert Freeman, Markus Brill, and Vincent Conitzer. General tiebreaking schemes for computational social choice. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1401–1409, 2015. [→ p. 3]
- Bernard Grofman and Scott L. Feld. If you like the alternative vote (a.k.a. the instant runoff), then you ought to know about the Coombs rule. *Electoral Studies*, 23(4):641–659, 2004. [→ p. 13]
- Edith Hemaspaandra, Holger Spakowski, and Jörg Vogel. The complexity of Kemeny elections. *Theoretical Computer Science*, 349(3):382–391, 2005. [→ p. 2]
- Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *Journal of Computer and System Sciences*, 63(4):512–530, 2001. [→ p. 13]
- David S. Johnson and Mario Szegedy. What are the least tractable instances of max independent set? In *Proceedings of the 10th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 927–928, 1999. [→ p. 26]
- John G. Kemeny. Mathematics without numbers. *Daedalus*, 88:577–591, 1959. [→ p. 2]
- Claire Kenyon-Mathieu and Warren Schudy. How to rank with few errors. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing (STOC)*, pages 95–103, 2007. [→ p. 2]
- Christian Komusiewicz. Tight running time lower bounds for vertex deletion problems. *ACM Transactions on Computation Theory*, 10(2), 2018. [→ p. 26]
- Arthur B. Levenglick. *Characterizations of Social Decision Functions*. PhD thesis, City University of New York, 1977. [→ p. 3, 6]
- Colin L. Mallows. Non-null ranking models. i. *Biometrika*, 44(1/2):114–130, 1957. [→ p. 7]
- Nicholas Mattei, Nina Narodytska, and Toby Walsh. How hard is it to control an election by breaking ties? In *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI)*, pages 1067–1068, 2014. [→ p. 3, 9, 12, 13]
- David C. McGarvey. A theorem on the construction of voting paradoxes. *Econometrica*, 21(4): 608–610, 1953. [→ p. 36, 45, 46]
- John H. Smith. Aggregation of preferences with variable electorate. *Econometrica*, 41(6): 1027–1041, 1973. [→ p. 3, 6]
- T. Nicolaus Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4(3):185–206, 1987. [→ p. 6, 19]

- Michel Truchon. Figure skating and the theory of social choice. *Cahier 98-16 du Centre de Recherche en Économie et Finance Appliquées, Québec*, 1998. [→ p. 2]
- Anke van Zuylen and David P. Williamson. Deterministic pivoting algorithms for constrained ranking and clustering problems. *Mathematics of Operations Research*, 34(3):594–620, 2009. [→ p. 2]
- Jun Wang, Sujoy Sikdar, Tyler Shepherd, Zhibing Zhao, Chunheng Jiang, and Lirong Xia. Practical algorithms for multi-stage voting rules with parallel universes tiebreaking. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 2189–2196, 2019. [→ p. 3]
- H. Peyton Young. Condorcet’s theory of voting. *The American Political Science Review*, 82(4): 1231–1244, 1988. [→ p. 6, 18]
- H. Peyton Young. Optimal voting rules. *Journal of Economic Perspectives*, 9(1):51–64, 1995. [→ p. 2]
- H. Peyton Young and Arthur Levenglick. A consistent extension of Condorcet’s election principle. *SIAM Journal on Applied Mathematics*, 35(2):285–300, 1978. [→ p. 2, 4]

A. Additional Material for Section 4

In the main body, in Section 4, we have given an informal overview of axiomatic properties satisfied by the rules in our three families. In this appendix, we give formal statements of these results. In particular, we will give formal definitions of the relevant axioms.

Let us introduce some additional notation. For two ranking profiles $P = (\succ_1, \dots, \succ_n)$ and $P' = (\succ'_1, \dots, \succ'_n)$, both defined over the same candidate set, we write $P + P' = (\succ_1, \dots, \succ_n, \succ'_1, \dots, \succ'_n)$ for the ranking profile obtained by concatenating the two lists. For an integer k , we write $kP = P + \dots + P$ obtained by concatenating k copies of P . For a set $S \subseteq \mathcal{L}(C)$ of rankings, we write $\text{cand}(S, r) = \{\text{cand}(\succ, r) : \succ \in S\}$ for the set of candidates that appear in position r in at least one of the rankings in S . If $\succ \in \mathcal{L}(C)$ is a ranking and $\rho : C \rightarrow C$ is a permutation of the candidate set, then $\rho(\succ)$ is the ranking where for all pairs $a, b \in C$ of candidates, we have $\rho(a) \rho(\succ) \rho(b)$ if and only if $a \succ b$. For a set $S \subseteq \mathcal{L}(C)$ of rankings, we write $\rho(S) = \{\rho(\succ) : \succ \in S\}$.

Let f be a social preference function, defined for profiles with any number of voters and over all possible candidate sets. (We assume this large domain to be able to state axioms that reason about variable agendas (i.e. different candidate sets) and about variable electorates (i.e., different numbers of voters).) Whenever we do not specify otherwise, in the following axioms we implicitly quantify over all possible finite sets C of candidates.

We begin with some basic axioms.

- The rule f is *anonymous* if for all profiles $P = (\succ_1, \dots, \succ_n)$ and all permutations $\sigma : [n] \rightarrow [n]$, we have $f(P) = f((\succ_{\sigma(1)}, \dots, \succ_{\sigma(n)}))$. Thus, reordering the rankings does not change the outcome.
- The rule f is *neutral* if for all profiles $P = (\succ_1, \dots, \succ_n)$ and all permutations $\rho : C \rightarrow C$, we have $f(\rho(P)) = \rho(f(P))$. Thus, a relabeling of candidates leads to the same relabeling of the output.
- The rule f is *unanimous* if for all rankings $\succ \in \mathcal{L}(C)$ and all profiles $P = (\succ, \dots, \succ)$, where all rankings in P are equal to \succ , we have $f(P) = \{\succ\}$.
- The rule f is *continuous* (sometimes known as the *overwhelming majority* axiom) if for any two profiles P and P' over the same candidate set, there exists an integer k such that $f(P + kP') \subseteq f(P')$.

The following are axioms about combining profiles.

- The rule f satisfies *reinforcement* if for all profiles P and P' over the same candidate set, we have $f(P + P') = f(P) \cap f(P')$ whenever the intersection is non-empty.
- The rule f satisfies *reinforcement at the top* if for all profiles P and P' over the same candidate set, we have $\text{cand}(f(P + P'), 1) = \text{cand}(f(P), 1) \cap \text{cand}(f(P'), 1)$ whenever the intersection is non-empty.
- The rule f satisfies *reinforcement at the bottom* if for all profiles P and P' over the same candidate set C , we have $\text{cand}(f(P + P'), |C|) = \text{cand}(f(P), |C|) \cap \text{cand}(f(P'), |C|)$ whenever the intersection is non-empty.

The following are independence axioms, describing that the output should not change when deleting certain candidates.

- The rule f satisfies *independence at the top* if for all profiles P and for all candidates $a \in \text{cand}(f(P), 1)$ that can appear in first position in $f(P)$, we have that for all rankings $\succ' \in \mathcal{L}(C \setminus \{a\})$, $\succ' \in f(P|_{C \setminus \{a\}})$ if and only if the ranking \succ'' , obtained by placing a at the top of ranking \succ' , is a member of $f(P)$.

- The rule f satisfies *independence at the bottom* if for all profiles P and for all candidates $a \in \text{cand}(f(P), |C|)$ that can appear in last position in $f(P)$, we have that for all rankings $\succ' \in \mathcal{L}(C \setminus \{a\})$, $\succ' \in f(P|_{C \setminus \{a\}})$ if and only if the ranking \succ'' , obtained by placing a at the bottom of ranking \succ' , is a member of $f(P)$.

Next, we introduce axioms that require a rule to follow the view of a majority of voters.

- The rule f places *Condorcet winners at top* if for all profiles $P = (\succ_1, \dots, \succ_n)$ where there exists a candidate $a \in C$ such that for all other candidates $b \in C \setminus \{a\}$, a majority of voters prefers a to b (i.e., $|\{i \in [n] : a \succ_i b\}| > |\{i \in [n] : b \succ_i a\}|$), we have $\text{cand}(f(P), 1) = \{a\}$. Thus, in profiles where a Condorcet winner exists, all output rankings must place it in the first position.
- The rule f *copies a majority ranking* if for all profiles $P = (\succ_1, \dots, \succ_n)$ where there exists a ranking \succ which makes up more than half the profile (i.e., $|\{i \in [n] : \succ_i = \succ\}| > n/2$), we have $f(P) = \{\succ\}$.

We will state Tideman’s independence of clones property later.

We will now state some axiomatic characterization results, taken or adapted from the literature. The first is a characterization of Kemeny’s rule.

Theorem A.1 (Young, 1988). *A social preference function f satisfies anonymity, neutrality, unanimity, reinforcement, independence at the top, and independence at the bottom, if and only if it is Kemeny’s rule.*

Next, there is an existing characterization of Sequential-Loser rules.

Theorem A.2 (Freeman et al., 2014, Lemma 1). *A social preference function f satisfies anonymity, neutrality, unanimity, continuity, reinforcement at the bottom, and independence at the bottom, if and only if there exists a scoring system $\mathbf{s} = (\mathbf{s}^{(m)})_{m \in \mathbb{N}}$ such that f equals Sequential- \mathbf{s} -Loser.*

Proof. This is exactly Lemma 1 from Freeman et al. [2014], except that instead of continuity they use a condition called “continuity at the bottom”, but this condition is weaker than continuity as we have defined it, because $f(P + kP') \subseteq f(P')$ implies that $\text{cand}(f(P + kP'), |C|) \subseteq \text{cand}(f(P'), |C|)$. \square

We will now “turn around” Theorem A.2 to obtain an axiomatic characterization of Sequential-Winner rules. To do so, we will use Lemma 3.5. For a social preference function f , let us write f^* for the social preference function defined as follows:

$$f^*(P) = \text{rev}(f(\text{rev}(P))) \quad \text{for every profile } P.$$

From Lemma 3.5, we know that if f is Seq- \mathbf{s} -Winner then f^* is Seq- \mathbf{s}^* -Loser, and if f is Seq- \mathbf{s} -Loser then f^* is Seq- \mathbf{s}^* -Winner. One can also easily deduce the following equivalences.

Lemma A.3. *Let f be a social preference function, and let f^* be defined as above. Then the following equivalences hold:*

- f satisfies anonymity (resp., neutrality, unanimity, continuity, reinforcement, copying a majority ranking) if and only if f^* satisfies the respective axiom.
- f satisfies reinforcement at the top (resp., at the bottom) if and only if f^* satisfies reinforcement at the bottom (resp., at the top).

- f satisfies independence at the top (resp., at the bottom) if and only if f^* satisfies independence at the bottom (resp., at the top).

Theorem A.4. *A social preference function f satisfies anonymity, neutrality, unanimity, continuity, reinforcement at the top, and independence at the top, if and only if there exists a scoring system $\mathbf{s} = (\mathbf{s}^{(m)})_{m \in \mathbb{N}}$ such that f equals Sequential- \mathbf{s} -Winner.*

Proof. It is routine to check that Seq.- \mathbf{s} -Winner satisfies the mentioned axioms for every scoring system \mathbf{s} . Now let f be a social preference function satisfying these axioms. By Lemma A.3, the social preference function f^* then satisfies anonymity, neutrality, unanimity, continuity, reinforcement at the bottom, and independence at the bottom. By Theorem A.2, there exists a scoring system \mathbf{s} such that f^* equals Seq.- \mathbf{s} -Loser. By Lemma 3.5, f equals Seq.- \mathbf{s}^* -Winner, as desired. \square

Freeman et al. [2014] also provided characterizations of specific rules. For example, they characterize Seq.-Borda-Loser (aka Baldwin) as the only Sequential-Loser rule that places Condorcet winners at the top.

Theorem A.5 (Freeman et al., 2014, Theorem 3). *A social preference function f satisfies anonymity, neutrality, unanimity, continuity, reinforcement at the bottom, independence at the bottom, and places Condorcet winners at the top if and only if f equals Sequential-Borda-Loser.*

If desired, one could similarly characterize Seq.-Borda-Winner as the only Sequential-Winner rule that places Condorcet losers at the bottom, using Lemma A.3.

Freeman et al. [2014] also characterize Seq.-Veto-Loser (aka Coombs) as the only Sequential-Loser rule that copies a majority ranking.

Theorem A.6 (Freeman et al., 2014, Theorem 2). *A social preference function f satisfies anonymity, neutrality, unanimity, continuity, reinforcement at the bottom, independence at the bottom, and copies majority rankings if and only if f equals Sequential-Veto-Loser.*

Using Lemmas 3.5 and A.3, we can deduce a characterization of Seq.-Plurality-Winner as the only Sequential-Winner rule that copies majority rankings.

Theorem A.7. *A social preference function f satisfies anonymity, neutrality, unanimity, continuity, reinforcement at the top, independence at the top, and copies majority rankings if and only if f equals Sequential-Plurality-Winner.*

Turning towards Tideman's [1987] independence of clones property, Freeman et al. [2014] adapt the axiom to the ranking context. Given a profile P , we say that $C' \subseteq C$ is a *clone set* if in every ranking in P , the candidates in C' appear consecutively (i.e., for each $\succ \in P$ and each $a \in C \setminus C'$ we either have $a \succ c'$ for all $c' \in C'$ or $c' \succ a$ for all $c' \in C'$).

- The rule f satisfies *independence of clones (with top replacement)* if the following property is satisfied. Let P be a profile and let C' be a clone set. Let P' be a profile obtained from P by replacing the set C' by a single candidate $a \notin C$. Then for all rankings $\succ \in \mathcal{L}(C)$, we have $\succ \in f(P)$ if and only if $\succ' \in f(P')$ where \succ' is the ranking obtained from \succ by deleting the candidates in C' and putting a at the position of the highest-ranked member of C' in \succ .

Freeman et al. [2014] prove that within the family of Sequential-Loser rules, Seq.-Plurality-Loser (aka STV) is characterized by this axiom.⁴

⁴In their statement of the characterization, they appear to have forgotten to include continuity.




Figure 2: Pairwise average normalized swap distance between rankings produced by different methods for Plurality (solid) and Borda (dashed) on Euclidean profiles with 10 candidates and 100 voters.

Theorem A.8 (Freeman et al., 2014, Theorem 1). *A social preference function f satisfies anonymity, neutrality, unanimity, continuity, reinforcement at the bottom, independence at the bottom, and independence of clones (with top replacement) if and only if f equals Sequential-Plurality-Loser.*

By invoking Lemmas 3.5 and A.3, we can again obtain a related characterization of Seq.-Veto-Winner, but using a slightly different version of the clones axioms. In particular, let us define the axiom *independence of clones (with bottom replacement)* exactly as before, except that the definition should end in saying “putting a at the position of the *lowest*-ranked member of C' in \succ .”

Theorem A.9. *A social preference function f satisfies anonymity, neutrality, unanimity, continuity, reinforcement at the top, independence at the top, and independence of clones (with bottom replacement) if and only if f equals Sequential-Veto-Winner.*

In the main body (and in particular Table 1), we have taken the “official” version of independence of clones to be the one with top replacement.

B. Additional Material for Section 5

B.1. Setup—Euclidean Model

In addition to ranking profiles generated from the Mallows model, we have also considered profiles generated from the *Euclidean model*. This model is parameterized by the dimension $d \geq 1$. To sample a Euclidean profile, for each candidate and voter we sample a point from the d -dimensional hypercube $[0, 1]^d$ uniformly at random. In the corresponding profile, each voter ranks the candidates in increasing order of their Euclidean ℓ_2 -distance to the voter.

B.2. Comparison of Ranking Methods—Euclidean Model

In the main body, in Section 5.2 and Section 5.3, we have analyzed the relation between rankings selected by the different rules on profiles generated using the Mallows model. To verify our results, we reran these experiments on profiles generated using the Euclidean model. Specifically, for each dimension $d \in \{1, 2, 3, 4, 5, 10, 15, 20\}$, we generated 10 000 profiles with 100 voters and 10 candidates. We depict the results in Figure 2 (which is analogous to Figure 1).

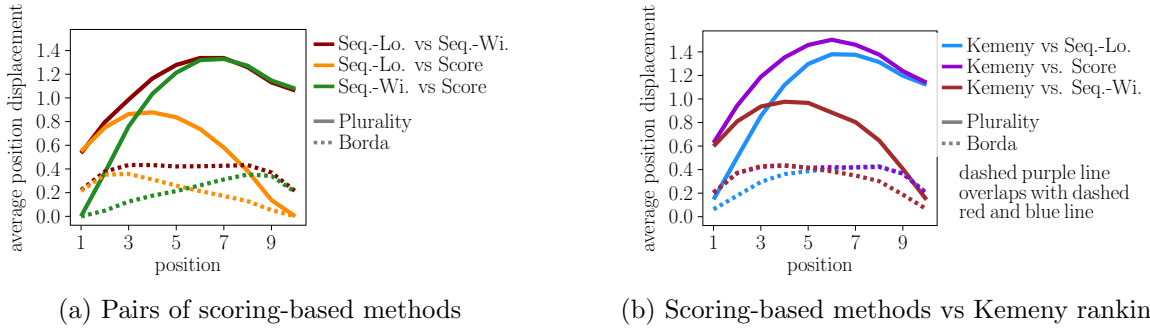


Figure 3: For pairs of rankings, average position displacement on each position for profiles generated using the Mallows model with $\text{norm-}\phi = 0.8$ with 10 candidates and 100 voters.

Comparison of Scoring-Based Ranking Methods We start by analyzing Figure 2(a), where we compare the rankings selected by our different scoring-based rules. The general trend here is quite similar to Mallows profiles for a large dispersion parameter (see Figure 1(a)): For Borda, the agreement of the three methods is much higher than for Plurality, with Borda-Score producing rankings close to the other two. For Plurality, Seq.-Plurality-Loser and Plurality-Score produce again similar results, whereas the ranking produced by Seq.-Plurality-Winner differs more. The general level of disagreement between the rules for Plurality is remarkably high here. For $d = 1$, the difference between Seq.-Plurality-Winner and the other two methods is around 0.4, which is almost 0.5 (the expected distance of two rankings drawn uniformly at random). Moreover, even for larger d , the level of disagreement remains high and is in particular around the level of disagreement for Mallows profiles with parameter $\text{norm-}\phi = 1$. This is somewhat surprising, as profiles produced by the Mallows model with $\text{norm-}\phi = 1$ are “maximally chaotic” and thus give the rules only limited information to distinguish the strength of candidates.

Comparison to Kemeny Ranking We now turn to the comparison of scoring-based ranking rules to Kemeny’s method (Figure 2(b)). For Plurality, like we have seen in Mallows profiles, Seq.-Plurality-Winner produces the best results, followed by Seq.-Plurality-Loser, and lastly Plurality-Score. Considering the influence of the dimension d , the difference between the method’s distance to the Kemeny ranking is more or less the same for all dimensions with $d = 1$ being the only exception: At $d = 1$, Seq.-Plurality-Loser and Plurality-Score are at normalized swap distance 0.43 from the Kemeny ranking, whereas Seq.-Plurality-Winner is at distance 0.27, highlighting again that Euclidean profiles with $d = 1$ are particularly challenging and that Seq.-Plurality-Winner does best.

In contrast, for Borda, the rankings produced by the three methods are all around the same small distance from the Kemeny ranking (mostly independently of the dimension).

B.3. Similarity in Different Vote Parts

To shed some further light on the relation of the different methods, we next analyze in which parts of the computed ranking the considered methods agree or disagree most. For this, for two rankings $\succ, \succ' \in \mathcal{L}(C)$, we define the *position displacement* in position $i \in [|C|]$ as

$$\frac{1}{2} (|i - \text{pos}(\succ, \text{cand}(\succ', i))| + |i - \text{pos}(\succ', \text{cand}(\succ, i))|).$$

The position displacement quantifies how far away the candidates ranked in position i in one ranking are ranked in the other ranking. Consider as an example the following two rankings:

$$\begin{aligned} a \succ b \succ c \succ d \\ d \succ' c \succ' a \succ' b \end{aligned}$$

Then, the position displacement on position 1 is $\frac{1}{2} \cdot (|1 - \text{pos}(\succ, d)| + |1 - \text{pos}(\succ', a)|) = \frac{1}{2} \cdot (|1 - 4| + |1 - 3|) = \frac{5}{2}$, whereas the position displacement on position 2 is $\frac{1}{2} \cdot (|2 - 3| + |2 - 4|) = \frac{3}{2}$. In [Figure 3](#), we show the average position displacement on 10 000 profiles with 100 voters and 10 candidates sampled from the Mallows model with $\text{norm-}\phi = 0.8$.⁵ (We chose $\text{norm-}\phi = 0.8$ in order to ensure that the tie-breaking rule plays no critical role, while still keeping some structure in the profile.)

First, the general picture for Plurality and Borda is similar in the sense that for all comparisons of rules, the shape of the respective curves is similar. Thus, all observations described in the following hold for both Plurality and Borda. Second, all four methods have a generally higher agreement on the top and bottom positions than on the middle positions. Third, focusing on the comparison of the different scoring-based methods ([Figure 3\(a\)](#)), by design Seq.-Winner and Score always place the same candidate in the first position. However, the agreement of the rules remains high in the second position and decreases continuously until position 7. This indicates the intuitive behavior that the more candidates are present in the current round, the higher is the correlation between the scores of the candidates in the initial profile and their score in this round. Comparing Seq.-Loser to Score, a reverse effect is present. For Seq.-Loser compared to Seq.-Winner, we almost have a symmetric curve with a generally slightly higher agreement on the top than on the bottom. Fourth, focusing on the comparison of the scoring-based methods and the Kemeny ranking ([Figure 3\(b\)](#)), interestingly, on the top positions the Kemeny ranking agrees most with Seq.-Loser while on the bottom it agrees most with Seq.-Winner. This suggests that one should use Seq.-Loser for identifying the best candidates and Seq.-Winner for identifying the worst candidates.

B.4. Number of Ties in Executions of the Rules

The main motivation for our complexity analysis is that ties might occur in the execution of our rules. To better understand whether ties actually occur in practice (and to provide evidence for the explanation we gave in the main body about the behavior of Plurality-Score and Seq.-Plurality-Loser for small dispersion parameters), we conducted the following experiment. We again sampled 10 000 profiles from the Mallows model, for each $\text{norm-}\phi \in \{0, 0.1, \dots, 1\}$. For each profile, we executed our rules, as usual breaking ties according to a tie-breaking order \succ_{tie} that we sampled randomly, and checked in each round whether a tie is present.⁶ The average number of rounds when a tie occurred is shown in [Figure 4](#).

For Plurality, we previously observed that on Mallows profiles with a small dispersion parameter, Seq.-Plurality-Loser and Plurality-Score produce very similar rankings. We mentioned that this could be explained by a large number of ties (that get resolved via the same tie-breaking order). This observation is clearly confirmed here (see [Figure 4\(a\)](#)). In contrast, for Seq.-Plurality-Winner there are only very few (or no) ties for a small dispersion parameter, while more ties appear if

⁵We rerun this experiment using profiles sampled from the Euclidean model with $d = 10$ producing a very similar picture.

⁶We use the following interpretation of the Score method. We start by computing the scores of the candidates in the initial profile. Subsequently, we eliminate a candidate with the highest number of points and add it at the first position in the selected ranking. However, we do not recompute the scores. In the second round, we eliminate the remaining candidate with the maximum number of points (in the initial profile) and add it in the second position. We repeat this process until all candidates are eliminated.

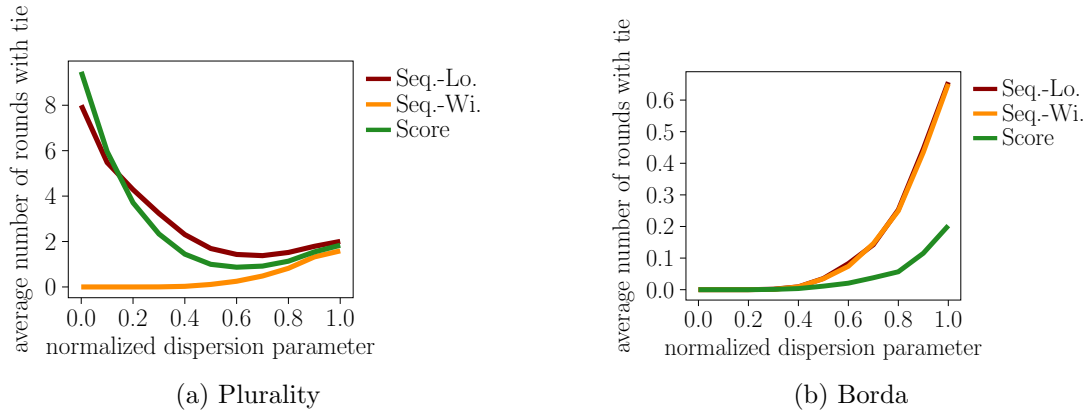


Figure 4: Average number of rounds in which a tie occurs for Seq.-Plurality-Winner, Seq.-Plurality-Loser, and Plurality-Score on profiles with 10 candidates and 100 voters sampled from the Mallows model.

the dispersion parameter is increased. In general, for each choice of the normalized dispersion parameter, Seq.-Plurality-Winner produces the smallest number of ties among the three methods based on Plurality. For $\text{norm-}\phi = 1$, for all three rules, there are ties in around 2 rounds on average (we were slightly surprised that this number is so high). Overall, these results show the importance of tie-breaking for ranking methods based on Plurality scores. They also provide an argument in favor of Seq.-Plurality-Winner over the other Plurality-based rules, as this rule leaves fewer decisions to the tie-breaking rule.

For Borda (Figure 4(b)), Borda-Score produces the least number of ties. Seq.-Borda-Winner and Seq.-Borda-Loser have more ties, at very similar rates to each other. The number of rounds with ties increases with a higher dispersion parameter. However, even for $\text{norm-}\phi = 1$, there are only very few ties for all Borda-based rules: on average around 0.6 rounds with a tie for Seq.-Borda-Winner and Seq.-Borda-Loser, and only 0.2 rounds for Borda-Score. (Compare this to 2 rounds for the Plurality-based rules.)

For Euclidean models, the number of rounds with a tie mostly does not vary with the dimension, and there are many fewer ties than for the Mallows model. In particular, for Sequential-Plurality-Winner, the average number of rounds with a tie is around 0.7, for Plurality-Score it is around 1, and for Sequential-Plurality-Loser it is around 1.4. Again we see that Sequential-Plurality-Winner produces the fewest ties among the Plurality-based rules. Moreover, for Borda the number of ties is again much lower than for Plurality, with Borda-Score producing the fewest ties, while Seq.-Borda-Winner and Seq.-Borda-Loser give results very similar to each other. Specifically, for Borda-Score the average number of rounds with a tie is around 0.05, and for Seq.-Borda-Winner and Seq.-Borda-Loser it is around 0.2.

B.5. Influence of Profile Size

So far, we have focused on profiles with $n = 100$ voters and $m = 10$ candidates. Now, we examine the influence of the size of our profile on the results. First, we will analyze the influence of varying the number of voters and second, we will analyze the influence of varying the number of candidates.

Varying the number of voters In Figure 5, we depict the pairwise difference of our ranking methods for Plurality and the Kemeny ranking, for profiles with 10 candidates and a varying number of voters. The profiles are generated using the Mallows model with $\text{norm-}\phi = 0.8$ (Figure 5(a)) and the Euclidean model with $d = 10$ (Figure 5(b)). For each $n \in \{25, 50, 100, 200, 300, 400, 500\}$,




Figure 5: Pairwise average normalized swap distance between the Kemeny method, Seq.-Plurality-Winner, Seq.-Plurality-Loser, and Plurality-Score rankings on profiles with 10 candidates and a varying number of voters.




Figure 6: Pairwise average normalized swap distance between the Kemeny ranking and rankings produced by our three scoring-based methods on profiles sampled from Mallows model with $\text{norm-}\phi = 0.8$ with 100 voters and a varying number of candidates.

we generated 10 000 profiles. For both generation models, the different methods become pairwise more similar with a higher number of voters. For Mallows the distance between the rankings decreases steeply, while for the Euclidean model the decrease (after $n = 100$) is slower. Generally speaking, increasing the number of voters gives us additional information about the strengths of the candidates and reduces the probability of artifacts. For Mallows profiles, there exists a clear ordering of the candidates in terms of their strengths (namely the central order), and additional voters clarify this situation. For Euclidean profiles, candidates are less clearly distinguishable; explaining why the four approaches do not all “converge” to the same ranking as the number of voters increases (unlike for Mallows). For both models, the ordering of pairs of our three scoring-based ranking methods in terms of their similarity is independent of the number of voters (the same is also the case for their ordering with respect to their similarity to the Kemeny ranking).

The above described general trends are also present if we use Borda instead of Plurality.

Varying the number of candidates We now turn to analyzing the influence of the number of candidates. In Figure 6, we depict the pairwise distances between our methods for Plurality (Figure 6(a)) and Borda (Figure 6(b)) in profiles with 100 voters and a varying number of candidates. The profiles are generated using the Mallows model with $\text{norm-}\phi = 0.8$. For the Euclidean model with $d = 10$, the results are similar, so we omit them. For each $m \in \{5, 10, 25, 50, 75\}$, we generated 100 profiles. Note that our use of *normalized* distances is particularly convenient when comparing results with differing numbers of candidates.

We start by examining the results for Plurality (Figure 6(a)). Here, as the number of candidates increases, the average number of Plurality points per candidate decreases and in particular more candidates get a Plurality score of zero. This leads to more ties in the execution




Figure 7: Pairwise average normalized swap distance between the Kemeny ranking and rankings produced by our three scoring-based methods for two different scoring systems on Mallows profiles with 100 voters and 10 candidates.

of Seq.-Plurality-Loser and Plurality-Score when determining the ordering of candidates at the bottom of the output ranking. Because both methods break ties using the same tie-breaking order \succ_{tie} , Seq.-Plurality-Loser and Plurality-Score become more similar as the number of candidates increases. On the flip side, when comparing Seq.-Plurality-Winner to Seq.-Plurality-Loser (or when comparing Seq.-Plurality-Winner to Plurality-Score), they become less similar as the number of candidates increases. This is because Seq.-Plurality-Winner is able to distinguish candidates who initially have low Plurality scores and does not need to rely on tie-breaking. Indeed, when Seq.-Plurality-Winner starts to rank the weak candidates, all other candidates have already been deleted and thus the average Plurality score of the weak candidates is higher and more informative.

Comparing our three scoring-based methods to the Kemeny ranking, we see that the distance increases with more candidates. This is because decreasing the average number of Plurality points per candidates makes it harder for our three scoring-based methods to distinguish the strengths of candidates; intuitively speaking, as we increase the number of candidates, the information provided by only examining the first position decreases (but, as discussed above, this effect is smaller for Seq.-Plurality-Winner).

For Borda (Figure 6(b)), the general trend is reversed. With small exceptions, the higher the number of candidates, the more similar are the rankings produced by the different methods. One effect that potentially contributes to this is that for a higher number of candidates the range of awarded points increases thereby allowing for a clearer distinction of the candidates. Nevertheless, for both Plurality and Borda, the ordering of the pairs of methods remains largely unaffected by changing the number of candidates.

Overall, the results from this section suggest that while the size of the profile in question influences the level of similarity of the different methods, the general trends observed in the previous sections hold mostly independent of the size of the profile.

B.6. Further Voting Rules

In this section, we briefly examine two additional scoring systems. In accordance with our theoretical analysis, we start by examining ranking methods based on the Veto scoring system with scoring vector $(0, \dots, 0, -1)$. Figure 7(a) shows the results of our experiment for the Mallows model (again we generated 10 000 profiles for each $\text{norm-}\phi \in \{0, 0.1, \dots, 0.9, 1\}$). Comparing Figure 7(a) to the analogous plots for Plurality (see Figure 1(a) and Figure 1(b)), we see that the results are very similar to each other after swapping the roles of Seq.-Winner and Seq.-Loser: For Plurality, Seq.-Plurality-Loser and Plurality-Score are closely related, while Seq.-Plurality-Winner produces results most similar to the Kemeny ranking. In contrast, for

Veto, Seq.-Veto-Winner and Veto-Score are closely related and Seq.-Veto-Loser is most similar to Kemeny. The switched role of Seq.-Winner and Seq.-Loser is to be expected, recalling our equivalence [Lemma 3.5](#). (For profiles sampled using the Euclidean model, similar conclusions apply.)

Recalling that for Veto, Seq.-Winner is similar to Score, whereas for Plurality Seq.-Loser is similar to Score, we wanted to check the behavior of our rules in between these two extremes. Thus, we introduce a new scoring system, which we call *Half*. This scoring system assigns one point to the first $\lfloor \frac{m}{2} \rfloor$ candidates and zero points to all other candidates, where m is the number of candidates. We depict the results for profiles sampled from the Mallows model in [Figure 7\(b\)](#) (again we generated 10 000 profiles for each norm- $\phi \in \{0, 0.1, \dots, 0.9, 1\}$). Indeed, in this case Seq.-Winner and Seq.-Loser are both at the same distance to Score and at the same distance to the Kemeny ranking. Nevertheless, naturally the rankings produced by Seq.-Half-Winner and Seq.-Half-Loser are still different. Remarkably, our three scoring-based ranking methods based on Half produce rankings that are closer to the Kemeny ranking than the rankings produced by any of our methods based on Plurality or Veto. This indicates the advantage of allowing voters to distinguish between many candidates to identify candidate strength.

C. Additional Material for [Section 6.2](#)

Theorem 6.1. *For every scoring system \mathbf{s} , POSITION- k DETERMINATION can be solved in*

- $\mathcal{O}(2^m \cdot nm^2)$ time and $\mathcal{O}(m^k \cdot nm^2)$ time for Sequential- \mathbf{s} -Winner, and
- $\mathcal{O}(2^m \cdot nm^2)$ time and $\mathcal{O}(m^{m-k} \cdot nm^2)$ time for Sequential- \mathbf{s} -Loser.

[Main]

Proof. It remains to prove the correctness of the algorithm described in the main body. For this, it is sufficient to prove that the recurrence relation is correct. To this end, assume a subset $C' \subseteq C$ of candidates is an elimination set as witnessed by the selected ranking \succ and let $c := \text{cand}(\succ, |C'|)$. Then, as witnessed by \succ , $C' \setminus \{c\}$ is an elimination set, and no candidate has a higher \mathbf{s} -score than c after deleting all candidates from $C' \setminus \{c\}$. If $T[C']$ is set to true because $C' \setminus \{c\}$ is an elimination set (as witnessed by \succ) for some candidate $c \in C'$, then C' is clearly an elimination set as we can eliminate candidates from $C' \setminus c$ in the first $|C'| - 1$ rounds (breaking ties according to \succ) and c in round $|C'|$. \square

C.1. Plurality

Before we present our ETH-based lower bound for the parameter m , we state some relevant results.

Proposition C.1 ([Amiri, 2021](#), extending results about subcubic vertex cover by [Johnson and Szegegy, 1999](#) and [Komusiewicz, 2018](#)). *If the Exponential Time Hypothesis (ETH) is true, there does not exist an algorithm solving CUBIC VERTEX COVER in time $2^{o(n)} \cdot \text{poly}(n)$, where n is the number of vertices.*

Corollary C.2. *If the ETH is true, there does not exist an algorithm solving REGULAR CLIQUE (the CLIQUE problem restricted to graphs where every vertex has the same degree) in time $2^{o(n)} \cdot \text{poly}(n)$, where n is the number of vertices.*

Proof. Suppose there was such an algorithm. Let $G = (V, E)$ be a cubic graph, and k be a target size for a vertex cover. Then the complement graph \overline{G} is regular. Recall that a set $T \subseteq V$ is a vertex cover in G if and only if $V \setminus T$ is an independent set in G if and only if $V \setminus T$ is a clique in \overline{G} . Hence by applying the algorithm to find a clique of size at least $n - k$ in \overline{G} , we can find a vertex cover of size at most k in G in time $2^{o(n)} \cdot \text{poly}(n)$, contradicting [Proposition C.1](#). \square

Theorem 6.3. *If the ETH is true, then WINNER DETERMINATION for Sequential-Plurality-Loser (aka STV) cannot be solved in $2^{o(m)} \cdot \text{poly}(n, m)$ time.*

[Main]

Proof. We reduce from CUBIC VERTEX COVER (given a graph where each vertex has degree three and an integer t , the problem asks whether there is a *vertex cover* of size t , that is, a set of t vertices such that each edge is incident to at least one of these vertices).

Given a cubic graph G with n vertices and $3n/2$ edges, and some integer t , we will create an equivalent instance of WINNER DETERMINATION for Seq.-Plurality-Loser. Since our reduction will only use $2n + 3n/2 + 3$ candidates, this implies the claimed ETH-result by [Proposition C.1](#).

Candidates For each vertex i , we create two *vertex candidates* v_i and v'_i . For each edge j , we create one *edge candidate* e_j . Moreover, we reference the edge candidates corresponding to edges incident to vertex i by e_1^i, e_2^i , and e_3^i . Finally we have the candidates d, w , and q .

Idea The idea is as follows. In the first n rounds, for each $i \in [n]$ either v_i or v'_i is eliminated. Eliminating v_i will correspond to selecting the i th vertex as part of the vertex cover. In round $n+1$, in order to make candidate d win the election, we must eliminate candidate w , because this is the latest point where it has smaller Plurality score than d . Eliminating any other candidate that can be eliminated in round $n+1$ will result in w having score larger than the score of d and this can never change afterwards. To be able to eliminate w in round $n+1$, however, we must ensure that the candidates $e_j, j \in [3n/2]$, and candidate q have at least the same score as w . This can only happen if the set $\{i \mid v_i \text{ is eliminated before round } n+1\}$ is a vertex cover. (Whenever some v_i is eliminated, the edge candidate corresponding to edges incident to vertex i gain one point and reach at least the same score as w .) Moreover the set $\{i \mid v'_i \text{ is eliminated before round } n+1\}$ must be of size at least $n-t$, ensuring that we have only selected t vertices to be part of the vertex cover. (Whenever some v'_i is eliminated, candidate q gains three points and needs in total at least $3(n-t)$ additional points to reach at least the same score as w in round $n+1$.)

Voters We have the following voters.

105n voters	$d \succ w \succ \dots$
99n voters	$w \succ d \succ \dots$
99n - 1 voters	$e_j \succ w \succ d \succ \dots \quad \forall j \in [m]$
99n - 3(n-t) voters	$q \succ w \succ d \succ \dots$
60n - 3 voters	$v_i \succ v'_i \succ w \succ d \succ \dots \quad \forall i \in [n]$
1 voter	$v_i \succ e_1^i \succ w \succ d \succ \dots \quad \forall i \in [n]$
1 voter	$v_i \succ e_2^i \succ w \succ d \succ \dots \quad \forall i \in [n]$
1 voter	$v_i \succ e_3^i \succ w \succ d \succ \dots \quad \forall i \in [n]$
60n - 3 voters	$v'_i \succ v_i \succ w \succ d \succ \dots \quad \forall i \in [n]$
3 voters	$v'_i \succ q \succ w \succ d \succ \dots \quad \forall i \in [n]$

This completes the construction.

Key Observations Observe that in round $\ell, \ell \in [n]$ either v_i or v'_i for some $i \in [n]$ must be eliminated, since these candidates have score $60n$, while every other candidate has score at least $99n - 3(n-t) > 96n$. If some vertex candidate is eliminated, then the respective other vertex candidate gains $60n - 3$ additional points and, hence, will not be eliminated before all edge

candidates. In round $n + 1$ either some edge candidate e_j (with score between $99n - 1$ and $99n + 1$), candidate q (with score between $96n$ and $102n$), or candidate w (with score $99n$) is eliminated. If any candidate different from w is eliminated in round $n + 1$, than w gains more than $95n$ additional points and will finally win the election. If candidate w is eliminated in round $n + 1$, then candidate d wins the election. Independent of whether we have eliminated candidate w in round $n + 1$ or we are still in round $n + 1$, next, all edge candidates and candidate q will be eliminated. Candidate w or, if w is eliminated than d receives the votes of all candidates eliminated after round n . If w is not eliminated, than d has still score $105n$ and will be eliminated next. Then, the remaining vertex candidates are eliminated and finally either w or d wins.

Correctness We show that d is a winner for Seq.-Plurality-Loser of the constructed profile if and only if graph G contains a vertex cover of size t .

For the “if”-part, assume that there is a vertex cover of size t . Without loss of generality, let the first t vertices denote such a vertex cover. To see that d is a winner of the election, consider the following elimination order. In round $\ell, \ell \in [t]$, eliminate candidate v_ℓ . In round $\ell, t + 1 \leq \ell \leq n$, eliminate candidate v'_ℓ . Now, each edge candidate has score at least $99n$ and also candidate q has score $99n$. Thus, we next eliminate w , then q and then the edge candidates in an arbitrary order that is consistent with their scores (some may have score $99n$ being covered once, some have score $99n + 1$ being covered twice). Finally eliminate the remaining vertex candidates (with scores between $120n - 6$ and $120n - 3$) so that only candidate d remains and wins. It is easy to verify that this elimination ordering is indeed consistent with the Plurality scores in the respective rounds.

For the “only if”-part, recall the idea and key observations. Assuming d wins, it must be that w is eliminated in round $n + 1$. To do this, $V^* = \{i \mid v_i \text{ is eliminated before round } n + 1\}$ must be a vertex cover, since each edge candidate must have gained at least one additional point in the first n rounds to have a score of $99n$ in round $n + 1$ where w is eliminated. Moreover, since candidate q also needs score at least $99n$ in round $n + 1$ to allow w to be eliminated in this round, it must hold that $\{i \mid v'_i \text{ is eliminated before round } n + 1\}$ is of size at least $n - t$. Thus, V^* is a vertex cover of size at most t . \square

C.2. Veto

Theorem 6.5. WINNER DETERMINATION for Sequential-Veto-Loser (aka. Coombs) is NP-complete. If the ETH is true, then the problem cannot be solved in $2^{o(m)} \cdot \text{poly}(n, m)$ time. [Main]

Proof. We reduce from REGULAR CLIQUE, i.e. CLIQUE restricted to regular graphs (where all vertices have the same degree). Our reduction will imply NP-hardness and also the claimed ETH-result by Corollary C.2. Let (G, k) be an instance of REGULAR CLIQUE, where $G = (V, E)$ is regular and each vertex has degree r . We construct a profile on candidate set $C = \{d, w\} \cup V \cup \{s_v : v \in V\}$ (the candidates from $\{s_v : v \in V\}$ act as dummy candidates). The question is whether d is a Coombs winner.

$$\begin{array}{ll}
k(k - 2) + r + 1 \text{ voters} & \dots \succ d \\
r + 1 \text{ voters} & \dots \succ w \\
k(k - 2) \text{ voters} & \dots \succ s_v \succ v \quad \text{for each } v \in V \\
1 \text{ voter} & \dots \succ d \succ v \quad \text{for each } v \in V \\
1 \text{ voter} & \dots \succ w \succ u \succ v \quad \text{for each } v \in V \text{ and each } u \in V \text{ with } \{v, u\} \in E
\end{array}$$

In these votes, “ \dots ” is replaced by all unmentioned candidates according to some common canonical order in which d is ranked first. To avoid doubt, for each edge $\{v, u\} \in E$, we introduce

two votes of the bottom type, one with $v \succ u$ and one with $u \succ v$. For convenience (to avoid talking about negative numbers), we say that the *bottom count* of a candidate is the number of times the candidate is ranked in last position. Thus the bottom count is the negative of the veto score, and Coombs proceeds by eliminating candidates with the highest bottom count. Throughout the proof, for a vertex $v \in V$, we write $\text{Nghbhd}(v) = \{u \in V : \{u, v\} \in E\}$ for the neighborhood of v in G .

The intuition behind the construction is that to avoid eliminating d (so as to make d a winner), we need to first eliminate w , despite its initially low bottom count. The way to increase its count is to eliminate vertices forming a dense subgraph, since then many ‘edge’ votes (of the bottom type) are transferred to w . The rule will start by eliminating vertices. Once some are eliminated, it is then only possible to eliminate vertices which are adjacent to all previously eliminated vertices. Thus, elimination sequences encode cliques.

We start by proving the forward direction. Suppose G contains a clique $T = \{v_1, \dots, v_k\}$ of size k . Note that in the initial profile, the bottom counts are:

- d has count $k(k - 2) + r + 1$,
- w has count $r + 1$,
- every $v \in V$ has count $k(k - 2) + r + 1$,
- every $s_v, v \in V$, has count 0.

Thus, we can eliminate v_1 . After that,

- d has count $k(k - 2) + r + 2$,
- w has count $r + 1$,
- every $v \in \text{Nghbhd}(v_1)$ has count $k(k - 2) + r + 2$,
- every $v \notin \text{Nghbhd}(v_1)$ has count $k(k - 2) + r + 1$,
- s_{v_1} has count $k(k - 2)$,
- every $s_v, v \neq v_1$, has count 0.

Thus, we can eliminate v_2 (since it is a neighbor of v_1). After this, since we have eliminated both endpoints of the edge $\{v_1, v_2\}$, the votes of the two voters corresponding to that edge are transferred to w .

In fact, we can eliminate candidates in the order v_1, \dots, v_k . After we have eliminated candidates $v_1, \dots, v_p, p \leq k$, we have the following situation.

- d has count $k(k - 2) + r + 1 + p$,
- w has count $r + 1 + p(p - 1)$,
- every $v \in V$ has count $k(k - 2) + r + 1 + |\{v_1, \dots, v_p\} \cap \text{Nghbhd}(v)|$,
- s_{v_1}, \dots, s_{v_p} have count $k(k - 2)$,
- every other s_v has count 0.

In particular, if $p < k$, then v_{p+1} can be eliminated in the next step, since its bottom count is $k(k - 2) + r + 1 + p$. After all members of the clique have been eliminated, we reach bottom counts

- d has count $k(k - 2) + r + 1 + k = k(k - 1) + r + 1$,
- w has count $r + 1 + k(k - 1)$,

- every $v \in V$ has count $k(k-2) + r + 1 + |\{v_1, \dots, v_k\} \cap \text{Nghbhd}(v)| \leq k(k-1) + r + 1$,
- s_{v_1}, \dots, s_{v_k} have count $k(k-2)$,
- every other s_v has count 0.

Hence, w is a candidate with highest bottom count, and thus now be eliminated. All its votes are transferred to the last (uneliminated) candidate in the canonical order, call it x . Because previously w had maximum bottom count, x now has maximum bottom count and can be eliminated. This argument applies repeatedly, and hence from now on the remaining candidates are eliminated in the canonical order, finishing with d (which was placed first in the canonical order), so d is a Coombs winner.

Conversely, suppose d is a Coombs winner. Thus, there is a way to eliminate candidates so that d is eliminated last. Now, at each of the first k elimination steps (assuming that d is eliminated last) the set of eliminateable candidates is a subset of $\{d\} \cup V$. To see this by induction, note that it holds at the first step. Further, after we have eliminated some candidates $W \subseteq V$ with $|W| = p < k$, writing e_W for the number of edges between vertices in W , the resulting bottom counts are

- d has count $k(k-2) + r + 1 + p$,
- w has count $r + 1 + 2 \cdot e_W$,
- every $v \in V$ has count $k(k-2) + r + 1 + |W \cap \text{Nghbhd}(v)|$,
- for $v \in W$, s_v has count $k(k-2)$,
- every other s_v has count 0.

Since $2 \cdot e_W \leq 2 \cdot \binom{p}{2} = p(p-1) < k(k-2) + p$ (since $p < k$), the bottom count of w is smaller than the score of d , so the candidate w cannot be eliminated in round p ; also clearly no candidate s_v can be eliminated. Thus, only a subset of $\{d\} \cup V$ can be eliminated at this step. In fact, we see that a vertex $v \in V \setminus W$ can be eliminated at this step if and only if v is adjacent to all candidates from W . It follows that if during the first k steps we do not eliminate d , then we eliminate k vertices. These vertices must form a clique in G . \square

Theorem 6.6. WINNER DETERMINATION for Sequential-Veto-Loser (aka. Coombs) is W[1]-hard with respect to the number n of voters. [Main]

Proof. We prove this statement by proving the hardness of an equivalent problem for Seq-Plurality-Winner: Specifically, we show that it is W[1]-hard parameterized by n to decide whether given a ranking profile P and a designated candidate d there is a selected ranking by Seq-Plurality-Winner where d is ranked last. Then by applying Lemma 3.5 the theorem follows.

In some round, we call the elimination of a candidate c *valid* if c is a Plurality winner in the election from this round.

We say that a candidate is *present* in some round if the candidate has not been deleted in some previous round.

For a set S and an element $s \in S$, we write $S - s$ as a shorthand notation for $S \setminus \{s\}$

Construction We prove hardness by a reduction from the W[1]-hard MULTICOLORED INDEPENDENT SET problem parameterized by the solution size ℓ . In MULTICOLORED INDEPENDENT SET, we are given a ℓ -partite graph $(V^1 \cup V^2 \cup \dots \cup V^\ell, E)$ and the question is whether there is an independent set X of size ℓ with $X \cap V^j \neq \emptyset$ for all $j \in [\ell]$. To simplify notation, we assume that $V^j = \{v_1^j, \dots, v_\nu^j\}$ for all $j \in [\ell]$. We refer to the elements of $[\ell]$ as *colors* and say that a vertex v has color $j \in [\ell]$ if $v \in V^j$.

$$\begin{aligned}
& g_1 \succ g_2 \succ g_3 \succ g_4 \succ [T] \succ c_\nu^1 \succ c_{\nu-1}^1 \succ \dots \succ c_1^1 \succ c_\nu^2 \succ \dots \succ c_1^2 \succ \dots \succ c_\nu^\ell \succ \dots \succ c_1^\ell \succ [B] \succ d \\
g_2 \succ g_3 \succ g_4 \succ [T] \succ c_1^1 \succ q_1^1 \succ c_2^1 \succ q_2^1 \succ \dots \succ c_\nu^1 \succ q_\nu^1 \succ c_1^2 \succ q_1^2 \succ \dots \succ c_\nu^2 \succ q_\nu^2 \succ \dots \succ c_1^\ell \succ q_1^\ell \succ \dots \succ c_\nu^\ell \succ q_\nu^\ell \succ [B] \succ d \\
& g_3 \succ g_4 \succ [T] \succ c_1^1 \succ q_1^1 \succ c_2^1 \succ q_2^1 \succ \dots \succ c_\nu^1 \succ q_\nu^1 \succ c_1^2 \succ q_1^2 \succ \dots \succ c_\nu^2 \succ q_\nu^2 \succ \dots \succ c_1^\ell \succ q_1^\ell \succ \dots \succ c_\nu^\ell \succ q_\nu^\ell \succ [B] \succ d \\
& g_4 \succ [T] \succ c_1^1 \succ q_1^1 \succ c_2^1 \succ q_2^1 \succ \dots \succ c_\nu^1 \succ q_\nu^1 \succ c_1^2 \succ q_1^2 \succ \dots \succ c_\nu^2 \succ q_\nu^2 \succ \dots \succ c_1^\ell \succ q_1^\ell \succ \dots \succ c_\nu^\ell \succ q_\nu^\ell \succ [B] \succ d
\end{aligned}$$

Figure 8: Global rankings from the construction for [Theorem 6.6](#).

For each $j \in [\ell]$ and $i \in [\nu]$, we introduce two *vertex candidates* c_i^j and q_i^j . Moreover, for each edge $e \in E$, we introduce an *edge candidate* f_e . Let $F = \cup_{e \in E} f_e$. For $j \in [\ell]$ and $i \in [\nu]$, let F_i^j be the set of all edge candidates corresponding to edges incident to v_i^j . Moreover, we introduce ℓ *blocker candidates* $B = \{b^1, \dots, b^\ell\}$. Lastly, we add *dummy candidates* $G = \{g_1, \dots, g_4\}$ and $T = \{t_1, \dots, t_4\}$, and add the designated candidate d .

For a subset C' of candidates, let $[C']$ be the lexicographic strict ordering of the candidates in C' . In particular let $[T]$ be $t_1 \succ t_2 \succ t_3 \succ t_4$, $[B]$ be $b^1 \succ \dots \succ b^\ell$, and $[B - b^j]$ be $b^1 \succ \dots \succ b^{j-1} \succ b^{j+1} \succ \dots \succ b^\ell$.

We now describe the ranking profile. We complete each ranking by appending the remaining candidates in an arbitrary order. First for each color $j \in [\ell]$, we introduce two *color rankings* and refer to them as the *first* and *second* color ranking:

$$\begin{aligned}
& c_1^j \succ [F_1^j] \succ q_1^j \succ c_2^j \succ [F_2^j] \succ q_2^j \succ \dots \succ \\
& c_\nu^j \succ [F_\nu^j] \succ q_\nu^j \succ b^j \succ [B - b^j] \succ d \\
q_\nu^j \succ [F_\nu^j] \succ c_\nu^j \succ q_{\nu-1}^j \succ [F_{\nu-1}^j] \succ c_{\nu-1}^j \succ \dots \succ \\
& q_1^j \succ [F_1^j] \succ c_1^j \succ b^j \succ [B - b^j] \succ d
\end{aligned}$$

Moreover, we introduce four *global rankings* as depicted in [Figure 8](#) (note that the last three rankings only “differ” in the beginning).

Lastly, we add five *dummy rankings* (these will be the rankings that contribute to the Plurality score of d at some point):

$$\begin{aligned}
& t_1 \succ t_2 \succ t_3 \succ t_4 \succ d \\
& t_2 \succ t_3 \succ t_4 \succ d \\
& t_3 \succ t_4 \succ d \\
& t_4 \succ d \\
& d
\end{aligned}$$

All rankings are completed arbitrarily. The general intuition is that we need to eliminate g_1 in some “early” round where all candidates still have a Plurality score of at most one, as g_1 is only ranked before d in one ranking. The elimination of g_1 then triggers immediately some uniquely determined follow-up eliminations (specifically of all candidates in $G \cup T$). After that d has a Plurality score of five. As d needs to be eliminated last, all candidates that are still present after this need to have at least Plurality score five when they are eliminated. As each edge candidate appears only in four rankings before d , this implies that all edge candidates must have already been deleted. Moreover, in Statement 5 of [Claim C.3](#), we prove that for each color $j \in [\ell]$, there is some $i \in [\nu]$ such that $c_{i_j}^j$ and $q_{i_j}^j$ are still present in the round where g_1 is eliminated. As all edge candidates must have been deleted before, from this one can conclude that $\{v_{i_1}^1, \dots, v_{i_\ell}^\ell\}$ is an independent set.

Forward Direction For the forward direction, assume that we are given a multicolored independent set $V' = \{v_{i_1}^1, \dots, v_{i_\ell}^\ell\}$.

We now describe a valid elimination order of the candidates for which d is eliminated last. We proceed in four phases.

In the first phase every candidate has Plurality score at most one in each round and it is thus valid to delete a candidate if it is ranked first in at least one ranking. For each $j \in [\ell]$, we do the following: We delete all candidates that are ranked before $c_{i_j}^j$ in the first color ranking for color j in the order in which they appear in this ranking. After that we delete all candidates that are ranked before $q_{i_j}^j$ in the second color ranking for color j in the order in which they appear in this ranking.

As V' is an independent set, after this, for each $j \in [\ell]$ the color rankings for color j are:

$$\begin{aligned} c_{i_j}^j &\succ q_{i_j}^j \succ b^j \succ [B - b^j] \succ d \\ q_{i_j}^j &\succ c_{i_j}^j \succ b^j \succ [B - b^j] \succ d \end{aligned}$$

At the end of phase one, every candidate has a Plurality score of at most one.

Subsequently in the second phase, we start by eliminating candidate g_1 . After that, for the next few rounds the elimination order is unique because one candidate is the unique Plurality winner. Specifically, afterwards we eliminate g_2 , then g_3 , then g_4 , then t_1 , then t_2 , then t_3 , and lastly t_4 . Afterwards, d is ranked in the first position in the five dummy rankings. The global rankings became:

$$c_{i_1}^1 \succ c_{i_2}^2 \succ \dots \succ c_{i_\ell}^\ell \succ [B] \succ d,$$

and three times:

$$c_{i_1}^1 \succ q_{i_1}^1 \succ c_{i_2}^2 \succ q_{i_2}^2 \succ \dots \succ c_{i_\ell}^\ell \succ q_{i_\ell}^\ell \succ [B] \succ d.$$

In the third phase, for $j = 1$ to $j = \ell$, we first eliminate $c_{i_j}^j$ and then $q_{i_j}^j$ (and subsequently increment j by one). To argue why all these eliminations are valid observe that in each round in this phase the Plurality score of all candidates is at most five: For d this follows directly from the fact that d appears after the blocker candidates in all but five rankings. For the blocker candidates, observe that in each round in the third phase at least one vertex candidate is present. Thus, no blocker candidate is ever ranked in first place in one of the last three global rankings. Ignoring the last three global rankings and all rankings where the blocker candidates appear after d , each blocker candidate appears in at most three rankings before all other blocker candidates. Thus, each blocker candidate has at most a Plurality score of three in some round in this phase. To see why vertex candidates have a Plurality score of at most five, fix some $j \in [\ell]$. For candidate $q_{i_j}^j$ the statement clearly holds, as $q_{i_j}^j$ only appears in five rankings before d . For candidate $c_{i_j}^j$, observe that $c_{i_j}^j$ appears before d in six rankings; however, in one of these rankings $c_{i_j}^j$ is ranked after $q_{i_j}^j$. As $q_{i_j}^j$ is eliminated after $c_{i_j}^j$, the Plurality score of $c_{i_j}^j$ is at most five in each round.

Using that the Plurality score of all candidates in each round in the third phase is upper bounded by five, we now argue why all eliminations are valid. For this, let us examine the situation for $j = 1$ (so the first round in the third phase): As $c_{i_1}^1$ is ranked first by the four global rankings and the first color ranking for color 1, $c_{i_1}^1$ has Plurality score five and eliminating it is valid. Subsequently, $q_{i_1}^1$ is ranked first by the last three global rankings and both color rankings for color 1. Thus, $q_{i_1}^1$ has Plurality score five and eliminating it is valid. The same argument also applies for increasing j , establishing the validity of this phase. After the third phase only blocker candidates and d remain.

In the fourth phase, we eliminate b^i for $i = 1$ to $i = \ell$. The designated candidate d has a Plurality score of five in each round of the fourth phase. Note that after the third phase, b^1 has six Plurality points and is thus the unique Plurality winner. After eliminating b^1 , b^2 has eight Plurality points and is thus the unique Plurality winner. The same reasoning applies until $i = \ell$. Afterwards, only d is left, which completes the argument.

Backward direction For the backward direction, assume that there is an execution of Seq.-Plurality-Winner such that d is eliminated in the last round. We will now reason about the elimination order of the candidates in this execution of Seq.-Plurality-Winner. Let x be the round in which candidate g_1 is eliminated. We will now prove a series of claims that are the cornerstone of the proof of correctness:

Claim C.3. *Assume that d is eliminated last and let x be the round in which g_1 is eliminated. Then,*

1. *Every candidate has at most Plurality score one in round x .*
2. *All candidates from $G \cup T$ are present in round x . Every candidate which is not part of $G \cup T$ and that is present in round x has a Plurality score of at least five in the round in which it is eliminated.*
3. *All candidates from F have been eliminated before round x .*
4. *For each $j \in [\ell]$, there are $i, i' \in [\nu]$ such that c_i^j and $q_{i'}^j$ are present in round x .*
5. *For each $j \in [\ell]$, there is some $i \in [\nu]$ such that both c_i^j and q_i^j are present in round x .*

Proof. **Proof of Statement 1.** As g_1 is only ranked once before d and d is eliminated last, g_1 has a Plurality score of one in round x . Thus, all candidates have a Plurality score of at most one in round x .

Proof of Statement 2. We prove the following from which Statement 2 directly follows, as d is the last candidate to be eliminated: In rounds x to $x + 8$ exactly the candidates $G \cup T$ are eliminated. The Plurality score of d is at least five in every round after round $x + 8$.

Assume that no candidate from $G \cup T$ has been eliminated before round x . Then, after g_1 is eliminated, in round $x + 1$ g_2 has Plurality score two and all other candidates have a Plurality score of at most one. Thus, g_2 will be eliminated. Following this reasoning, g_3 will be eliminated in round $x + 2$, g_4 in round $x + 3$, t_1 in round $x + 4$, t_2 in round $x + 5$, t_3 in round $x + 6$, and t_4 in round $x + 7$. Afterwards d is ranked first in the five dummy rankings.

For the sake of contradiction, assume that some candidate from $G \cup T$ is eliminated before round x . Let h be the first candidate from $G \cup T$ that is eliminated. Then in the round where h is eliminated h has a Plurality score of one, as in all rankings where h is not ranked first and appears before d it is ranked after some candidates from $G \cup T$. As the elimination of h distributes one Plurality point to some other candidate from $G \cup T$, the elimination of h triggers the above described elimination procedure from this candidate onwards ultimately leading to d gaining at least one Plurality point. Accordingly, d has a Plurality score of at least two in round x in this case, a contradiction to Statement 1.

Proof of Statement 3. Assume that for some $e \in E$, f_e is eliminated after round $x - 1$. Observe that each edge candidate appears only in four rankings before d (the four color rankings of colors of its endpoints). Thus, before d is eliminated the Plurality score of f_e is at most four, a contradiction to Statement 2 and d being eliminated last.

Proof of Statement 4. Fix some $j \in [\ell]$. We prove the statement in three steps by first excluding that no vertex candidate for color j is present, then excluding that only vertex candidates of the form q_i^j are present, and finally excluding that only vertex candidates of the form c_i^j are present. From these three parts the statement directly follows.

First, assume for the sake of contradiction that for all $i \in [\nu]$, c_i^j and q_i^j have been eliminated before round x . By this and Statement 3, it follows that either some blocker candidate or d is ranked first in the two color rankings for color j in round x . However, this implies that this candidate has Plurality score at least two in round x , a contradiction to Statement 1.

Second, assume for the sake of contradiction that for all $i \in [\nu]$ c_i^j has been eliminated before round x but that there is some $i' \in [\nu]$ such that $q_{i'}^j$ is present in round x .

We make a case distinction based on whether $q_{i'}^j$ is the only present vertex candidate for this color or not: We claim that if there is an $i'' \in [\nu] - i'$ such that $q_{i''}^j$ is present in round x , then the Plurality score of both $q_{i'}^j$ and $q_{i''}^j$ is at most four in any round before d is eliminated. To see this note that there are only five rankings in which $q_{i'}^j$ and $q_{i''}^j$ are ranked before d and that in fact these are the same five rankings for both. However, for the two color rankings for color j it holds that in one of them $q_{i'}^j$ is ranked before $q_{i''}^j$ and in the other $q_{i''}^j$ is ranked before $q_{i'}^j$. Thus, as long as both $q_{i'}^j$ and $q_{i''}^j$ are not eliminated, both of them can have at most four Plurality points. As $q_{i'}^j$ and $q_{i''}^j$ are present in round x , using Statement 2 it follows that none of $q_{i'}^j$ and $q_{i''}^j$ can reach a Plurality score of five before d 's elimination, a contradiction to d being eliminated last.

Otherwise for all $i'' \in [\nu] - i'$ $q_{i''}^j$ has been eliminated before round x (and by our initial assumption for all $i \in [\nu]$ c_i^j has also been eliminated before round x). Using Statement 3, it follows that $q_{i'}^j$ is ranked first in both color rankings for color j in round x , a contradiction to Statement 1.

Third, assume for the sake of contradiction that for all $i \in [\nu]$ q_i^j has been eliminated before round x but that there is some $i' \in [\nu]$ such that $c_{i'}^j$ has not been eliminated before round x . We make a case distinction similar as above. We claim that if there is a $i'' \in [\nu] - i'$ such that $c_{i''}^j$ is present in round x , then both $c_{i'}^j$ and $c_{i''}^j$ have at most four Plurality points in any round before d is eliminated. Assume without loss of generality that $i' < i''$. Both $c_{i'}^j$ and $c_{i''}^j$ are only ranked before d in the two color rankings for color j and the four global rankings. However, in four of these six rankings $c_{i'}^j$ is ranked before $c_{i''}^j$ (i.e., the first color ranking for color j and the last three global rankings), while in the other two $c_{i''}^j$ is ranked before $c_{i'}^j$. Thus, as long as $c_{i'}^j$, $c_{i''}^j$, and d are present, both $c_{i'}^j$ and $c_{i''}^j$ have at most four Plurality points. With the help of Statement 2, we again reach a contradiction to d being eliminated last.

Otherwise for all $i'' \in [\nu] - i'$, $c_{i''}^j$ has been eliminated before round x . Using Statement 2, it follows that $c_{i'}^j$ is ranked first in both color rankings for color j in round x , a contradiction to Statement 1.

Proof of Statement 5. Fix some $j \in [\ell]$. Let i' be the smallest i such that c_i^j is present in round x (from Statement 4 we know that such an i needs to exist). Assume for the sake of contradiction that there is some $i'' < i'$ such that $q_{i''}^j$ is present in round x . We argue that in this case neither $c_{i'}^j$ nor $q_{i''}^j$ can ever reach five Plurality points before d is eliminated, a contradiction to Statement 2 and d being eliminated last: The only rankings in which one of $c_{i'}^j$ and $q_{i''}^j$ appears before d are the two color rankings for color j and the four global rankings. However, in the first color ranking and the last three global rankings $q_{i''}^j$ appears before $c_{i'}^j$, whereas in the other two rankings $c_{i'}^j$ appears before $q_{i''}^j$. Thus, as long as the other candidate is present none of the two can get a Plurality score of five.

Finally, assume for the sake of contradiction that $q_{i'}^j$ is not present in round x , implying that it was eliminated in round y for some $y < x$. Together with our above observation that excludes the presence of $q_{i''}^j$ for all $i'' < i'$ in round x and Statement 4, this implies that there is some $i'' > i'$ such that $q_{i''}^j$ is present in round x . Clearly, $q_{i''}^j$ was ranked in first position by at least one ranking in round y . As all candidates from $G \cup T$ are still present in round x (Statement 2), $q_{i'}^j$ is not ranked first in one of the global rankings in round y . The only remaining two rankings where $q_{i'}^j$ appears before d are the two color rankings for color j . However, as $c_{i'}^j$ and $q_{i''}^j$ are present in round y and $q_{i'}^j$ is ranked behind $c_{i'}^j$ in the first color ranking and behind $q_{i''}^j$ in the second color ranking for color j , $q_{i'}^j$ is also not ranked first in these rankings in round y , a contradiction. \blacklozenge

By Statement 5 of [Claim C.3](#), we get that for each color $j \in [\ell]$, there is some $i_j \in [\nu]$ such that

$c_{i_j}^j$ and $q_{i_j}^j$ are present in round x . We claim that $V' = \{v_{i_1}^1, \dots, v_{i_\ell}^\ell\}$ is an independent set in the given graph. Assume for the sake of contradiction that there are two colors $j \neq j' \in [\ell]$ such that $e = \{v_{i_j}^j, v_{i_{j'}}^{j'}\} \in E$. From Statement 3 of [Claim C.3](#), we get that f_e has been eliminated before round x . The only rankings where f_e is ranked before d are the four color rankings for colors j and j' . In each of these rankings f_e is ranked between $c_{i_j}^j$ and $q_{i_j}^j$ or between $c_{i_{j'}}^{j'}$ and $q_{i_{j'}}^{j'}$. As all these four candidates are still present in round x , it follows that f_e was never ranked in the first position in some ranking before round x , a contradiction to f_e being eliminated before round x . Thus, V' is an independent set, which is clearly multicolored. \square

Theorem 6.7. POSITION- k DETERMINATION for Sequential-Veto-Loser is in XP with respect to the number n of voters. [Main]

Proof. For a ranking profile $P' = (\succ'_1, \dots, \succ'_{n'})$ over m' candidates, the *bottom-list* ($\text{cand}(\succ_1, m'), \dots, \text{cand}(\succ_{n'}, m')$) of P' contains for each voter its least preferred candidate.

Assume we are given a set C of m candidate and a ranking profile $P = (\succ_1, \dots, \succ_n)$ of n voters. Knowing for each round in an execution of Sequential-Veto-Loser the bottom-list is clearly sufficient for us to reconstruct the selected ranking. Even more, if we only know the bottom-list in some round, then we can reconstruct which candidates have been deleted in previous rounds. These are exactly the candidates that appear behind the currently bottom candidate of a voter in its original vote (note that they clearly need to be deleted and that not more candidates could have been deleted because they have never appeared in the last place and thus never had the lowest Veto score). To formalize this, for a tuple $\mathbf{x} = (x_1, \dots, x_n) \in C^n$, let $D(\mathbf{x})$ be the set of candidates that are ranked behind x_i for $i \in [n]$ in \succ_i , i.e., $D(\mathbf{x}) = \{c \in C \mid \exists i \in [n] : \text{pos}(x_i, \succ_i) < \text{pos}(c, \succ_i)\}$. Intuitively speaking, we need to delete all candidates from $D(\mathbf{x})$ to make \mathbf{x} a bottom-list of our profile. However, in case that $D(\mathbf{x})$ contains some of the candidates appearing in \mathbf{x} , \mathbf{x} will actually not be the bottom list of the resulting profile. Accordingly, we call \mathbf{x} *valid* if \mathbf{x} is the bottom list of $P|_{C \setminus D(\mathbf{x})}$.

Using this notation, we solve the problem via dynamic programming. For this we introduce a table $T[i, c, c_1, \dots, c_n]$ for $i \in [k]$ and $c, c_1, \dots, c_n \in C$. Moreover, we add a dummy cell $T[0, \emptyset, \text{cand}(\succ_1, m), \dots, \text{cand}(\succ_n, m)]$, which we set to true. An entry $T[i, c, c_1, \dots, c_n]$ is true if there is an execution of Sequential-Veto-Loser resulting in ranking \succ such that c is ranked in position i in \succ and the bottom list of the profile after round i is (c_1, \dots, c_n) .

For increasing $i \in [k]$, we fill table T by setting $T[i, c, c_1, \dots, c_n]$ to true if (c_1, \dots, c_n) is valid and there exists candidates $c', c'_1, \dots, c'_n \in C$ such that

- $D((c_1, \dots, c_n)) = D((c'_1, \dots, c'_n)) \cup \{c\}$
- no candidate appears more often than c among c'_1, \dots, c'_n , and
- $T[i-1, c', c'_1, \dots, c'_n]$ is true.

After T is filled, we can simply check whether there are candidates $c_1, \dots, c_n \in C$ such that $T[k, d, c_1, \dots, c_n]$ is true in which case we return yes and no otherwise. \square

C.3. Borda

We now prove that WINNER DETERMINATION for Sequential-Borda-Loser is NP-hard, even for a constant number of voters. We begin with a useful observation for reasoning about Sequential-Borda-Loser.

Remark C.4 (Weighted majority graph, C2-Borda scores). Every ranking profile induces a *weighted majority graph* (aka the C2-graph) which is an edge-weighted directed graph whose




Figure 9: A bilevel graph as described in Lemma C.5.

vertex set is the set of candidates, and for $c, d \in C$, the weight of the edge $c \rightarrow d$ is $w_{cd} = |\{i \in N : c \succ_i d\}| - |\{i \in N : d \succ_i c\}|$. Given a ranking profile and its induced weighted majority graph, the *C2-Borda score* of an alternative $c \in C$ is $\sum_{d \in C \setminus \{c\}} w_{cd}$. It is well-known that the Borda score of an alternative is an affine transformation of its C2-Borda score (indeed, the C2-Borda score is equivalent to the difference between the candidates' Borda score and the average Borda score of candidates). Hence, to obtain the output of Sequential-Borda-Loser at a given profile, we only need to know the profile's weighted majority graph.

In the construction of our reduction we build a weighted majority graph to reason about the scores of candidates. Using well-known results of McGarvey [1953] and Debord [1987], any arc-weighted digraph (in which all arc weights have the same parity; in our reductions we only ever use weights 0 and 2) can be realized as the weighted majority graph of a ranking profile, and this profile can be constructed in polynomial time. In some of our reductions, we will need to prove that a particular digraph can be constructed using a small number of rankings. For this, we use the following lemma:

Lemma C.5 (Erdős and Moser, 1964, Bachmeier et al., 2019). *We call an edge-weighted digraph $G = (V, A)$ bilevel if we can partition its vertex set as $V = (C_1 \cup \dots \cup C_s) \cup (D_1 \cup \dots \cup D_s)$, where all subsets are pairwise disjoint but some of them may be empty, such that*

$$A = (C_1 \times D_1) \cup \dots \cup (C_s \times D_s)$$

with all arcs having weight 2 (see Figure 9).

If G is bilevel, then it can be induced as the weighted majority graph of a 2-voter profile.

Proof. Consider the sets $C_1, \dots, C_s, D_1, \dots, D_s$ as sets that are linearly ordered in some arbitrary way. Then construct the following two rankings \succ_1 and \succ_2 :

$$\begin{aligned} C_1 \succ_1 D_1 \succ_1 C_2 \succ_1 D_2 \succ_1 \dots \succ_1 C_s \succ_1 D_s, \\ \text{rev}(C_s) \succ_2 \text{rev}(D_s) \succ_2 \dots \succ_2 \text{rev}(C_1) \succ_2 \text{rev}(D_1). \end{aligned}$$

It is easy to check that these two voters induce D as their weighted majority graph. □

We are now ready to prove the theorem.

Theorem 6.8. *Let $n \geq 8$ be a fixed even integer. Then WINNER DETERMINATION for Sequential-Borda-Loser (aka Baldwin), restricted to instances with exactly n voters, is NP-complete. In addition, if the ETH is true, then the problem cannot be solved in $2^{o(m)} \cdot \text{poly}(m)$ time.* [Main]

Proof. We will give the proof for the case $n = 8$. It can be extended to any larger even number of voters by repeatedly adding 2 opposite rankings (one the reverse of the other) to the profile constructed in the reduction. Adding opposite rankings does not change the induced weighted




Figure 10: An illustration of the reduction of [Theorem 6.8](#). All arcs have weight 2. Red superscripts denote the difference between the weight of outgoing and ingoing arcs for a candidate.

majority scores. Thus, the C2-Borda scores of the candidates do not change, and hence this does not change the result of Sequential-Borda-Loser.

We reduce from CUBIC VERTEX COVER. (Then the ETH-based claim follows using [Proposition C.1](#).) Let $G = (V, E)$ be a graph with q vertices where each vertex $v \in V$ is incident to exactly 3 edges, and let t be the target vertex cover size. We construct an instance of the WINNER DETERMINATION problem as follows.

The candidate set consists of one candidate for each vertex, one candidate for each edge, a designated candidate d , and dummy candidates $B = \{b_1, b_2, b_3, b_4\}$, $F = \{f_1, \dots, f_{q-t+8}\}$, $G = \{g_1\}$, $H = \{h_1, h_2, h_3, h_4\}$, and $K = \{k_1, \dots, k_5\}$. Let $C = V \cup E \cup \{d\} \cup B \cup F \cup G \cup H \cup K$.

The ranking profile will be constructed so as to induce a desired weighted majority graph, where all arcs will have weight 2. The arcs are as follows:

$$A = \{(e, v) \in E \times V : v \in e\} \cup (B \times V) \cup (G \times E) \cup (H \times E) \cup (F \times B) \cup (K \times H) \cup (H \times \{d\})$$

The constructed weighted majority graph is depicted in [Figure 10](#). We now describe how to write the weighted majority graph (C, A) as a sum of 4 bilevel graphs (as defined in [Lemma C.5](#)). For each vertex v , we label the three edges incident to it arbitrarily as e_v^1, e_v^2, e_v^3 . Consider the following arc sets:

$$\begin{aligned} A_1 &= (B \times V) \cup ((H \cup G) \times E), \\ A_2 &= \{(e_v^1, v) : v \in V\} \cup (F \times B), \\ A_3 &= \{(e_v^2, v) : v \in V\} \cup (K \times H), \\ A_4 &= \{(e_v^3, v) : v \in V\} \cup (H \times \{d\}). \end{aligned}$$

It is clear that each of these sets describe bilevel graphs, that they are pairwise disjoint, and that $A = A_1 \cup A_2 \cup A_3 \cup A_4$. By invoking [Lemma C.5](#), we get a profile P containing 8 voters with the depicted weighted majority graph.

The C2-Borda scores (see [Remark C.4](#)) in this profile are:

- d has score -8 (since it is beaten by 4 H -candidates)
- each $b \in B$ has score $2t - 16$
- each $v \in V$ has score -14 (since v is beaten by 3 edge candidates, and by 4 B -candidates)

- each $e \in E$ has score -6 (since e beats 2 vertex candidates but is beaten by 4 H -candidates and 1 G -candidate).
- each $h \in H$ has score $2|E| - 8$ (since it beats $|E|$ edge candidates and d , but is beaten by 5 K -candidates)
- Candidates in F , G , and K will always have non-negative scores (since they are beaten by no candidates)

Suppose that $T = \{v_1, \dots, v_t\}$ is a vertex cover of G . Then the following is a valid elimination ordering according to Seq.-Borda-Loser, where d is eliminated last and thus ranked first in the corresponding selected ranking.

- Eliminate T : In the first t rounds, vertex candidates have the lowest C2-Borda score (-14), so we can eliminate T in an arbitrary ordering. Each time we eliminate a vertex candidate, the score of the B -candidates goes down by 2, and the score of all incident edge candidates also goes down by 2.
- Eliminate B : Starting in round $t + 1$, candidates in B have the uniquely lowest C2-Borda score of -16 , so we can eliminate B in an arbitrary ordering. As we do so, the scores of the remaining vertex candidates go up.
- Eliminate E : The remaining vertex candidates in $V \setminus T$ currently have score -6 because B has been eliminated. Edge candidates have a C2-Borda score of either -8 or -10 depending on whether T contains one or both of the endpoints of the edge (note that for each edge at least one endpoint has been deleted because T is a vertex cover). These are the lowest C2-Borda scores, so we can eliminate all edge candidates (in an arbitrary ordering except that the -10 edges get eliminated first). As we eliminate edge candidates, the scores of H -candidates go down.
- Eliminate H : Just after we finish eliminating E , the score of H has dropped to -8 , which is the lowest C2-Borda score. So we can eliminate H in an arbitrary order.
- Eliminate $F \cup G \cup K$: At this point, all remaining candidates have C2-Borda score 0 (and they have no arcs between them, so eliminating candidates does not change the scores). So we can eliminate all candidate except d in an arbitrary order.
- Eliminate d .

Conversely, suppose there is a ranking selected by Seq.-Borda-Winner where d is eliminated last. In the first $t - 1$ rounds, only vertex candidates can be eliminated since only they have the lowest C2-Borda score of -14 . But then, in round t , the B -candidates also have score -14 . We distinguish two cases: whether another vertex is eliminated, or whether a B -candidate is eliminated.

- Case 1: In round t , a vertex candidate is eliminated. Define T to be the set of vertices whose candidates were eliminated in the first t rounds. Starting in round $t + 1$, the B -candidates have score -16 , which is uniquely lowest, so they will then all be consecutively eliminated.
- Case 2: In round t , a B -candidate is eliminated. Then the C2-Borda score of the remaining vertex candidates goes up to -12 while the remaining B -candidates have score -14 , which makes them the candidates with the uniquely lowest C2-Borda score. Thus, in the following rounds, all the B -candidates will be consecutively eliminated. Define T to be the set of vertices whose candidates were eliminated in the first $t - 1$ rounds.

In either of the two cases, we have now reached a stage where for a set T of at most t vertices (to be precise, either t or $t - 1$ vertices), the corresponding vertex candidates have been eliminated, followed by the elimination of set B . We will prove that T is a vertex cover.

After the above-described eliminations, since B is eliminated, the remaining vertex candidates in $V \setminus T$ have score -6 . Candidate d has score -8 . Edge candidates have score -6 if neither of their endpoints were contained in T , and otherwise they have a score of -8 or -10 . All other candidates have score -6 or higher. Thus, in the next rounds, edges that were covered have their candidates eliminated. While this happens, the score of H -candidates goes down, but they do not become eliminateable before all covered edge candidates are eliminated. If in this way all edge candidates are eliminated, then T was a vertex cover and we are done. If there is an edge that is not covered by T , then after all covered edge candidates are eliminated, we end up in a situation where d has score -8 , but all other candidates have score -6 or higher, so d would need to be eliminated next, a contradiction. \square

D. Additional Material for Section 6.3

We start by observing that for each scoring system \mathbf{s} our general problem POSITION- k DETERMINATION is in XP, as we can simply guess which candidates are ranked on the first k positions (in which ordering) in the selected ranking and then verify whether it gives rise to a valid execution of Sequential- \mathbf{s} -Winner.

Observation D.1. *For every scoring system \mathbf{s} , POSITION- k DETERMINATION for Sequential- \mathbf{s} -Winner is in XP.*

D.1. Plurality

In this section, we consider Seq.-Plurality-Winner. We start by showing that TOP- k DETERMINATION is NP-hard and W[1]-hard with respect to k .

Proposition D.2. *TOP- k DETERMINATION for Sequential-Plurality-Winner is NP-hard and W[1]-hard with respect to k .*

Proof. We reduce from INDEPENDENT SET, which is W[1]-hard when parameterized by the solution size. Given a graph $G = (V, E)$ with $|V| = \nu$, and an integer ℓ , INDEPENDENT SET asks whether there is an independent set of size ℓ in G (we assume without loss of generality that $\nu > 2$, $\ell > 2$, $|E| > 1$, and $\ell < \nu$). From an instance of INDEPENDENT SET, we construct an instance of TOP- k DETERMINATION as follows. We add a candidate c_v for each vertex $v \in V$. Moreover, we introduce $2\nu^3(\ell - 1) + \nu^3$ dummy candidates, a designated candidate d , a blocker candidate b , and an edge candidate e .

We now turn to the description of the ranking profile. We first add the following rankings:

$$\begin{array}{lll} 2\nu^2 \text{ voters} & c_v \succ b \succ d \succ \dots & \forall v \in V \\ 2\nu^2 \text{ voters} & c_v \succ e \succ \dots & \forall v \in V \\ 1 \text{ voter} & c_v \succ b \succ c_w \succ e \succ \dots & \forall \{v, w\} \in E \end{array}$$

Moreover, for each $v \in V$, as long as c_v has less than $2\nu^2(\ell - 1) + \nu^2$ Plurality points in the current profile, we add a ranking where c_v is ranked first and some so far never in second place appearing dummy candidate is ranked second. We set $k := \ell + 2$.

Note that dummy candidates will clearly not be eliminated in the first k rounds so we can ignore them.

Assume that $V = \{v_1, \dots, v_\ell\}$ is an independent set. Then, we eliminate in the first $\ell - 1$ rounds candidates $c_{v_1}, \dots, c_{v_{\ell-1}}$. Note that by this eliminations only the Plurality scores of

dummy candidates, of b , and of e have changed. After round $\ell - 1$, e has score $2\nu^2(\ell - 1)$. Moreover, b has score at most $2\nu^2(\ell - 1) + (\ell - 1) \cdot \nu$. Thus, the vertex candidates still have the highest score and we can eliminate c_{v_ℓ} . After round ℓ , candidate b clearly has the highest Plurality score, so we eliminate b . The elimination of b redistributes $2\ell\nu^2$ points to d , at most ν points to each vertex candidates, and no point to e (because V is an independent set and thus in each of the voters of the third type at least one vertex candidate ranked before e is still present). Thus, we can eliminate d in the next round $\ell + 2$.

Conversely, assume that there is an execution of Seq.-Plurality-Winner such that d is eliminated in round $\ell + 2$ or before. As argued before clearly in the first ℓ rounds vertex candidates need to be deleted. Let $V' \subseteq V$ be the subset of vertices corresponding to these vertex candidates. We claim that V' needs to be an independent set. As argued above, b will be eliminated in round $\ell + 1$. Thus, d has a Plurality score of $2\ell\nu^2$ in round $\ell + 2$. In round $\ell + 2$, candidate e is also ranked in the first position in the $2\ell\nu^2$ votes of the second type. Moreover, if there is an edge $\{v, u\} \in E$ with $v, u \in V'$, then we have eliminated all candidates ranked before e in the corresponding vote, giving e a Plurality score of at least $2\ell\nu^2 + 1$. This leads to a contradiction, as this implies that e has a higher score than d in round $\ell + 2$. \square

Moreover, also when we parameterized by the number of n of voters we still get W[1]-hardness by reusing some ideas for the reduction for WINNER DETERMINATION for Seq.-Veto-Loser from [Theorem 6.6](#).

Theorem D.3. TOP- k DETERMINATION for Sequential-Plurality-Winner is W[1]-h. with respect to the number n of voters.

Proof. We prove hardness by a reduction from the W[1]-hard MULTICOLORED INDEPENDENT SET problem parameterized by the solution size ℓ .

Construction In MULTICOLORED INDEPENDENT SET, we are given an ℓ -partite graph $(V^1 \cup V^2 \cup \dots \cup V^\ell, E)$ and the question is whether there is an independent set X of size ℓ with $X \cap V^j \neq \emptyset$ for all $j \in [\ell]$. To simplify notation, we assume that $V^j = \{v_1^j, \dots, v_\nu^j\}$ for all $j \in [\ell]$. We refer to the elements of $[\ell]$ as *colors* and say that a vertex v has color $j \in [\ell]$ if $v \in V^j$. Moreover, let $|E| = \mu$.

We construct an instance of our problem by setting $k := \mu + \ell \cdot (\nu - 1) + 1$. We start by describing the candidate set. For each $j \in [\ell]$ and $i \in [\nu + 1]$, we introduce a *vertex candidates* c_i^j (notably candidates $c_{\nu+1}^j$ do not correspond to a vertex but act more like a dummy candidate). Moreover, for each edge $e \in E$, we introduce an *edge candidate* f_e . For $j \in [\ell]$ and $i \in [\nu]$, let F_i^j be the set of all edge candidates corresponding to edges incident to v_i^j . Further, for $j \neq j' \in [\ell]$, let $F^{j,j'}$ be the set of all edge candidates corresponding to edges between a vertex of color j and a vertex of color j' (note that there are no edges between vertices of the same color). Moreover, we introduce k *blocker candidates* $B = \{b_1, \dots, b_k\}$. Lastly, we add our designated candidate d . For a subset C' of candidates, let $[C']$ be the lexicographic strict ordering of the candidates in C' .

Turning to the input rankings, we introduce for each color $j \in [\ell]$, $\binom{\ell}{2}$ copies of two types of *color rankings*:

$$\begin{aligned} c_1^j \succ [F_1^j] \succ c_2^j \succ [F_2^j] \succ \dots \succ c_\nu^j \succ [F_\nu^j] \succ c_{\nu+1}^j \succ [B] \succ d, & \quad \forall t \in [\binom{\ell}{2}] \\ c_{\nu+1}^j \succ [F_\nu^j] \succ c_\nu^j \succ [F_{\nu-1}^j] \succ c_{\nu-1}^j \succ \dots \succ [F_1^j] \succ c_1^j \succ [B] \succ d. & \quad \forall t \in [\binom{\ell}{2}] \end{aligned}$$

Moreover, for each pair of colors $j \neq j' \in [\ell]$, we introduce the following *critical rankings*:

$$[F^{j,j'}] \succ d \succ [B].$$

We complete all rankings arbitrarily.

Forward Direction For the forward direction, assume that we are given a multicolored independent set $V' = \{v_{i_1}^1, \dots, v_{i_\ell}^\ell\}$.

We now describe a valid elimination order of the candidates for which d is eliminated in round k . For $j = 1$ to $j = \ell$, we do the following: We delete all candidates that are ranked before $c_{i_j}^j$ in the first type of color rankings for color j in the order in which they appear in this type of ranking. Afterwards, we delete all candidates that are ranked before $c_{i_{j+1}}^j$ in the second type of color ranking for color j in the order in which they appear in this type of ranking. Thus, all candidates ranked before the blocker candidates in the color rankings, except $c_{i_j}^j, c_{i_{j+1}}^j$, and candidates from $F_{i_j}^j$ got deleted. Thus, as V' is an independent set, all edge candidates got deleted. As for each color two vertex candidates remain, this implies that $\mu + \ell \cdot (\nu - 1) = k - 1$ candidates got deleted. The resulting profile looks as follows:

$$\begin{aligned} c_{i_j}^j &\succ c_{i_{j+1}}^j \succ [B] \succ d, \quad \forall j \in [\ell], t \in \binom{[\ell]}{2} \\ c_{i_{j+1}}^j &\succ c_{i_j}^j \succ [B] \succ d, \quad \forall j \in [\ell], t \in \binom{[\ell]}{2} \\ d &\succ [B], \quad \forall j \neq j' \in [\ell] \end{aligned}$$

Notably, d is a Plurality winner in this profile so we can eliminate d in round k .

Backward Direction For the backward direction, assume that there is an execution of Seq.-Plurality-Winner for which d is ranked in one of the first k positions in the selected ranking. Let $k^* \leq k$ be the round in which d is eliminated. Notably, as there exist $\binom{[\ell]}{2}$ identical rankings, the Plurality winner in each round, needs to have Plurality score at least $\binom{[\ell]}{2}$. Moreover, as d only appears in the $\binom{[\ell]}{2}$ critical rankings in one of the first k^* positions, this implies that d needs to be ranked first in all critical rankings in round k^* and that all candidates have a Plurality of score at most $\binom{[\ell]}{2}$ in round k^* . From the first part of this observation it follows that all edge candidates need to be deleted before round k^* . Using this and that all candidates have Plurality score at most $\binom{[\ell]}{2}$ in round k^* , we get that for each color $j \in [\ell]$, at least two vertex candidates corresponding to vertices of this color are present in round k^* : If there is a color with no remaining vertex candidates in round k^* , then some blocker candidate will be ranked in the first position in the $2\binom{[\ell]}{2}$ rankings corresponding to this color in round k^* . If there is a color with only one vertex candidate from this color remaining, then this candidate is ranked first in the first position in the $2\binom{[\ell]}{2}$ rankings corresponding to this color in round k^* .

Now, for each color $j \in [\ell]$, let i_j be the smallest i such that c_i^j is present in round k^* and let $t_j \neq i_j$ be some other index such that $c_{t_j}^j$ is present in round k^* (by our above observation both of these need to exist). We claim that $\{v_{i_1}^1, \dots, v_{i_\ell}^\ell\}$ is an independent set in the given graph. Assume for the sake of contradiction that there are $j \neq j' \in [\ell]$ with $\{v_{i_j}^j, v_{i_{j'}}^{j'}\} = e \in E$. We claim that in this case f_e has not been eliminated before round k^* : Recall that a candidate can only get eliminated if its Plurality score is at least $\binom{[\ell]}{2}$. Moreover, as f_e only appears in one of the first k^* positions in one of the critical rankings, this means that in the round in which f_e is eliminated it needs to be ranked in the first position in one of the color rankings for either color j or j' . However, note that in all of these rankings f_e is either ranked between $c_{i_j}^j$ and $c_{t_j}^j$ or between $c_{i_{j'}}^{j'}$ and $c_{t_{j'}}^{j'}$. As all these four candidates are still present in round k^* (and thus also in all previous rounds), it follows that f_e was never ranked in the first position in one of the color rankings, a contradiction to all edge candidates being deleted before round k^* . \square

From [Theorem 6.7](#), using the equivalence between Seq.-Loser and Seq.-Winner from [Lemma 3.5](#), we can directly conclude the following:

Corollary D.4. POSITION- k DETERMINATION for Sequential-Plurality-Winner is in XP with respect to the number n of voters.

Moreover, if we combine the two parameter n and k , for each of which we have proven W[1]-hardness, then we can obtain fixed-parameter tractability. The core observation here is that in the first k rounds of an execution of Seq.-Plurality-Winner at most $n \cdot k$ candidates can have a non-zero Plurality in one of the rounds (these are the candidates which are ranked in one of the first k positions of some voter). All other candidates will be eliminated immediately (in some arbitrary order). Subsequently we can apply our algorithm from [Theorem 6.1](#).

Observation D.5. POSITION- k DETERMINATION for Sequential-Plurality-Winner is solvable in $\mathcal{O}(2^{nk} \cdot nm^2)$ time.

D.2. Veto

We now turn to Seq.-Veto-Winner and start by proving that TOP- k DETERMINATION is NP-hard and W[2]-hard with respect to k for this rule.

Proposition D.6. TOP- k DETERMINATION for Sequential-Veto-Winner is NP-hard and W[2]-hard with respect to k .

Proof. We reduce from HITTING SET, where given a universe U and a family of sets \mathcal{S} and an integer ℓ , the question is whether there is an ℓ -subset of the universe containing at least one element from each set from \mathcal{S} , i.e., $U' \subseteq U$ with $|U'| = \ell$ and $S \cap U' \neq \emptyset$ for all $S \in \mathcal{S}$. HITTING SET is W[2]-hard when parameterized by ℓ . Let $\nu := |U|$ and $\mu := |\mathcal{S}|$. For an element $u \in U$, let \mathcal{S}_u denote the family of sets in which u appears. That is, $\mathcal{S}_u = \{S \in \mathcal{S} \mid u \in S\}$.

We construct an instance of TOP- k DETERMINATION as follows. For each element $u \in U$, we introduce an *element candidate* c_u . For each set $S \in \mathcal{S}$, we introduce a *set candidate* e_S . Lastly, we introduce a *blocker candidate* b and the designated candidate d . We set $k := \ell + 1$.

For convenience (to avoid talking about negative numbers), we say that the *bottom count* of a candidate is the number of times the candidate is ranked in last position. Thus the bottom count is the negative of the veto score, and Sequential-Veto-Winner proceeds by eliminating candidates with the lowest bottom count.

We now turn to the description of the ranking profile. We first add the following rankings:

$$\begin{aligned} \dots \succ b \succ e_S \succ c_u, & \quad \forall u \in U \text{ and } S \in \mathcal{S}_u \\ \dots \succ b \succ c_{u'} \succ c_u, & \quad \forall u \in U \text{ and } u' \in U \setminus \{u\} \\ \dots \succ b \succ e_S, & \quad \forall S \in \mathcal{S} \text{ and } i \in [\nu + \mu + \ell - 1] \\ \dots \succ d, & \quad \forall i \in [\nu + \mu + \ell] \\ \dots \succ b, & \quad \forall i \in [\nu + \mu + \ell + 1] \end{aligned}$$

Note that with these rankings, for each element $u \in U$, the candidate c_u has bottom count at most $\nu + \mu$. For each element $u \in U$, we add several copies of the ranking $\dots \succ b \succ c_u$ until c_u has bottom count exactly $\nu + \mu$. Thus, the bottom counts of the candidates in the initial profile are as follows:

- For each $u \in U$, c_u has a count of $\nu + \mu$.
- For each $S \in \mathcal{S}$, e_S has a count of $\nu + \mu + \ell - 1$.
- Candidate b has a count of $\nu + \mu + \ell + 1$.
- The designated candidate d has a count of $\nu + \mu + \ell$.

Initially, all element candidates c_u are Veto winners, having the lowest bottom count. Note that the bottom count of uneliminated candidates can only increase over time. Thus, for d to be a Veto winner in some round, each element candidate needs to be either deleted or ranked last in at least ℓ additional rankings, and also each set candidate needs to be either deleted or ranked last in at least one additional vote.

Let $U' = \{u_1, \dots, u_\ell\}$ be a hitting set of size ℓ . Then, in the first round we eliminate c_{u_1} , increasing the count of each other element candidate by 1, and increasing the score of each set candidate corresponding to a set from \mathcal{S}_{u_1} by 1. All remaining element candidates are still veto winners and we continue eliminating c_{u_i} for $i = 2, \dots, \ell$. After round ℓ , each remaining element candidate has count $\nu + \mu + \ell$. Moreover, as U' is a hitting set, each set candidate also has count at least $\nu + \mu + \ell$. This means that d is a Veto winner in round $\ell + 1 = k$ and we can eliminate it.

Conversely, assume that there is an execution of Seq.-Veto-Winner such that d is eliminated in round $\ell + 1$ or earlier. As each element candidate either needs to be ranked last in ℓ additional rankings or deleted, in the first ℓ rounds element candidates need to be eliminated. Let $U' \subseteq U$ be the subset of elements that correspond to the eliminated element candidates. Then in case U' does not form a hitting set, there is a set candidate that still has count $\nu + \mu + \ell - 1$ in round $\ell + 1$, and thus in particular a lower count than the designated candidate d , a contradiction. \square

By applying [Lemma 3.5](#) to [Observation 6.4](#), we get that POSITION- k DETERMINATION for Seq.-Veto-Winner is fixed-parameter tractable with respect to n .

Corollary D.7. POSITION- k DETERMINATION for Sequential-Veto-Winner is solvable in $\mathcal{O}(2^n \cdot nm^2)$ time.

D.3. Borda

We conclude by studying Seq.-Borda-Winner. Recall [Remark C.4](#) which showed that it suffices to reason about the weighted majority graph induced by a profile.

Theorem D.8. TOP- k DETERMINATION for Sequential-Borda-Winner is NP-complete for $n = 8$.

Proof. We reduce from CUBIC VERTEX COVER. Let $G = (V, E)$ be a graph with q vertices where each vertex $v \in V$ is incident to exactly 3 edges, and let t be the target vertex cover size. We construct an instance of the TOP- k DETERMINATION problem as follows.

The candidate set consists of one candidate for each vertex, one candidate for each edge, a designated candidate d , and dummy candidates f , candidates $B = \{b_1, \dots, b_6\}$, and candidates $H = \{h_1, \dots, h_{q+3-t}\}$. Let $C = V \cup E \cup \{d\} \cup \{f\} \cup B \cup H$. We set $k = t + 7$. The ranking profile will be constructed so as to induce a desired weighted majority graph, where all arcs will have weight 2. The arcs are as follows:

$$A = \{(e, v) \in E \times V : e \in E, v \in e\} \cup \{(d, f)\} \cup (V \times B) \cup (B \times H).$$

The constructed weighted majority graph is depicted in [Figure 11](#). We now describe how to write the weighted majority graph (C, A) as a sum of 4 bilevel graphs (as defined in [Lemma C.5](#)). For each vertex v , we label the three edges incident to it arbitrarily as e_v^1, e_v^2, e_v^3 . Consider the following arc sets:

$$\begin{aligned} A_1 &= (V \times B), \\ A_2 &= \{(e_v^1, v) : v \in V\} \cup (B \times H), \\ A_3 &= \{(e_v^2, v) : v \in V\} \cup \{(d, f)\}, \\ A_4 &= \{(e_v^3, v) : v \in V\}. \end{aligned}$$




Figure 11: An illustration of the reduction of [Theorem D.8](#). All arcs have weight 2. Red superscripts denote the difference between the weight of outgoing and ingoing arcs for a candidate.

It is clear that each of these sets describe bilevel graphs, that they are pairwise disjoint, and that $A = A_1 \cup A_2 \cup A_3 \cup A_4$. By invoking [Lemma C.5](#), we get a profile P containing 8 voters with the depicted weighted majority graph.

The C2-Borda scores (for a definition, see [Remark C.4](#)) in this profile are:

- d has score 2
- each $b \in B$ has score $6 - 2t$
- each $v \in V$ has score 6 (since v is incident to 3 edges, and beats 6 b -candidates)
- each $e \in E$ has score 4 (since e is incident to 2 vertices).
- each $h \in H$ as well as f have negative scores (and can only have non-positive scores throughout the elimination process because they do not beat any candidates) and will not be selectable in the first k rounds.

Suppose that $T = \{v_1, \dots, v_t\}$ is a vertex cover of G . Then the following is a valid start of an elimination ordering, with d eliminated in round $k = t + 7$: First, we eliminate all candidates from T in some arbitrary ordering, then all candidates from B in some arbitrary ordering, and then d . Explicitly, in the first t rounds, the maximum Borda score of a candidate is 6 and all vertex candidates have a Borda score of 6 (no other candidates have a Borda score of 6 in these rounds). Thus we can select members of the vertex cover T in each of these rounds. In round $t + 1$, the remaining vertex candidates and the candidates in B have the maximum Borda score of 6. Thus, we can eliminate candidates in B (while doing so, the Borda scores of vertex candidates decrease, so candidates in B continue having the maximum Borda score). After all 6 candidates in B are eliminated, we are in the following situation with respect to the remaining candidates' C2-Borda score:

- d has score 2
- each e has score at most 2 (since we have eliminated a vertex cover, and thus have eliminated at least one candidate that e beats)
- each remaining $v \in V$ has score -6 ; each $h \in H$ and f have non-positive scores.

Hence, at this point, candidate d has the highest Borda score and can be eliminated.

Conversely, suppose there is a ranking selected by Sequential-Borda-Winner where d is eliminated in round $k = t + 7$ or earlier. As observed above, in the first t rounds, only vertex




Figure 12: An illustration of the reduction of [Theorem D.9](#). All arcs have weight 2. Red superscripts denote the C2-Borda score of the candidates (after adding dummy candidates).

candidates can be eliminated. Let $T = \{v_1, \dots, v_t\}$ be the set of vertices whose candidates are eliminated in these rounds. From round $t + 1$ until $t + 6$, all candidates in B have score at least 6 (it cannot go lower because the candidates that B beats, namely H , cannot be eliminated). Because d has score only 2, all the 6 candidates in B are eliminated before d . This brings us to round $t + 7$ where by assumption d is eliminated. Hence at this point, no candidate has score higher than 2. In particular, for every edge $e \in E$, its score is less than 4. This can only have happened if at least one of the vertices incident to e has been eliminated and is thus part of T . It follows that T is a vertex cover. \square

Theorem D.9. TOP- k DETERMINATION for Sequential-Borda-Winner is $W[2]$ -hard with respect to k .

Proof. We reduce from HITTING SET, using a similar construction as in [Theorem D.8](#). Let U be a given universe of elements and let \mathcal{S} be a given family of subsets of U . We are also given an integer t , and the question is whether there is a t -subset of the universe containing at least one element from each set from \mathcal{S} , i.e., $U' \subseteq U$ with $|U'| = t$ and $S \cap U' \neq \emptyset$ for all $S \in \mathcal{S}$. Let $q = |U|$.

We construct an instance of the TOP- k DETERMINATION problem as follows.

We first give an incomplete description of the constructed instance. Later, we will add some dummy candidates that have no influence except that they increase the Borda scores of some of the candidates to a desired level. The candidate set consists of one candidate for each element, one candidate for each set $S \in \mathcal{S}$, a designated candidate d , and a set $B = \{b_1, b_2\}$ of two blocking candidates. Let $C_{\text{base}} = V \cup E \cup \{d\} \cup B$ (again, we will add to this set later). The ranking profile will be constructed so as to induce a desired weighted majority graph, where all arcs will have weight 2 (using standard arguments; [McGarvey, 1953](#), [Debord, 1987](#)). The arcs are as follows:

$$A_{\text{base}} = \{(S, e) \in \mathcal{S} \times U : e \in S\} \cup (U \times B)$$

The constructed weighted majority graph is depicted in [Figure 12](#).

The C2-Borda scores (for a definition, see [Remark C.4](#)) in this profile are:

- d has score 0
- each $S \in \mathcal{S}$ has score $2|S|$
- each $e \in U$ has score $2|\{S : e \in S\}| - 4$
- each $b \in B$ has score $-2|U|$.

Choose α to be an even integer such that the number $\alpha + 2 - 2t$ is larger than all the Borda scores just mentioned. (Clearly we can take an α that is polynomial size.)

Now we go through every candidate $c \in C_{\text{base}}$ and add dummy candidates D_c and arcs $\{c\} \times D_c$ to increase the C2-Borda score so that we now have the following scores (see [Figure 12](#)):

- d has score $\alpha - 2$
- each $S \in \mathcal{S}$ has score α
- each $e \in U$ has score $\alpha + 2$
- each $b \in B$ has score $\alpha + 2 - 2t$.

Explicitly, if a candidate's target C2-Borda score is R and its current score is $S < R$, then we need to add $|D_c| = (R - S)/2$ dummy candidates. Thus, the complete candidate set is $C = C_{\text{base}} \cup \bigcup_{c \in C_{\text{base}}} D_c$ and the final arc set is $A = A_{\text{base}} \cup \bigcup_{c \in C_{\text{base}}} (\{c\} \times D_c)$. This weighted majority graph can be realized by a polynomial-size ranking profile [[McGarvey, 1953](#)].⁷ We set $k = t + 3$.

We now prove that our reduction is correct. Note that throughout the argument, all the dummy candidates have non-positive C2-Borda score and it will never be possible to eliminate any of them, so they can essentially be ignored.

Suppose that $T = \{e_1, \dots, e_t\}$ is a hitting set. Then the following is a valid start of an elimination ordering, with d eliminated in round $k = t + 3$: First, we eliminate all candidates from T in some arbitrary ordering, then the two candidates from B in some arbitrary ordering, and then d . Explicitly, in the first t rounds, the maximum Borda score of a candidate is $\alpha + 2$ and all element candidates have a Borda score of $\alpha + 2$ (no other candidates have a Borda score of $\alpha + 2$ in these rounds). Thus we can select members of the hitting set T in each of these rounds. In round $t + 1$, the remaining element candidates and the candidates in B have the maximum Borda score of $\alpha + 2$. Thus, we can eliminate candidates in B (while doing so, the Borda scores of element candidates decrease, so candidates in B continue having the maximum Borda score). After all 2 candidates in B are eliminated, we are in the following situation with respect to the remaining candidates' C2-Borda score:

- d has score $\alpha - 2$
- each $S \in \mathcal{S}$ has score at most $\alpha - 2$ (since we have eliminated a hitting set, and thus have eliminated at least one candidate that S beats)
- each remaining $e \in U$ has score $\alpha - 2$
- dummy candidates have non-positive score

Hence, at this point, candidate d has the highest Borda score and can be eliminated.

Conversely, suppose there is a ranking selected by Sequential-Borda-Winner where d is eliminated in round $k = t + 3$ or earlier. As observed above, in the first t rounds, only element candidates can be eliminated. Let $T = \{e_1, \dots, e_t\}$ be the set of elements whose candidates are eliminated in these rounds. In rounds $t + 1$ and $t + 2$, all candidates in B have score at least $\alpha + 2$ (it cannot go lower because the candidates that B beats are all dummy candidates which cannot be eliminated). Because d has score only $\alpha - 2$, the two candidates in B must be eliminated before d . This brings us to round $t + 3$ where by assumption d is eliminated. Hence at this point,

⁷Note that we are not claiming that this weighted majority graph can be realized by constantly many voters. The reason that a construction like in [Theorem D.8](#) does not translate is that we do not have a constant upper bound on the number of occurrences of each element. Such bounds typically cannot be imposed while retaining W[1]-hardness.

no candidate has score higher than $\alpha - 2$. In particular, for every set $S \in \mathcal{S}$, its score is less than α . This can only have happened if at least one of the elements of S has been eliminated and is thus part of T . It follows that T is a hitting set. \square