

# Dynamic Proportional Rankings

Jonas Israel and Markus Brill

Research Group Efficient Algorithms  
Technische Universität Berlin  
10587 Berlin, Germany  
{j.israel, brill}@tu-berlin.de

## Abstract

Proportional ranking rules aggregate approval-style preferences of agents into a collective ranking such that groups of agents with similar preferences are adequately represented. Motivated by the application of live Q&A platforms, where submitted questions need to be ranked based on the interests of the audience, we study a dynamic extension of the proportional rankings setting. In our setting, the goal is to maintain the proportionality of a ranking when alternatives (i.e., questions)—not necessarily from the top of the ranking—get selected sequentially. We propose generalizations of well-known aggregation rules to this setting and study their monotonicity and proportionality properties. We also evaluate the performance of these rules experimentally, using realistic probabilistic assumptions on the selection procedure.

## 1 Introduction

From “ask-me-anything” sessions to panel discussions and town hall meetings, an increasing number of both virtual and in-person discussion formats are enhanced by digital tools that aim to make the event more interactive and responsive to the audience. Using live Q&A platforms such as *slido* ([www.sli.do](http://www.sli.do)), *Mentimeter* ([www.mentimeter.com](http://www.mentimeter.com)) or *Pigeonhole Live* ([www.pigeonholelive.com](http://www.pigeonholelive.com)), participants in the audience can submit questions and upvote questions submitted by others; a moderator then selects the most popular questions for the discussion. By reducing barriers to participation (e.g., by allowing anonymous submissions), these tools aim to better represent the diversity in the audience.

The moderator of the discussion is presented with an aggregated list, in which audience questions are ranked by popularity (i.e., number of upvotes). Based on this ranking, the moderator then picks the next question. When selecting a question, it is usually not required to follow the ranking strictly; rather, the choice is at the moderator’s discretion, allowing him or her to take into account other factors such as discussion flow, etc. That being said, it is generally expected that questions at the top of the ranking are more likely to be selected than questions further down in the list. After a question has been selected, it is removed from the ranking.

Ranking questions solely by popularity, though intuitively appealing, has a major downside: minority opinions might go completely unrepresented, even when the minority makes up a substantial proportion of the audience. To illustrate this phenomenon, which is often referred to as “tyranny of the majority,” consider a situation in which the audience is composed of two groups. One group makes up 60% of the entire audience and is only interested in questions related to topic *A*; the remaining 40% of participants are only interested in questions on a different topic *B*. Now, assuming that sufficiently many questions on topic *A* have been submitted, and that participants only upvote questions related to their own interest, questions on topic *B* are unlikely to appear anywhere near the top of the ranking, which is populated exclusively by questions on topic *A*. As a consequence, questions on topic *B* are very unlikely to be selected, despite the fact that these questions are supported by 40% of the audience.

In this paper, we propose an approach to avoid the problem of underrepresenting minority opinions. Specifically, we model the scenario described above as a *proportional representation* problem and employ ranking algorithms based on (approval-based) proportional voting rules [Aziz *et al.*, 2017; Brill *et al.*, 2017]. The algorithms we consider aggregate the upvotes of the participants into a *proportional* ranking over questions, such that each minority (i.e., group of participants with similar preferences) is represented in the ranking to an extent that is proportional to the group’s size. Whenever a question is selected by the moderator, our methods dynamically recompute the ranking, pushing questions supported by underrepresented groups closer to the top.

At a technical level, our point of departure is the theory of proportional rankings [Skowron *et al.*, 2017], which studies how a collective ranking over a set of alternatives can be constructed in such a way that majority and minority opinions are represented adequately. The question we are interested in is how proportional ranking algorithms can be adapted to the dynamic setting. More specifically, we ask:

*How can the proportional representativeness of a collective ranking be maintained in a dynamic setting, where alternatives get selected sequentially?*

To answer this question, we consider two well-known aggregation rules dating back to the late 19th century: sequential Phragmén [Phragmén, 1894] and sequential PAV [Thiele, 1895]. These two rules, together with a few variants of the

latter, performed best in the analysis conducted by Skowron *et al.* [2017]. For both rules, we propose two distinct generalizations to the sequential selection setting: a *dynamic* variant and a *myopic* variant (see Section 3 for details). As a benchmark, we also consider the rule that simply orders questions by the number of received upvotes.

**Our contribution.** In this paper, (i) we formalize the setting of dynamic ranking rules and generalize the rules of Phragmén and Thiele to this setting (Section 3); (ii) we define a notion of satisfaction monotonicity and analyze to what extent the considered rules satisfy it (Section 4); (iii) we provide theoretical bounds regarding two different proportionality notions (Section 5); and (iv) we experimentally evaluate our dynamic ranking rules (Section 6). Omitted proofs and further details can be found in the full version of this paper [Israel and Brill, 2021].

**Related work.** Proportional representation is a fundamental desideratum in multiwinner elections [Monroe, 1995; Faliszewski *et al.*, 2017; Lackner and Skowron, 2020]. For approval preferences in particular, a wide variety of proportionality axioms have been studied [Aziz *et al.*, 2017; Sánchez-Fernández *et al.*, 2017; Janson, 2018; Peters and Skowron, 2020]. Proportionality in the context of rankings has been considered in the aforementioned paper by Skowron *et al.* [2017] and (for linear preferences) by Schulze [2011]. Notions of fairness over multiple elections among a fixed set of voters have received considerable attention in previous years. This line of work includes, e.g., the study of long-term fairness over different decisions [Freeman *et al.*, 2017; Lackner, 2020], single decisions under changing preferences [Tennenholtz, 2004; Boutilier and Procaccia, 2012; Parkes and Procaccia, 2013; Oren and Lucier, 2014; Hemaspaandra *et al.*, 2017], and storable votes [Casella, 2005; Casella, 2012]. In a practical attempt to avoid the underrepresentation of minorities, the live Q&A app *SpeakUp* (www.speakup.digital) allows audience members to add *attributes* (relating to, e.g., gender or education) to submitted questions. The moderator can then manually filter questions with attributes that have been underrepresented in the discussion. Requiring organizers to identify relevant attributes poses the risk of overlooking important subgroups or introducing unwanted biases; it also presumes the willingness of participants to reveal potentially sensitive information. In contrast, the ranking algorithms considered in this paper do not require attributes in order to ensure the representation of minority opinions.

## 2 Preliminaries

We briefly introduce some basic concepts from the theory of approval-based preference aggregation; for details, see the survey by Lackner and Skowron [2020]. Let  $C$  be a finite set of candidates and  $N = \{1, \dots, n\}$  a finite set of voters. An (approval) profile  $A = (A_1, \dots, A_n)$  is a list that contains, for each  $i \in N$ , the approval set  $A_i \subseteq C$  of voter  $i$ . Given an approval profile  $A$  and a candidate  $c \in C$ , we let  $N_c = \{i \in N : c \in A_i\}$  denote the *supporters* of  $c$ . The *approval score* of  $c$  is given by  $|N_c|$ . In the motivating application,  $C$  consists of all submitted questions and  $A_i$  contains

the questions that have been upvoted by participant  $i$ .

For a finite set  $S$ , we let  $\mathcal{L}(S)$  denote the set of all linear orders, or *rankings*, over  $S$ . We often write a ranking  $r \in \mathcal{L}(S)$  as a sequence  $r = (r_1, r_2, \dots, r_{|S|})$ , and for  $j \leq |S|$ , we let  $r_{\leq j}$  denote the set  $\{r_1, r_2, \dots, r_j\}$  of the first  $j$  elements in  $r$ .

An approval-based *ranking rule* maps an approval profile  $A$  to a ranking  $r \in \mathcal{L}(C)$  of all candidates. We will make use of the following three (non-dynamic) ranking rules.<sup>1</sup>

*Approval Voting (AV)* ranks the candidates according to their approval score. This rule is not proportional and we use it mainly as a benchmark.

*Sequential PAV (seqPAV)* ranks candidates iteratively, in each iteration choosing an unranked candidate maximizing the marginal contribution in terms of weighted voter satisfaction. Formally, for a subset  $S \subseteq C$  of candidates, define  $\text{sc}(S) = \sum_{i \in N} \sum_{j=1}^{|A_i \cap S|} \frac{1}{j}$ . If  $k$  candidates have already been ranked, the marginal contribution of an unranked candidate  $c$  is given by  $\text{mc}(c) = \text{sc}(r_{\leq k} \cup \{c\}) - \text{sc}(r_{\leq k})$ .

*Sequential Phragmén* can be described in terms of voters buying candidates.<sup>2</sup> Every candidate costs 1 credit. All voters start without any credits but earn them continuously over time (at a constant and identical rate). As soon as a group of voters who all approve the same candidate  $c$  together own 1 credit, they immediately buy that candidate; at this point, their balance is reset to 0 and candidate  $c$  is added in the next position of the ranking. This is done until all candidates are ranked.

## 3 Dynamic Ranking Rules

In this section, we formally introduce the setting of dynamic ranking rules and we adapt existing (non-dynamic) ranking rules to this setting. The input of a dynamic ranking rule consists of two parts: an approval profile and a (potentially empty) sequence of candidates that have already been “implemented” or “executed”; the output is a ranking of all not-yet-implemented candidates. To formalize this notion, we let  $X = (x_1, x_2, \dots, x_j)$  denote the sequence of implemented candidates (where  $j \in \{0, \dots, |C| - 1\}$ ); whenever the order of elements in  $X$  does not matter, we slightly abuse notation and treat  $X$  as the set  $X = \{x_1, x_2, \dots, x_j\}$ .

**Definition 1.** An (approval-based) dynamic ranking rule  $\mathcal{R}$  maps a profile  $A$  and a sequence  $X = (x_1, x_2, \dots, x_j)$  of candidates to a ranking  $\mathcal{R}(A, X) \in \mathcal{L}(C \setminus X)$ .

Applying a dynamic ranking rule to a sequential selection process (as outlined in the introduction) is now straightforward: At the beginning, when no candidate has yet been implemented,  $X = ()$  and the ranking  $\mathcal{R}(A, ())$  ranks all candidates in  $C$ . Given this ranking, a *decision maker (DM)* selects an alternative  $x_1 \in C$  to be implemented. The updated ranking of the remaining candidates is then given by  $\mathcal{R}(A, (x_1))$ , and the process is repeated. At iteration  $t \in \mathbb{N}$ , when  $t - 1$  candidates have been implemented and

<sup>1</sup> All rules may encounter ties; we assume that a priority ordering over candidates is used as a tiebreaker. In the motivating example, the submission time of a question yields a natural priority ordering.

<sup>2</sup> Another (equivalent) formulation of this method is in terms of a load balancing procedure [Janson, 2016; Brill *et al.*, 2017].

thus  $X = (x_1, x_2, \dots, x_{t-1})$ , we let  $r^t$  denote the ranking  $\mathcal{R}(A, X) \in \mathcal{L}(C \setminus X)$  from which the DM can make a choice.

We will sometimes make the assumption that the DM only ever implements alternatives that appear near the top of the ranking. In this *depth-restricted setting*, we are given a natural number  $h$  and we assume that  $x_t \in r_{\leq h}^t$  for all time steps  $t$ . This setting models situations in which the DM does not have the resources (or the ability) to consider the whole ranking.

The straightforward ranking rule AV trivially translates to the dynamic setting: When a candidate is implemented, it is simply removed from the ranking; the order between the remaining candidates does not change. AV is used in all of the live Q&A platforms mentioned in Section 1. In the following, we propose dynamic variants of other ranking rules.

**Dynamic seqPAV.** For this straightforward dynamization of seqPAV, we modify the notion of marginal contribution to also take into account the satisfaction derived from previously implemented candidates:

$$\text{mc}_{\text{dyn}}(c) = \text{sc}(X \cup r_{\leq k} \cup \{c\}) - \text{sc}(X \cup r_{\leq k}).$$

Dynamic seqPAV ranks candidates iteratively, adding in each round a candidate  $c$  maximizing  $\text{mc}_{\text{dyn}}(c)$ .

**Dynamic Phragmén.** Our first dynamization of sequential Phragmén works in two phases. As before, voters buy candidates and every candidate has a cost of 1 credit. Voters do not start with 0 credits, however; they may have an initial *debt* due to previously implemented candidates they approve. The debts of voters are determined in the *first phase*, which iterates through the sequence  $X$  (starting with  $x_1$ ) and, for each implemented candidate  $x_j \in X$ , divides the cost of 1 among the voters in  $N_{x_j}$ . More precisely, this assignment of debts is done in such a way that, in each iteration  $j$ , the maximum total debt across all voters in  $N_{x_j}$  is as small as possible. (The assignment of debts, therefore, mimics the assignment of loads in the load-balancing formulation of sequential Phragmén.) We let  $d_i \geq 0$  denote the total debt of voter  $i \in N$  resulting from this first phase. In the *second phase*, we run sequential Phragmén to obtain the desired ranking of candidates in  $C \setminus X$ . At the beginning of this phase, each voter  $i$  has a credit balance of  $-d_i \leq 0$ . As in sequential Phragmén, voters continuously earn credits, and voters starting with debts can only participate in the purchase of a candidate once they have a positive balance. Voters are not allowed to go into debt for buying candidates.

These dynamic rules rank candidates in the same fashion as their non-dynamic counterparts, while taking the sequence  $X$  of previously implemented candidates into account. (Note that the implementation *order* matters for dynamic Phragmén, but not for dynamic seqPAV.) In particular, both dynamic rules coincide with their non-dynamic counterpart when  $X = ()$ . Moreover, the ranking among the remaining candidates does not change whenever the top-ranked candidate is implemented: if  $r^t = (r_1, r_2, r_3, \dots)$  and  $x_t = r_1$ , then  $r^{t+1} = (r_2, r_3, \dots)$ .

We also consider two “myopic” dynamic ranking rules.

**Myopic seqPAV.** In this myopic dynamization of seqPAV, we compute the marginal contribution of each candidate  $c \in$

$C \setminus X$  only with respect to the set  $X$  of previously implemented candidates, i.e.,  $\text{mc}_{\text{myopic}}(c) = \text{sc}(X \cup \{c\}) - \text{sc}(X)$ . Then, we simply rank those candidates according to decreasing  $\text{mc}_{\text{myopic}}(c)$ -value.

**Myopic Phragmén.** In this myopic dynamization of sequential Phragmén, we first run the first phase of dynamic Phragmén in order to determine the debts  $\{d_i\}_{i \in N}$  of voters. Then, for each candidate  $c \in C \setminus X$ , we compute the voter debts that would result from adding candidate  $c$  to  $X$  (and running the first phase for one more iteration). Let the debts induced by candidate  $c$  be  $\{d_i^c\}_{i \in N}$ . Myopic Phragmén ranks the candidates in  $C \setminus X$  according to increasing  $\max_{i \in N_c} d_i^c$ , breaking ties according to the second highest debt, and so on.

Intuitively, myopic seqPAV and myopic Phragmén rank candidates according to their suitability of being the next implemented candidate. In contrast to dynamic seqPAV and dynamic Phragmén, this way of comparing candidates does not lead to rankings that are representative by themselves. In particular, both myopic rules coincide with AV when  $X = ()$ .

We illustrate these rules with a simple example.

**Example 2.** Let  $C = \{a, b, c, d, e\}$  and assume alphabetic tiebreaking. Consider a set of 9 voters with the following approval sets:

$$5 \times \{a, b\}, \quad 3 \times \{c, d\}, \quad 1 \times \{e\}.$$

Let  $V$  denote the group consisting of the 5  $\{a, b\}$ -voters and  $V'$  the group consisting of the 3  $\{c, d\}$ -voters. First, consider dynamic seqPAV and dynamic Phragmén. In the first iteration, both rules output  $r^1 = (a, c, b, d, e)$ , effectively alternating between candidates supported by voter groups  $V$  and  $V'$ . Let us assume that the DM implements candidate  $x_1 = b$ , i.e.,  $X = (b)$ . Then, the two rules output  $r^2 = (c, a, d, e)$ .

Next, consider myopic seqPAV and myopic Phragmén. In the first iteration, both rules (and AV) rank the candidates according to their approval scores:  $r^1 = (a, b, c, d, e)$ . After the implementation of  $b$ , both rules output  $r^2 = (c, d, a, e)$ , which differs from the AV ranking  $r^2 = (a, c, d, e)$ .

In this example, all of our ranking rules demote candidate  $a$  in  $r^2$  because voter group  $V$  is already (partially) satisfied with  $X = (b)$ . The myopic rules even rank *both*  $c$  and  $d$  higher than  $a$  in  $r^2$ , since implementing either  $c$  or  $d$  would yield a more proportional sequence  $X$  than choosing  $a$  would.

All presented ranking rules can be computed in polynomial time; see the full version of this paper for an asymptotic runtime analysis and pseudocode.

## 4 Monotonicity of Voter Satisfaction

We start our analysis of dynamic ranking rules by considering the satisfaction of voters during the sequential selection process. In doing so, we assume that voters derive satisfaction not only from implemented candidates they approve, but also—possibly to a lesser extent—from approved candidates appearing near the top of the ranking: high positions in the ranking come with increased attention (and, presumably, high selection probabilities in future iterations) for the respective

candidates. In particular, improved ranking positions of supported candidates can be viewed as a kind of compensation for (groups of) voters who are not (yet) well-represented by the implemented alternatives. To make this concrete, consider an iteration  $t$ , where the DM is confronted with ranking  $r^t$  and chooses to implement candidate  $x_t$ . Following the logic outlined above, it might be natural to expect that voters *not* approving  $x_t$  (or, more precisely, the candidates approved by these voters) should get a “boost” in the ranking. At the very least, it seems reasonable to expect that the satisfaction of such voters with the new ranking  $r^{t+1}$  is at least as high as with the old ranking  $r^t$ . AV trivially satisfies this property, which we informally refer to as *satisfaction monotonicity*. Somewhat surprisingly, however, the following simple example demonstrates that this intuitive monotonicity notion is not achievable for dynamic ranking rules that satisfy a minimal degree of representativeness.<sup>3</sup>

**Example 3.** Consider the following profile with 7 voters:

$$1 \times \{a\}, \quad 3 \times \{b\}, \quad 3 \times \{a, c\}.$$

All rules considered in this paper rank the approval winner  $a$  first in  $r^1$ . If the DM chooses to implement candidate  $x_1 = c$ , all of our rules—except AV—output  $r^2 = (b, a)$  in the second iteration. Intuitively, the rules give more voting power to the 3 supporters of  $b$  (all of which are unrepresented by  $c$ ) than to the 4 supporters of  $a$  (3 of which are already partially represented). Observe that the satisfaction of the voter approving  $a$  decreases when going from  $r^1$  to  $r^2$ , despite the fact that this voter does not approve the candidate being implemented.

The following definition is motivated by the question whether monotonicity failures can be prevented by moving to the depth-restricted setting and putting lower bounds on the size of voter groups for which monotonicity should hold. To measure satisfaction of a group  $V \subseteq N$  of voters with a set  $S \subseteq C$  of candidates, we use the *average satisfaction* of  $V$  with  $S$ , i.e.,  $\text{avg}_V(S) = \frac{1}{|V|} \cdot \sum_{i \in V} |A_i \cap S|$ .

**Definition 4.** For  $h \geq 1$  and  $\alpha \in (0, 1]$ , a dynamic ranking rule satisfies  $(h, \alpha)$ -monotonicity if, for all profiles and all groups of voters  $V \subseteq N$  of size  $|V| \geq \alpha \cdot |N|$ , the following holds for every iteration  $t$ :

$$\text{If } x_t \notin \bigcup_{i \in V} A_i, \text{ then } \text{avg}_V(r_{\leq h}^{t+1}) \geq \text{avg}_V(r_{\leq h}^t).$$

That is,  $(h, \alpha)$ -monotonicity requires that satisfaction monotonicity holds for groups that make up at least an  $\alpha$ -fraction of the electorate, and when measuring average satisfaction with respect to the first  $h$  positions in a ranking.

AV trivially satisfies  $(h, \alpha)$ -monotonicity for all  $h$  and all  $\alpha$ . On the other hand, all other considered rules violate this notion unless we consider rather large groups of voters.

**Proposition 5.** Consider the depth-restricted setting for some  $h \geq 3$ . Then, dynamic seqPAV and dynamic Phragmén fail to satisfy  $(h, \alpha)$ -monotonicity for all  $\alpha < \frac{6}{2h+5}$ .

<sup>3</sup>Example 3 can be turned into an impossibility result: Every dynamic ranking rule that (i) ranks the approval winner at the top in the first iteration and (ii) gives priority to less satisfied voter groups fails satisfaction monotonicity.

Furthermore, myopic seqPAV and myopic Phragmén fail to satisfy  $(h, \alpha)$ -monotonicity for all  $\alpha < \frac{1}{h}$ .

The monotonicity requirement can be weakened further by only requiring satisfaction monotonicity in cases in which the implemented candidate is never co-approved with any candidate that is approved by a member of the group under consideration (i.e., there is no  $c \in \bigcup_{k \in V} A_k$  with  $\{c, x_t\} \subseteq A_i$  for some  $i \in N$ ). Both myopic rules satisfy this weak implementation monotonicity, whereas the two dynamic versions fail it. For a thorough discussion of this weaker version of monotonicity, we refer to the full version of this paper.

Despite the negative results in this section, we rarely found monotonicity violations of any kind in our experiments (see Section 6).

## 5 Proportional Representation

We now turn to analyzing the proportional representativeness that is provided by our dynamic ranking rules. The following two sections capture different perspectives on representation, focusing on the representativeness of the ranking  $r^t$  at any given iteration  $t$  (Section 5.1) and on the representativeness of the set  $X$  of implemented candidates (Section 5.2).

### 5.1 Proportionality of Rankings

In certain applications of dynamic ranking rules, such as the live Q&A platforms mentioned in the introduction, it is desirable for the ranking  $r^t$  to provide a representative overview of the opinions of the voters at any given iteration  $t$ . In this section, we prove proportionality guarantees that are satisfied by ranking  $r^t$  for any fixed iteration  $t$ .

Measures for the proportionality of a ranking have been proposed by Skowron *et al.* [2017]. In particular,  $\kappa$ -group representation measures, informally speaking, how far down in the ranking a group of voters needs to look in order to obtain a given amount of satisfaction. In order to adapt the notion of  $\kappa$ -group representation to the dynamic ranking setting, we need the following notation. For iteration  $t$ , let  $X^t = \{x_1, \dots, x_{t-1}\}$  denote the set of candidates implemented in the first  $t-1$  rounds and, for a group  $V \subseteq N$  of voters, let  $\lambda^t(V) = |\bigcap_{i \in V} A_i \setminus X^t|$  denote the cohesiveness of  $V$  with respect to the remaining candidates  $C \setminus X^t$ .

**Definition 6** (Group representation). Let  $\kappa(\alpha, \lambda)$  be a function from  $((0, 1] \cap \mathbb{Q}) \times \mathbb{N}$  to  $\mathbb{N}$ . A dynamic ranking rule satisfies  $\kappa$ -group representation if the following holds for all profiles  $A$ , groups of voters  $V \subseteq N$ , rational numbers  $\alpha \in (0, 1]$ , and integers  $\lambda, t \leq |C|$ :

$$\text{If } |V| \geq \alpha \cdot n \text{ and } \lambda^t(V) \geq \lambda, \text{ then } \text{avg}_V(r_{\leq \kappa(\alpha, \lambda)}^t) \geq \lambda.$$

In words: If a group  $V$  of voters makes up an  $\alpha$ -fraction of the electorate and has at least  $\lambda$  commonly approved candidates remaining at iteration  $t$ , then this group derives an average satisfaction of at least  $\lambda$  from the candidates ranked in the top  $\kappa(\alpha, \lambda)$  positions of ranking  $r^t$ .<sup>4</sup>

<sup>4</sup>A natural lower bound for  $\kappa(\alpha, \lambda)$  is given by  $\lceil \lambda/\alpha \rceil$ . Note that the  $\kappa$  functions used in this section not only depend on  $\alpha$  and  $\lambda$ , but also on the set  $V$  and on the sequence  $X$  of previously implemented candidates. In an attempt to simplify notation, we decided to not make this dependencies explicit in Definition 6.

Our first result in this section is for dynamic Phragmén. We recall that  $d_i$  denotes the initial debt of voter  $i$  at the end of the first phase of the method, and let  $d_{\text{avg}}^V = \frac{1}{|V|} \sum_{i \in V} d_i$  denote the average debt of voters in  $V$ .

**Theorem 7.** *Dynamic Phragmén satisfies  $\kappa$ -group representation for*

$$\kappa(\alpha, \lambda) = \left\lceil \frac{2(\lambda + m + 1) + s \cdot |V|}{\alpha} \right\rceil,$$

where  $m = |\bigcup_{i \in V} A_i \cap X|$  and  $s = \sum_{i \in V} (d_i - d_{\text{avg}}^V)^2$ .

Observe that this function is increasing both in the number  $m$  of already implemented candidates that are approved by some voter in  $V$  and in the variance  $s$  of debts of voters in  $V$ . For the special case  $X = ()$ , Theorem 7 implies a group representation of  $\lceil \frac{2\lambda+2}{\alpha} \rceil$  for (non-dynamic) sequential Phragmén. For  $\lambda \geq 2$ , this is an improvement over the  $\kappa$ -group representation bound of  $\lceil \frac{5\lambda}{\alpha^2} + \frac{1}{\alpha} \rceil$  proved by Skowron *et al.* [2017].

The proof of Theorem 7 employs the notion of *proportionality degree* [Skowron, 2018]. In particular, we first prove a bound on the proportionality degree of dynamic Phragmén, using a potential function approach that is similar to the one used by Skowron [2018] for the non-dynamic setting. Then, we establish a relationship between the proportionality degree and group representation, and use it to translate the bound on the former into a bound on the latter.

For dynamic seqPAV we prove the following generalisation of Theorem 3 by Skowron *et al.* [2017]; the latter theorem corresponds to the special case  $X = ()$ .

**Theorem 8.** *Dynamic seqPAV satisfies  $\kappa$ -group representation for*

$$\kappa(\alpha, \lambda) = \left\lceil \frac{2(\lambda + 1 + \text{avg}_V(X))^2}{\alpha^2} \right\rceil.$$

AV does not perform any different in the dynamic ranking setting compared to the non-dynamic one. Thus, it satisfies the same bounds on group representation as those stated in Theorem 2 by Skowron *et al.* [2017]. Since myopic seqPAV and myopic Phragmén both agree with AV in the case  $X = ()$ , the same bounds hold for these two rules.

**Proposition 9.** *Myopic seqPAV and myopic Phragmén fail  $\kappa$ -group representation for  $\kappa(\alpha, \lambda) \leq \left\lceil \frac{\lambda \cdot \alpha}{2\alpha - 1} \right\rceil - 1$  and for all functions  $\kappa(\alpha, \lambda)$  if  $\alpha \leq \frac{m+1}{m+2}$ , where  $m = |\bigcup_{i \in V} A_i \cap X|$ .*

## 5.2 Proportionality of Implemented Candidates

In this section, we study worst-case bounds on the representativeness of the set  $X$  of implemented candidates. Clearly, no non-trivial bounds are obtainable without restricting the selection behavior of an adversarial DM. Therefore, we will make the following two assumptions throughout this section:

- (A1) The DM is depth-restricted and always implements a candidate from the top  $h$  positions of the ranking.
- (A2) Every candidate  $c \in C$  has sufficiently<sup>5</sup> many “clones,” i.e., candidates  $c'$  with identical supporter set  $N_{c'} = N_c$ .

<sup>5</sup>Given an upper bound  $T$  on the number of iterations,  $h + T - 1$  clones suffice (as at least  $h$  clones will always remain in the ranking).

Assumptions (A1) and (A2) together ensure that the DM can be forced to implement a candidate approved by a voter, by populating the top  $h$  positions exclusively with such candidates. Arguably the most natural way to ensure (A2) is to assume that we are in the *party-approval* setting [Brill *et al.*, 2020], where candidates are interpreted as parties and can be selected arbitrarily often. In the motivating example of live Q&A platforms, party-approval preferences could result from assigning attributes to questions and eliciting participants’ approval preferences over attributes.

Recall that  $X^{t+1}$  denotes the set containing the implemented candidates from the first  $t$  rounds. The following property is a natural adaption of the well-studied proportionality axiom *proportional justified representation* (PJR) [Sánchez-Fernández *et al.*, 2017].

**Definition 10.** *A dynamic ranking rule satisfies proportional justified selection (PJS) if the following holds for all  $t, \ell \in \mathbb{N}$  and for all groups  $V \subseteq N$  of voters: If  $|V| \geq \frac{\ell}{t} \cdot |N|$  and  $|\bigcap_{i \in V} A_i| \geq \ell$ , then  $|X^{t+1} \cap \bigcup_{i \in V} A_i| \geq \ell$ .*

A weaker version of this axiom is obtained by fixing  $\ell = 1$ ; in analogy to a well-known notion due to Aziz *et al.* [2017], we refer to the resulting property as *justified selection* (JS).

We prove the following theorem by treating the set  $X^{t+1}$  of implemented alternatives as a committee and using the fact that sequential Phragmén satisfies PJR [Brill *et al.*, 2017].

**Theorem 11.** *Under assumptions (A1) and (A2), myopic Phragmén satisfies PJS.*

Analogously, we can translate a representation guarantee for sequential PAV [Sánchez-Fernández *et al.*, 2017] into a guarantee for myopic seqPAV.

**Proposition 12.** *Under assumptions (A1) and (A2), myopic seqPAV satisfies JS for  $t \leq 5$ .*

Similar positive results are not possible for the other rules we consider.

**Proposition 13.** *Dynamic seqPAV and dynamic Phragmén fail to satisfy JS, even under assumptions (A1) and (A2) and for  $t = 2$ .*

## 6 Experimental Evaluation

In order to better understand the behavior of the dynamic ranking rules considered in this paper, we conducted computational experiments using randomly generated approval profiles. Since we were mainly interested in the proportional representation of groups of voters with similar preferences, we generated profiles according to two probabilistic models that lead to polarized electorates with easily identifiable groups. We measured (1) how the satisfaction of a voter group with the set of implemented candidates varies with the size of the group, and (2) how the satisfaction of a voter group with the current ranking varies over time.

**Setup.** All of our profiles consist of 60 voters and 20 candidates, and the approval sets are generated according to two different models. In the *blurred parties model*, we assign each voter and each candidate to one of *two* parties. The size of the voter group  $V$  associated with the first party varies over the experiments; the candidates are always divided equally. Each

voter approves a candidate from their own party with 95% probability and a candidate of the other party with 5% probability. The *spatial model* is an adaption of the 4-Gaussian model used by Elkind *et al.* [2017] for linear preferences. Voters and candidates correspond to points in the Euclidean plane and voters approve nearby candidates. There are *three* parties with equidistant center locations, and candidates as well as voters get sampled as points around the party centers according to a normal distribution. We let the size of the voter group  $V$  associated with the first party grow, and divide the remaining voters equally among the two remaining parties. There are 7 candidates associated with the first party. The selection behavior of the DM is modeled via Google click-through rates. In particular, the probability of selection decreases when going down the ranking. In Figure 1 we plot the averages of 100 generated elections. More details about the setup can be found in the full version of this paper.

**Satisfaction with implemented candidates.** We measure the average satisfaction of voter group  $V \subseteq N$  with  $X^{k+1}$ , where  $k$  is the number of candidates associated with that group (i.e.,  $k = 10$  for the blurred parties model and  $k = 7$  for the spatial model). We plot this value against the relative size  $\alpha = |V|/|N|$  of the group  $V$ . The graphs in the first row of Figure 1 show that for both models, AV is not proportional:  $\text{avg}_V(X^{k+1})$  starts out very low and only jumps up as soon as  $V$  becomes the biggest group (which happens at  $\alpha = 1/2$  and  $\alpha = 1/3$ , respectively). In other words, AV underrepresents minorities and overrepresents majorities. The performance of the other four rules are indistinguishable, as all yield proportionally increasing satisfaction values.

**Satisfaction with rankings.** The graphs in the second row of Figure 1 depict the average satisfaction of a group  $V$  of size  $\alpha = 1/4$  with the first 5 candidates of the ranking over the first 11 iterations. Again, AV behaves poorly, as it gives satisfaction to  $V$  only once the larger groups have been satisfied. The satisfaction values under the two myopic rules jump heavily from one iteration to the next, as these rules tend to mainly represent one group of voters per iteration. On the other hand, the two dynamic rules keep the satisfaction of  $V$  relatively constant at around one fourth of the maximum possible satisfaction. These rules provide proportional representation in each single iteration, which is in line with the theoretical results in Section 5.1.

## 7 Conclusion

Motivated by the problem of how submitted questions in a live Q&A session can be ranked in a more representative way, we have introduced dynamic ranking rules. We proposed two paradigms of dynamizing existing ranking rules: under the *dynamic paradigm*, we target proportional representation of voter interests at each individual time step; under the *myopic paradigm*, we try to make the set of implemented candidates as representative as possible. While the former approach lends more flexibility for the decision maker and guarantees a proportional exposure of candidates in each ranking, the latter approach is computationally more efficient and yields stronger selection guarantees. Our experimental




Figure 1: Experimental results for the blurred parties model (left) and the spatial model (right). The graphs in the first row show the average satisfaction of  $V$  with the first  $k$  implemented candidates, for relative group size  $\alpha \in [0, 1]$ . The graphs in the second row show the average satisfaction of  $V$  with  $r_{\leq 5}^t$ , for  $1 \leq t \leq 11$ .

results illustrate the difference between the two approaches, and verify that both approaches lead to proportional results.

The application of live Q&A platforms gives rise to some interesting extensions of our model. In realistic scenarios, neither the electorate nor the set of candidates is static, as people enter or leave the audience and new questions come up continuously. Moreover, participants can change their approval preferences throughout the event. Our approach can take these dynamic aspects into account in a straightforward manner: After each implementation, we can apply our ranking rules to the current set of not-yet-implemented candidates and to the current approval preferences—the only necessary information from previous iterations is the sequence of implemented candidates.

The dynamic ranking rules proposed in this paper are applicable to a wide variety of sequential selection procedures in which proportional representation is desired and, at the same time, some flexibility on the part of the decision maker is necessary (e.g., think of human-in-the-loop decision support systems for hiring or budgeting decisions). Other applications of dynamic ranking rules include committee election scenarios in which some part of the committee is fixed (e.g., due to external constraints) and the remaining seats need to be filled in such a way that the committee as a whole is representative.

## Acknowledgments

We would like to thank the anonymous reviewers for their insightful comments. This work was partially supported by the Deutsche Forschungsgemeinschaft under grant BR 4744/2-1. We thank Paula Blechschmidt, Cristina Cornelio, Benny Kimelfeld, Phokion Kolaitis, Kunal Relia, and Julia Stoyanovich for helpful comments and discussions.

## References

- [Aziz *et al.*, 2017] Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 48(2):461–485, 2017.
- [Boutilier and Procaccia, 2012] Craig Boutilier and Ariel Procaccia. A dynamic rationalization of distance rationalizability. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1278–1284. AAAI Press, 2012.
- [Brill *et al.*, 2017] Markus Brill, Rupert Freeman, Svante Janson, and Martin Lackner. Phragmén’s voting methods and justified representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 406–413. AAAI Press, 2017.
- [Brill *et al.*, 2020] Markus Brill, Paul Gözl, Dominik Peters, Ulrike Schmidt-Kraepelin, and Kai Wilker. Approval-based apportionment. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1854–1861. AAAI Press, 2020.
- [Casella, 2005] Alessandra Casella. Storable votes. *Games and Economic Behavior*, 51(2):391–419, 2005.
- [Casella, 2012] Alessandra Casella. *Storable Votes: Protecting the Minority Voice*. Oxford University Press, 2012.
- [Elkind *et al.*, 2017] Edith Elkind, Piotr Faliszewski, Jean-François Laslier, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. What do multiwinner voting rules do? An experiment over the two-dimensional euclidean domain. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 494–501. AAAI Press, 2017.
- [Faliszewski *et al.*, 2017] Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. Multiwinner voting: A new challenge for social choice theory. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 2. 2017.
- [Freeman *et al.*, 2017] Rupert Freeman, Seyed Majid Zahedi, and Vincent Conitzer. Fair social choice in dynamic settings. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 4580–4587. IJCAI, 2017.
- [Hemaspaandra *et al.*, 2017] Edith Hemaspaandra, Lane A. Hemaspaandra, and Jörg Rothe. The complexity of controlling candidate-sequential elections. *Theoretical Computer Science*, 678:14–21, 2017.
- [Israel and Brill, 2021] Jonas Israel and Markus Brill. Dynamic proportional rankings. Technical report, arXiv:2105.08043 [cs.GT], 2021.
- [Janson, 2016] Svante Janson. Phragmén’s and Thiele’s election methods. Technical report, arXiv:1611.08826 [math.HO], 2016.
- [Janson, 2018] Svante Janson. Thresholds quantifying proportionality criteria for election methods. Technical report, arXiv:1810.06377 [cs.GT], 2018.
- [Lackner and Skowron, 2020] Martin Lackner and Piotr Skowron. Approval-based committee voting: Axioms, algorithms, and applications. Technical report, arXiv:2007.01795 [cs.GT], 2020.
- [Lackner, 2020] Martin Lackner. Perpetual voting: Fairness in long-term decision making. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, pages 2103–2110. AAAI Press, 2020.
- [Monroe, 1995] Burt L. Monroe. Fully proportional representation. *The American Political Science Review*, 89(4):925–940, 1995.
- [Oren and Lucier, 2014] Joel Oren and Brendan Lucier. On-line (budgeted) social choice. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1456–1462. AAAI Press, 2014.
- [Parkes and Procaccia, 2013] David Parkes and Ariel Procaccia. Dynamic social choice with evolving preferences. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI)*, pages 767–773. AAAI Press, 2013.
- [Peters and Skowron, 2020] Dominik Peters and Piotr Skowron. Proportionality and the limits of welfarism. In *Proceedings of the 21st ACM Conference on Economics and Computation (ACM-EC)*, pages 793–794. ACM Press, 2020.
- [Phragmén, 1894] L. Edvard Phragmén. Sur une méthode nouvelle pour réaliser, dans les élections, la représentation proportionnelle des partis. *Öfversigt af Kongliga Vetenskaps-Akademiens Förhandlingar*, 51(3):133–137, 1894.
- [Sánchez-Fernández *et al.*, 2017] Luis Sánchez-Fernández, Edith Elkind, Martin Lackner, Norberto Fernández García, Jesús A. Fisteus, Pablo Basanta Val, and Piotr Skowron. Proportional justified representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 670–676. AAAI Press, 2017.
- [Schulze, 2011] Markus Schulze. Free riding and vote management under proportional representation by the single transferable vote. <http://m-schulze.9mail.de/schulze2.pdf>, 2011. Accessed: 2021-01-21.
- [Skowron *et al.*, 2017] Piotr Skowron, Martin Lackner, Markus Brill, Dominik Peters, and Edith Elkind. Proportional rankings. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 409–415. IJCAI, 2017.
- [Skowron, 2018] Piotr Skowron. Proportionality degree of multiwinner rules. Technical report, arXiv:1810.08799 [cs.GT], 2018.
- [Tennenholtz, 2004] Moshe Tennenholtz. Transitive voting. In *Proceedings of the 5th ACM Conference on Electronic Commerce (ACM-EC)*, pages 230–231. ACM Press, 2004.
- [Thiele, 1895] Thorvald N. Thiele. Om flerfoldevalg. *Översigt over det Kongelige Danske Videnskabernes Selskabs Forhandlingar*, pages 415–441, 1895.