

Online Approval Committee Elections

Virginie Do^{1,2}, Matthieu Hervouin², Jérôme Lang³ and Piotr Skowron⁴

¹Meta AI Research, Paris, France

²Université Paris Dauphine, PSL, CNRS, LAMSADE, France

³CNRS, Université Paris Dauphine, PSL, LAMSADE, France

⁴University of Warsaw, Poland

{virginie.do, matthieu.hervouin}@dauphine.eu, jerome.lang@lamsade.dauphine.fr,
p.skowron@mimuw.edu.pl

Abstract

Assume k candidates need to be selected. The candidates appear over time. Each time one appears, it must be immediately selected or rejected—a decision that is made by a group of individuals through voting. Assume the voters use approval ballots, i.e., for each candidate they only specify whether they consider it acceptable or not. This setting can be seen as a voting variant of choosing k secretaries. Our contribution is twofold. (1) We assess to what extent the committees that are computed online can proportionally represent the voters. (2) If a prior probability over candidate approvals is available, we show how to compute committees with maximal expected score.

1 Introduction

In the vast majority of elections, the set of candidates is known upfront. Yet, there are contexts where candidates appear over time. A paradigmatic example is hiring for a job: candidates come every day to pass an interview, they are evaluated by members of a jury, *and then it must be decided immediately whether to hire them or not*. When we must hire only one candidate and when the evaluation is performed by a single agent (human or algorithm), we get the classic *secretary problem*. The variant where several employees must be hired, is called the *multiple secretary problem*. These problems have numerous variants (for instance, depending on whether we want to optimize the rank or the value of the selected candidate, whether the distribution over candidates' values is known, etc.). However, when the candidates are evaluated by a set of voters, we obtain a *voting* version of the *secretary problem*, or equivalently, an *online* version of *multiwinner elections* (also called *committee elections*).

A typical example is the recruitment of a team of researchers for a research unit, from a pool of candidates who are being interviewed one at a time. In this example, the jury member may have different backgrounds and may have heterogeneous preferences over candidates. A second example is another hiring scenario where the voters are explicit criteria (e.g., skills or demographic attributes such as gender).

In generalized secretary problems, each candidate is usually evaluated by a single number [Babaioff *et al.*, 2008]. In voting,

numerical evaluations are not very convenient, and simpler types of ballots are typically used. In this paper we assume each voter submits an approval ballot. Deciding whether to approve or disapprove the current candidate is cognitively easy and the landscape of approval-based committee rules (ABC rules) is relatively well understood. In particular, ABC rules are well studied for *proportionality*, which guarantees that voters' approvals are fairly represented in the elected committee. In contexts where voters are evaluation criteria, it ensures that these are well covered by the set of hired candidates.

Formally, we have a set of n voters; at each time t , a new candidate c_t is observed; c_t is then approved or disapproved (possibly after being interviewed) by each voter; and we have to decide immediately whether to include c_t in the committee. We have to select a committee of size k , and exactly m candidates can be observed.

If we could wait until all candidates have been interviewed, then we would be in the classic setting of approval-based committee elections. In this setting there is a plethora of well-understood rules with different distinctive properties [Lackner and Skowron, 2020]; we would then pick one of these rules, say f , and compute its outcome. We cannot do this, because if candidate c_t appears at time t , then it must be decided at time t whether to hire it or not.¹ Still, the rule f can serve as a reference: the set of candidates that would have been computed by f if we had been able to wait can be considered the optimal set of winners, and can be used for measuring the quality of an online selection algorithm.

We consider two paradigms for evaluating rules.

First, we examine ABC rules f that aim at maximizing certain scoring functions f -sc. Given a fixed function f -sc, we evaluate online rules by comparing the scores of returned committees with the scores of the optimal ones. We explain how the existing mechanisms for the multiple secretary problem can be applied in order to obtain rules that perform well. Yet, our first contribution lies in designing optimal algorithms for the case, where a prior distribution over candidate approvals is available. The goal becomes to design algorithms that maxi-

¹We could consider intermediate contexts where we can wait some amount of time before deciding to hire a candidate or not, but in this first study we will simply assume that the decision to hire a candidate or not must be done immediately and irrevocably. This is often realistic: a good candidate has good chances to find another job if not hired immediately.

mize the expected score of the selected committees. We show that for selected scoring functions f -sc or under some assumptions, this can be done in polynomial time.

Second, we look at two *axioms of proportionality*. These properties require that each group of voters with cohesive preferences must be represented in the selected committee proportionally to its size. We give a polynomial-time computable online selection policy that satisfies the axiom of *proportional justified representation* [Sánchez-Fernández *et al.*, 2017]. We further show that the stronger axiom of *extended justified representation* is not satisfiable in the online setting, but it can be approximated. We give two such approximation algorithms, one which gives the best possible approximation but is computationally inefficient, and the other one, which can be computed in polynomial time.

After discussing related work (Section 2), we define the online committee elections problem (Section 3). Then, we discuss the construction of policies maximizing the expected score (Section 4), and the construction of policies providing some exact or approximate proportionality guarantees (Section 5). We conclude in Section 6.

2 Related Work

Online selection problems Our setting is close to generalized secretary problems [Babaioff *et al.*, 2008], where the goal is to hire the best possible subset of candidates among a finite set of candidates arriving one at a time. A candidate’s value is revealed upon arrival, and the hiring decision must be taken immediately and cannot be changed afterwards. The connection between our work and this class of problems is detailed in Section 4.1. [Yu *et al.*, 2019] consider a bi-criteria secretary problem with multiple choices, where each criterion can be seen as a voter, the subset of choices as the elected committee, and their objective function as the multiwinner approval voting objective. We consider instead a multiplicity of voters, and various objectives corresponding to general Thiele rules in multiwinner voting. In the single secretary problem variant of [Bearden *et al.*, 2005], there are multiple independent attributes, which can be seen as voters in a single-winner election. In the context of search engines, [Panigrahi *et al.*, 2012] aim to find a diverse set of items from an input stream, by maximizing a coverage function of multiple features. In our committee election framework, items can be seen as candidates and features as voters, yet their objective function is different from the committee scoring functions we consider, and proportional representation is not studied.

Social choice in online settings Our work is mainly related to the study of proportional representation in committee elections, and in particular approval-based committee elections, which are surveyed in [Lackner and Skowron, 2020]. In recent years, there has been increased interest in studying online versions of voting problems. Proportionality is studied in [Dey *et al.*, 2017] who formalize voting streams, a setting in which alternatives are fixed, but voters arrive in an online manner, which is the opposite to ours. In [Freeman *et al.*, 2017], the sets of voters and alternatives are fixed, but the valuations of each voter for an alternative varies over time. Utilities are defined at each timestep as the cumulative reward of each

agent given past decisions, and the goal is to maximize Nash social welfare. [Lackner, 2020] consider a similar setting with ordinal preferences instead of cardinal valuations, and study voting rules that weight agents according to their past satisfaction. [Hossain *et al.*, 2021] address the partial observability of voters’ preferences.

While these works study (repeated) single-winner elections, the only existing work on online multiwinner elections to our knowledge is [Oren and Lucier, 2014]. The difference with our setting is that they consider online random arrival of voters rather than candidates, and they do not study proportionality axioms. [Do *et al.*, 2021] also study a close online committee selection problem to ours, yet a major difference is the absence of voters in their case. Proportionality is then defined based on multiple demographic attributes and a distance to target proportions on these attributes. In independent work, [Banerjee *et al.*, 2022] study a similar setting of fair online allocation where each public good can be assimilated to an election candidate. While they focus on a quantitative notion of proportional fairness, we study welfare guarantees and qualitative proportionality axioms.

Recent studies address fairness in online versions of other public decision-making problems, such as dynamic proportional rankings [Israel and Brill, 2021] and participatory budgeting [Lackner *et al.*, 2021].

3 Preliminaries

For each $i \in \mathbb{N}$ we write $[i]$ to denote the set $\{1, \dots, i\}$. By $H(i)$ we denote the i -th harmonic number, i.e., $H(i) = \sum_{j=1}^i 1/j$. By $w(\cdot)$ we denote the inverse function² of $x \mapsto x^x$, i.e., $w(i) = x$ if $x^x = i$; clearly $w(i) = O(\log(i))$. Further, it holds that $\log(i) = O(w(i)^2)$.

An *approval-based election* (in short, an election) is a triple $E = (C, N, k, (A_i)_{i \in N})$, where $C = \{1, \dots, m\}$ is the set of *candidates*, $N = \{1, 2, \dots, n\}$ is the set of *voters*, k is the desired size of the committee, and for each $i \in N$, $A(i) \subseteq C$ is the approval ballot associated with i , that is, the set of candidates that i finds acceptable. Conversely, we let $N(c) = \{i \in N : c \in A(i)\}$ denote the set of voters who approve candidate c .

We refer to k -elements subsets of C as to size- k *committees*. An *approval-based committee election rule* (in short, an ABC rule) is a function \mathcal{R} that takes as input an election $E = (C, N, k)$ and returns a nonempty set of committees; we call the elements of $\mathcal{R}(E)$ winning committees. Typically we are interested in selecting a single winning committee, but we allow for ties.

An *online ABC rule* is an algorithm that iterates over the candidates according to the sequence c_1, c_2, \dots, c_m , and in each step makes the decision whether to include a candidate at hand, c_t , in the winning committee, or not. When making such a decision we assume that the algorithm does not know the preferences of the voters over the candidates $c_{t'}$ with $t' > t$. In other words, we assume the candidates appear one after another over time. When a candidate c_t appears, the voters’

²We have $w(i) = \exp(W(\ln(i)))$ where W is the Lambert function, i.e. the inverse multivalued function of $x \mapsto xe^x$.

preferences regarding c_t are revealed, and the algorithm needs to make an irrevocable decision of whether c_t is selected or not.

In the following sections, we evaluate winning committees $W \in \mathcal{R}(E)$ based on the approval ballots $A(i)$ of the voters, either assigning a value to W to measure aggregated satisfaction (Section 4), or analyzing the proportionality axioms satisfied by W (Section 5). Importantly, in the online setting we consider, the approval ballots are not available beforehand, but only once all candidates have been seen and approvals revealed. It is possible to analyze an ABC rule *ex-ante*, by measuring the quality of W in terms of $\mathbb{E}[f(|A(i) \cap W|)]$ for some f , as we do in Section 4. It is also possible to evaluate it *ex-post*, once all approval ballots are available and a committee has been elected, as we do in Section 5.

4 Maximizing Aggregated Satisfaction

In this section we look at the problem of maximizing the aggregated voters' satisfaction. Consider a voter i , who approves r members of the elected committee W , i.e., $r = |A(i) \cap W|$. Given a function $f: \mathbb{N} \rightarrow \mathbb{R}$, we define the f -utility of i from W as $f(r)$. An f -Thiele method [Thiele, 1895; Lackner and Skowron, 2020] is an ABC rule that maximizes the total f -utility of the voters: given an election $E = (C, N, k)$ it elects committees W that maximize the following score:

$$f\text{-sc}(W) = \sum_{i \in N} f(|A(i) \cap W|).$$

Examples of Thiele methods commonly studied in the literature include (1) Multiwinner Approval Voting rule (MAV), with $f_{\text{MAV}}(r) = r$, (2) Proportional Approval Voting (PAV) with $f_{\text{PAV}}(r) = H(r)$, and (3) Approval Chamberlin–Courant rule (CC), with $f_{\text{CC}}(r) = \min(1, r)$.

4.1 Unknown Distributions

For MAV, the problem of finding the best approximate committee can be casted as a *multiple choice secretary problem* [Kleinberg, 2005]. This is because we can define the value of a candidate c as $N(c)$, the number of voters approving c , and because f_{MAV} is additive in the values.

More generally, for all concave utility functions, we can directly apply the results for the *submodular secretary problem* [Bateni *et al.*, 2013]. This includes not only MAV but also CC and PAV, because f_{CC} and f_{PAV} are submodular set functions. We obtain an online ABC rule with a constant-factor approximation guarantee, which works as follows. It first divides the sequence of candidates into k roughly equal-size parts—the size of each part is between $\lfloor m/k \rfloor$ and $\lceil m/k \rceil$. From each part we select exactly one candidate as follows. Consider the i -th part, and assume the set of $(i-1)$ candidates W_{i-1} has been already selected. To select the i -th candidate we first observe the first $\lceil m/k \rceil$ candidates in the i -th part of the sequence, and find one, call it a_i , that maximizes $f(W_{i-1} \cup \{a_i\})$. Next, we select the first candidate c such that $f(W_{i-1} \cup \{c\}) \geq f(W_{i-1} \cup \{a_i\})$. If we found no such candidate in the i -th part of the sequence, we pick the last candidate from the i -th part, and move to

the next part. This way, we select exactly one candidate from each part of the sequence. By the result of [Bateni *et al.*, 2013], this algorithm returns a committee W such that $\mathbb{E}[f\text{-sc}(W)] \geq \frac{1-1/e}{7} f\text{-sc}(W_{\text{opt}}) \approx 0.09 f\text{-sc}(W_{\text{opt}})$.

The above results are valid for unknown distributions of values of (subset of) candidates, and do not depend on how these values are defined. The interest of the following Section 4.2 is to leverage the specificity of our problem, where the objective function $f\text{-sc}$ depends on multiple voters and is additively decomposable in their utilities. Precisely, the input at each time t is richer than the single value $N(c)$, and more specific than an abstract function $f\text{-sc}(S)$. Rather, we observe at each time the approvals of each voter for the arriving candidate c_t , i.e., n multiple binary valuations. Using this specific problem structure, and an additional assumption of prior knowledge of approval probabilities, we show how dynamic programming can be used to go beyond the off-the-shelf solutions offered by generalized secretary problems.

4.2 Known Distributions

We now assume that we have a known prior distribution: we know the probability $p(i) = p(x \in A_i)$ that voter i approves the next observed candidate x . These probabilities may depend on i , and the events $x \in A_i$ and $x \in A_j$ for different $i \neq j$ need not be independent.

Whether it is realistic to assume we know $p(i)$ depends on the context. If we have a database of past instances on similar problems, then we can compute the approval frequency of voters (or of a given voter, if she appears in several instances and the database is not anonymous).

For the sake of simplicity, we now assume that (UI) $p(i)$ has the same (*uniform*) value p for all voters, and that the events that voter i approves or not candidate j are *independent*. Assumption (UI) is classic in social choice. Note that (UI) is not necessary for some of our results to hold: what we need is only that the probability P_j that a candidate is approved by j voters is polynomial-time computable. Under (UI), $P_j = \binom{n}{j} p^j (1-p)^{n-j}$.

A policy is a function π that decides, at each step when a new candidate comes and once the approvals and disapprovals for this candidate are observed, whether the candidate should be selected or not. More rigorously, the policy maps a *state* to a decision; we postpone the definition of a state because it varies with the rule used.

A history is a sequence of candidates together with associated votes and actions: $h = \langle (c_t, N(c_t), a_t), t = 1, \dots, q \rangle$ for $q \leq m$. c_t is the candidate observed at time t ; $N(c_t)$ is the set of voters who approve c_t ; and $a_t \in \{yes, no\}$ (decision to selecting or not c_t). For instance, $h = \langle (a, \{1, 3, 4\}, yes), (b, \{1, 2\}), no \rangle$ is the history where a is observed, approved by voters 1, 3 and 4, selected, then b is observed, etc. A history is *terminal* if either $q = m$ or the number of selected candidates is k . A policy is *safe* if all its induced histories select exactly k candidates. Provided that $m \geq k$, safe policies exist. Each terminal history h of a safe policy has an associated set of selected candidates $W(h)$ of cardinality k and a reward $f\text{-sc}(W(h))$. A policy induces a probability distribution over histories, which in turns allows to

define its expected score. An *optimal safe policy* is one with maximal expected score:

$$V^* = \max_{\pi} \mathbb{E}_{h \sim \pi} [f\text{-sc}(W(h))]$$

We now show that we can express an optimal safe policy as a mapping from any state to an action *yes* or *no*, and that it can be computed by a dynamic programming algorithm. The exact definition of a state varies with the multiwinner voting rule. A state is *full* if it corresponds to k candidates having been selected already, and *tight* if its associated number of candidates already selected plus the number of candidates yet to be seen is equal to k .

Multiwinner Approval Voting

With MAV, we have $f_{\text{MAV}}(W) = \sum_{c \in W} |N(c)|$. In this case a state is a triple $s = (\alpha, \beta, \gamma)$, where

- $\alpha \in \{1, \dots, m\}$ is the number of candidates seen so far, including the currently observed candidate.
- $\beta \in \{0, \dots, \min(k, \alpha - 1)\}$ is the number of candidates selected so far.
- $\gamma \in \{0, \dots, n\}$ is the number of voters who approve the current candidate.

The number of states is $(n + 1)(k + 1)(m + 1 - \frac{k}{2})$. A state is full if $\beta = k$ and tight if $\beta + m - \alpha + 1 = k$. A safe policy must map every full state to *no* and every tight state to *yes*. Note that for $\alpha = m$, a state obtained by following a safe policy is either tight or full.

Let $V^*(\alpha, \beta, \gamma)$ be the expected score of an optimal safe policy from (α, β, γ) .

Under (UI), V^* satisfies the Bellman equations

$$\begin{aligned} V(\alpha, \beta, \gamma, no) &= \sum_{j=0}^n P_j V^*(\alpha + 1, \beta, j) \\ V(\alpha, \beta, \gamma, yes) &= \gamma + \sum_{j=0}^n P_j V^*(\alpha + 1, \beta + 1, j) \\ V^*(\alpha, \beta, \gamma) &= \max(V(\alpha, \beta, \gamma, no), V(\alpha, \beta, \gamma, yes)) \end{aligned}$$

$V(\alpha, \beta, \gamma, yes)$ and $V(\alpha, \beta, \gamma, no)$ are the expected utilities obtained when selecting (resp., not selecting) the current candidate in state (α, β, γ) and then following an optimal safe policy. (For a detailed justifications, see the Appendix). Thus, the optimal safe policy can be computed by dynamic programming by iterating on all states from $\alpha = m$ down to $\alpha = 1$. There are $O(n^2 km)$ states and each state needs a summation over n terms.

Proposition 4.1. *For MAV, under assumption (UI), an optimal safe policy can be computed in time $O(n^2 km)$.*

Chamberlin-Courant Approval Voting

Recall that $f_{\text{CC}}(W) = |\{i \in N : W \cap A(i) \neq \emptyset\}|$.

Proposition 4.2. *For CCAV, under assumption (UI), an optimal policy can be computed in time $O(n^3 km)$.*

The proof is similar to that of Proposition 4.1, except that the state space is larger. The marginal value of a candidate x when the current set of selected candidates is S is 1 if some of the voters who have no approved candidate in S approves x , and 0 otherwise. This means that to be able to determine the marginal value of a new candidate, it is necessary to know the number of voters, among those who have disapproved all

candidates selected so far, who approve the currently observed candidate. So now a state is a tuple $s = (\alpha, \beta, \gamma, \delta)$, where α and β are as before, $\delta \in \{0, \dots, n\}$ is the number of all “unsatisfied” voters, i.e., those that approve no candidate in the current selection, and $\gamma \in \{0, \dots, \delta\}$ is the number of those unsatisfied voters who approve the current candidate.

The optimal policy is computed by dynamic programming, iterating on all states, with the Bellman equations

$$\begin{aligned} V(\alpha, \beta, \gamma, \delta, no) &= \sum_{i=0}^{\delta} p(i, \delta) V^*(\alpha + 1, \beta, i, \delta) \\ V(\alpha, \beta, \gamma, \delta, yes) &= \\ &\quad \gamma + \sum_{i=0}^{\delta-\gamma} p(i, \delta - \gamma) V^*(\alpha + 1, \beta + 1, i, \delta - \gamma) \\ V^*(\alpha, \beta, \gamma, \delta) &= \max(V(\alpha, \beta, \gamma, \delta, no), V(\alpha, \beta, \gamma, \delta, yes)) \end{aligned}$$

where $p(i, \delta) = \binom{\delta}{i} p^i (1 - p)^{\delta - i}$. Now there are $\Theta(n^2 km)$ states, therefore the algorithm runs in $O(n^3 km)$.

PAV and General Thiele Rules

The states used for CCAV are no longer sufficient: to know the marginal gain for voter i , of the current candidate, relatively to the current selection, we must store the number of candidates already selected approved by i .

The number of states is now in the order of mk^n . The dynamic programming algorithm still works but runs in time exponential in n . Consequently, we get polynomial-time computability if the number of voters is constant.

Theorem 4.1. *For all Thiele rules, including PAV, f_{CCAV} , if the number of voters is constant then an optimal safe policy can be computed in polynomial time.*

Small values of n are realistic: we can think of a small jury, or of the interpretation of voters as criteria.

A subclass of Thiele rules for which the problem is still tractable consists of rules such that the number of values of the score vector is bounded by a constant. This is the case for MAV and CCAV, but also for other rules such as truncated PAV, defined by the vector $(1, 1/2, 0, \dots, 0)$.

5 Proportionality

In this section we focus on the concept of proportionality. Our goal is to design online ABC rules which would guarantee each minority of the voters the right to decide about a part of the elected committee.

Recall that k is the committee size. For an integer $\ell \in [k]$ we say that a group of voters $S \subseteq N$ is ℓ -cohesive if (1) it is large enough, $|S| \geq \ell \cdot n/k$, and (2) its members approve of at least ℓ common candidates, $|\bigcap_{i \in S} A(i)| \geq \ell$. We extend this notion, and define an approximate variant of ℓ -cohesiveness. Given an $\alpha > 1$ we say that a group S is α - ℓ -cohesive if (1) $|S| \geq \alpha \cdot \ell \cdot n/k$, and (2) $|\bigcap_{i \in S} A(i)| \geq \ell$.

The two notions of proportionality that are commonly considered in the literature are proportional justified representation (PJR) [Sánchez-Fernández *et al.*, 2017], and extended justified representation (EJR) [Aziz *et al.*, 2017]. Below we define their approximate variants.

Definition 5.1 (Proportional justified representation). *Given an $\alpha > 1$ we say that a committee W satisfies an α -Proportional Justified Representation (α -PJR) if for each $\ell \in [k]$ and each α - ℓ -cohesive group of voters S it holds that $|\bigcup_{i \in S} A(i) \cap W| \geq \ell$.*

Analogously, we define the axiom of α -EJR.

Definition 5.2 (Extended justified representation). *Given an $\alpha > 1$ we say that a committee W satisfies an α -Extended Justified Representation (α -EJR) if for each $\ell \in [k]$ and each α - ℓ -cohesive group of voters S there exists a voter $i \in S$ who approves of at least ℓ committee members, i.e., $|A(i) \cap W| \geq \ell$.*

We say that a committee election rule satisfies α -PJR if each committee returned by the rule satisfies α -PJR. Analogously, we define what it means that a rule satisfies α -EJR. These axioms form a hierarchy: if a rule satisfies α -EJR then it also satisfies α -PJR. If a rule satisfies α -EJR (respectively, α -PJR) for $\alpha = 1$ then we simply say that the rule satisfies EJR (respectively, PJR). EJR is a very strong axiom and for the time being it is known to be satisfied only by PAV [Aziz *et al.*, 2017] and Rule X [Peters and Skowron, 2020]. Further, Sequential Phragmén’s Rule satisfies 2-EJR [Skowron, 2021].

5.1 Proportional Justified Representation

Somehow surprisingly, it appears that the axiom of PJR can be satisfied in the online setting by the following Greedy Budgeting Rule. Each voter is initially given 1 dollar. When a candidate $c \in C$ arrives we look if the voters who approve c have at least n/k dollars in total. If so, we add c to the committee and ask the voters from $N(c)$ to pay n/k . The properties of the algorithm do not depend on how spread the cost of n/k among the voters from $N(c)$, but a fair policy would suggest to do it as evenly as possible. This way the rule would resemble the method of equal shares [Peters and Skowron, 2020; Pierczyński *et al.*, 2021].

Since the voters have in total n dollars, and buying each candidate costs n/k , it is clear that the rule cannot select more than k candidates. If it picks less, we can add the last candidates that appear, so that exactly k of them are selected.

Theorem 5.1. *The Greedy Budgeting rule satisfies PJR.*

Proof. Consider an election $E = (C, N, k)$, and towards a contradiction suppose the committee W returned by the Greedy Budgeting Rule does not satisfy PJR. Let S be an ℓ -cohesive group such that $|\bigcup_{i \in S} A(i) \cap W| < \ell$.

Each time we select a candidate, we ask the voters to pay exactly n/k . Since $|\bigcup_{i \in S} A(i) \cap W| \leq \ell - 1$ we asked the voters from S to pay at most $(\ell - 1) \cdot n/k$. Since $|S| \geq \ell \cdot n/k$ we get that the voters from S have at least n/k dollars at each step of the rule. Consequently, each time when a candidate from $\bigcap_{i \in S} A(i)$ appears, these voters have enough money to buy it. As a result, each candidate from $\bigcap_{i \in S} A(i)$ would be selected. There are at least ℓ such candidates. This gives a contradiction and completes the proof. \square

5.2 An Online Algorithm with $H(k)$ -EJR

We now move to the case of extended justified representation (EJR). We start by defining the Online Greedy Cohesive algorithm (OGCA), and next we will prove that OGCA satisfies $H(k)$ -EJR (note that OGCA favors large groups of voters and does not satisfy JR).

Consider a candidate $c \in C$ that arrives. If there exists a group of voters $S \subseteq N(c)$ with $|S| \geq H(k) \cdot \ell \cdot n/k$ such that each voter from S approves less than ℓ candidates selected

so far, then OGCA accepts c . Otherwise, c is rejected. If the rule were to select less than k candidates, the candidates that arrived last are accepted so that the committee seats are filled.

Theorem 5.2. *For each election E , OGCA returns a size- k committee that satisfies $H(k)$ -EJR.*

Proof. The fact that the algorithm satisfies $H(k)$ -EJR follows directly from its definition, we will prove that it selects at most k candidates using a budgeting argument.

With each candidate c we associate the price of n/k . When the algorithm accepts c , its cost is spread equally among the voters who approve it. Notice that for any $\ell \in [k]$, each voter buys at most ℓ candidates forming an $H(k)$ - ℓ -cohesive group. For any such candidate c , $|N(c)| \geq H(k) \cdot \ell \cdot n/k$.

For $\ell = 1$, each voter buys at most one candidate c forming an $H(k)$ -1-cohesive group, and in such a case the voter pays at most $n/k \cdot \frac{1}{H(k) \cdot 1 \cdot n/k} = \frac{1}{H(k)}$.

For $\ell = 2$, a voter buys at most two candidates forming an $H(k)$ -2-cohesive group. One of such candidates could have been bought before (as a candidate forming a $H(k)$ -1-cohesive group), and the voter would pay $\frac{1}{H(k)}$ for it. For the second candidate, the voter would pay at most $n/k \cdot \frac{1}{H(k) \cdot 2 \cdot n/k} = \frac{1}{2H(k)}$.

Repeating the reasoning for $\ell = 1, \dots, k$, we get that each voter paid at most: $\sum_{\ell=1}^k 1/\ell H(k) = 1$. Thus, the total amount of money paid is at most equal to n . Since each candidate costs $\frac{n}{k}$, our algorithm could have selected at most k candidates. \square

Interestingly, in terms of proportionality guarantees the Online Greedy Cohesive algorithm is optimal.

Theorem 5.3. *For each $\epsilon > 0$ there exists no online ABC rule that would satisfy $(1 - \epsilon)H(k)$ -EJR.*

Proof. For the sake of contradiction assume that there exists an algorithm \mathcal{A} that satisfies $(1 - \epsilon)H(k)$ -EJR for some rational $\epsilon > 0$.

Let us fix k , and assume the number of voters n is such that $(1 - \epsilon) \cdot n/k$ is an integer. In the first round there arrive candidates who are approved by $(1 - \epsilon)H(k) \cdot n/k$ voters. Each such candidate is approved by a disjoint group of voters. Assume that the number of such candidates equals m_1 , where:

$m_1 = \left\lfloor \frac{n}{(1-\epsilon)H(k) \cdot n/k} \right\rfloor \geq \frac{k}{(1-\epsilon)H(k)} - 1$. Note that each such a candidate must be selected by \mathcal{A} . Indeed, if one of them were not selected, the algorithm could violate $(1 - \epsilon)H(k)$ -EJR. This could happen, for example, if all the remaining candidates that have not yet arrived were approved by no voters.

In the second round there arrive m_2 candidates, each approved by a different group of $2(1 - \epsilon)H(k) \cdot n/k$ voters, where: $m_2 \geq k/2(1 - \epsilon)H(k) - 1$. By the same argument as before, we infer that \mathcal{A} must accept each such a candidate.

Analogously, in the i -th round, $i \leq \frac{k}{(1-\epsilon)H(k)}$ there arrive m_i candidates: $m_i \geq \frac{k}{i(1-\epsilon)H(k)} - 1$, and each of them must be accepted by \mathcal{A} . In total \mathcal{A} must have accepted the following number of candidates:

$$m = \sum_{i=1}^{\left\lfloor \frac{k}{(1-\epsilon)H(k)} \right\rfloor} m_i.$$

In the following sequence of estimations we use the fact that for each i it holds that $\log(i) \leq H(i) \leq \log(i) + 2$:

$$\begin{aligned}
 m &= \sum_{i=1}^{\lfloor \frac{k}{(1-\epsilon)H(k)} \rfloor} m_i \geq \sum_{i=1}^{\lfloor \frac{k}{(1-\epsilon)H(k)} \rfloor} \left(\frac{k}{i(1-\epsilon)H(k)} - 1 \right) \\
 &\geq \frac{k}{(1-\epsilon)H(k)} \cdot H\left(\left\lfloor \frac{k}{(1-\epsilon)H(k)} \right\rfloor\right) - \left\lfloor \frac{k}{(1-\epsilon)H(k)} \right\rfloor \\
 &\geq \frac{k}{(1-\epsilon)H(k)} \cdot \left(\log\left(\frac{k}{(1-\epsilon)H(k)}\right) - 2 \right) \\
 &\geq \frac{k}{(1-\epsilon)H(k)} \cdot \left(H(k) - \log((1-\epsilon)H(k)) - 4 \right) \\
 &\geq \frac{k}{(1-\epsilon)} \left(1 - \frac{\log(H(k)) + 4}{H(k)} \right).
 \end{aligned}$$

Note that $\lim_{k \rightarrow \infty} \frac{\log(H(k)) + 4}{H(k)} = 0$, thus for sufficiently large k the last expression in the sequence of inequalities is larger than k . We infer that \mathcal{A} would need to select more than k candidates, a contradiction. \square

5.3 A Polynomial-Time Algorithm Satisfying $w(k)^2$ -EJR

The OGCA algorithm presented in Section 5.2 cannot be computed in a polynomial time. This is because checking if there exists an ℓ -cohesive group is NP-hard [Skowron *et al.*, 2017]. In this section we define an algorithm that runs in polynomial time, and which offers only a slightly worse EJR guarantee than OGCA.

Our algorithm, which we call Subcommittees via Greedy Budgeting Rule (SGBR), is defined as follows. Let $\alpha = \lceil w(k) \rceil$. The idea is to independently elect α smaller committees, each of size $k' = \lfloor k/\alpha \rfloor$. We elect the i -th subcommittee, $i \in [\alpha]$, using the Greedy Budgeting Rule, but with a constraint that we can pick only the candidates who are approved by at least $\frac{n\alpha^i}{k}$ voters.

Formally, we assume that each voter is given an initial budget of $(1, \dots, 1) \in [0, 1]^\alpha$, that is α independent coins. The i -th coin can be used for buying the candidates who are approved by at least $\frac{n\alpha^i}{k}$ voters. Each candidate costs $\frac{n\alpha}{k}$ coins. When a candidate $c \in C$ arrives, we find the largest pair $i \in N$ and $S \subseteq N(c)$ (we first maximize i , and second $|S|$) such that: (1) $|S| \geq \frac{n\alpha^i}{k}$, (2) Each voter from S has at least $\frac{n\alpha}{k|S|}$ coins of type i left. That is, those voters can afford to buy candidate c assuming each of them paid with the coins of type i , and each would pay the same amount of money. If such pair (i, S) does not exist, we reject c . Otherwise, c is accepted and we ask each voter from S to pay $\frac{n\alpha}{k|S|}$ for c .

Since each voter has in total α coins, and buying each candidate costs $\frac{n\alpha}{k}$ the algorithm selects $\leq k$ candidates.

Theorem 5.4. *SGBR satisfies $\lceil w(k) \rceil^2$ -EJR.*

Proof. For the sake of contradiction assume that given an election $E = (C, N, k)$ Subcommittees via Greedy Budgeting returns a committee W that fails α^2 -EJR. Let S be a subset of α^2 - ℓ -cohesive voters such that $|S| = \alpha^2 \ell \frac{n}{k}$ and $|\cap_{i \in S}$

$A(i) \geq \ell$ for some $\ell \in [k]$, and that for all $v \in S$, we have $|A(v) \cap W| < \ell$.

There exists $j \in [\alpha]$ such that $\frac{n}{k} \alpha^j \leq |S| \leq \frac{n}{k} \alpha^{j+1}$. From that it follows that $\ell \leq \alpha^{j-1}$. Let W_j be the j -th subcommittee. For each elected candidate from W_j a single voter can pay at most:

$$\frac{\frac{n\alpha}{k}}{\frac{n}{k} \alpha^j} = \frac{1}{\alpha^{j-1}} \leq \frac{1}{\ell}.$$

Since each voter from S approves at most $(\ell - 1)$ candidates in W_j , they paid at most $(\ell - 1) \frac{1}{\ell}$ and their remaining budget is greater than $\frac{1}{\ell} \geq \frac{1}{\alpha^{j-1}}$. Thus, when a candidate from $\cap_{i \in S} A(i) \setminus W_j$ appears there are at least $|S|$ voters, $|S| \geq \frac{n}{k} \alpha^j$, each having at least $\frac{1}{\alpha^{j-1}}$ coins of type j left. Thus, their total budget is sufficient to buy the candidate: $|S| \cdot 1/\alpha^{j-1} \geq n\alpha^j/k \cdot 1/\alpha^{j-1} = n\alpha/k$.

Consequently, each candidate from $\cap_{i \in S} A(i)$ would be selected. There are at least ℓ such candidates, thus $|A(v) \cap W| \geq \ell$ for each $v \in S$. This gives a contradiction and completes the proof. \square

Remark. While EJR always implies PJR, α -EJR does not necessarily imply PJR. Still, α -EJR provides desirable guarantees even when PJR is not satisfied. PJR is a property that is often used both in the context of proportionality and of diversity. If our primary focus is to provide voters from large cohesive groups with multiple representatives, then α -EJR should be considered. For example, assume there are n voters, $2n$ candidates and $k = n$ is the desired committee size. Assume that each of the first n candidates that arrive is approved by exactly one voter, and each such candidate is approved by a different voter. Further, assume that each of the next n candidates is approved by all n voters. Then the committee that consists of the first n candidates satisfies PJR (in fact, this would be the committee selected by the Greedy Budgeting Rule). Such a committee would be bad from the perspective of α -EJR, since each voter would have only a single representative. In this case, our algorithms for α -EJR would return clearly better committees.

6 Conclusion

Our main message is that online approval-based committee elections are easier than we might have thought. First, for Thiele rules with submodular score functions we have a constant-factor approximation algorithm (obtained as a corollary of a known result), and a dynamic programming algorithm for maximizing the expected score, which runs in polynomial time for some rules or under some assumptions. Second, we have an algorithm that returns a committee satisfying PJR, and two that return an approximate EJR committee: one with an optimal ratio, the other one running in polynomial-time.

A further work direction consists in moving to ordinal preferences, where voters rank the current candidates in the list of candidates already observed.

Acknowledgements

This work was funded in part by the French government under management of Agence Nationale de la Recherche as part of the "Investissements d'avenir" program, reference ANR-19-P3IA-0001 (PRAIRIE 3IA Institute). We also thank the IJCAI-22 anonymous reviewers for useful comments and suggestions.

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. *48(2)*:461–485, 2017.
- [Babaioff *et al.*, 2008] Moshe Babaioff, Nicole Immorlica, David Kempe, and Robert Kleinberg. Online auctions and generalized secretary problems. *ACM SIGecom Exchanges*, 7(2):1–11, 2008.
- [Banerjee *et al.*, 2022] Siddhartha Banerjee, Vasilis Gkatzelis, Safwan Hossain, Billy Jin, Evi Micha, and Nisarg Shah. Proportionally fair online allocation of public goods. <http://www.cs.toronto.edu/nisarg/papers/fair-online-public-goods.pdf>, 2022. Accessed: 2022-05-05.
- [Bateni *et al.*, 2013] Mohammad Hossein Bateni, Mohammadtaghi Hajiaghayi, and Morteza Zadimoghaddam. Submodular secretary problem and extensions. *ACM Transactions on Algorithms (TALG)*, 9(4):1–23, 2013.
- [Bearden *et al.*, 2005] J. Neil Bearden, Ryan O. Murphy, and Amnon Rapoport. A multi-attribute extension of the secretary problem: Theory and experiments. *Journal of Mathematical Psychology*, 49(5):410–422, 2005.
- [Dey *et al.*, 2017] Palash Dey, Nimrod Talmon, and Otniel van Handel. Proportional representation in vote streams. *arXiv preprint arXiv:1702.08862*, 2017.
- [Do *et al.*, 2021] Virginie Do, Jamal Atif, Jérôme Lang, and Nicolas Usunier. Online selection of diverse committees. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 154–160. International Joint Conferences on Artificial Intelligence Organization, 8 2021. Main Track.
- [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)*, 2017.
- [Hossain *et al.*, 2021] Safwan Hossain, Evi Micha, and Nisarg Shah. Fair algorithms for multi-agent multi-armed bandits. *Advances in Neural Information Processing Systems*, 34, 2021.
- [Israel and Brill, 2021] Jonas Israel and Markus Brill. Dynamic proportional rankings. *arXiv preprint arXiv:2105.08043*, 2021.
- [Kleinberg, 2005] Robert Kleinberg. A multiple-choice secretary algorithm with applications to online auctions. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 630–631. Citeseer, 2005.
- [Lackner and Skowron, 2020] Martin Lackner and Piotr Skowron. Multi-winner voting with approval preferences. Technical Report arXiv:2007.01795 [cs.GT], arXiv.org, 2020.
- [Lackner *et al.*, 2021] Martin Lackner, Jan Maly, and Simon Rey. Fairness in long-term participatory budgeting. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 299–305. International Joint Conferences on Artificial Intelligence Organization, 8 2021. Main Track.
- [Lackner, 2020] Martin Lackner. Perpetual voting: Fairness in long-term decision making. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 2103–2110, 2020.
- [Oren and Lucier, 2014] Joel Oren and Brendan Lucier. Online (budgeted) social choice. In *Twenty-Eighth AAAI Conference on Artificial Intelligence*, 2014.
- [Panigrahi *et al.*, 2012] Debmalya Panigrahi, Atish Das Sarma, Gagan Aggarwal, and Andrew Tomkins. Online selection of diverse results. In *Proceedings of the fifth ACM international conference on Web search and data mining*, pages 263–272, 2012.
- [Peters and Skowron, 2020] Dominik Peters and Piotr Skowron. Proportionality and the limits of welfarism. In *Proceedings of the 2020 ACM Conference on Economics and Computation (ACM-EC-2020)*, pages 793–794, 2020.
- [Pierczyński *et al.*, 2021] Grzegorz Pierczyński, Piotr Skowron, and Dominik Peters. Proportional participatory budgeting with additive utilities. In *Proceedings of the 35th Annual Conference on Neural Information Processing Systems (NeurIPS-2021)*, 2021.
- [Sánchez-Fernández *et al.*, 2017] Luis Sánchez-Fernández, Edith Elkind, Martin Lackner, Norberto Fernández, Jesus A. Fisteus, Pablo Basanta Val, and Piotr Skowron. Proportional justified representation. In *Proceedings of the 31st Conference on Artificial Intelligence (AAAI-2017)*, pages 670–676, 2017.
- [Skowron *et al.*, 2017] Piotr Skowron, Martin Lackner, Markus Brill, Dominik Peters, and Elkind Elkind. Proportional rankings. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI-2017)*, pages 409–415, 2017.
- [Skowron, 2021] Piotr Skowron. Proportionality degree of multiwinner rules. In *Proceedings of the 2021 ACM Conference on Economics and Computation (ACM-EC-2021)*, 2021.
- [Thiele, 1895] Thorvald N. Thiele. Om flerfoldsvalg. In *Over-sigt over det Kongelige Danske Videnskabernes Selskabs Forhandlinger*, pages 415–441. 1895.
- [Yu *et al.*, 2019] Ge Yu, Sheldon Howard Jacobson, and Nigar Kiyavash. A bi-criteria multiple-choice secretary problem. *Iise Transactions*, 51(6):577–588, 2019.