

Proportional Aggregation of Preferences for Sequential Decision Making

NIKHIL CHANDAK, Max Planck Institute for Intelligent Systems, Tübingen AI Center, Germany

SHASHWAT GOEL, Max Planck Institute for Intelligent Systems, ELLIS Institute Tübingen, Germany

DOMINIK PETERS, CNRS, LAMSADE, Université Paris Dauphine - PSL, France

We study the problem of fair sequential decision making given voter preferences. In each round, a decision rule must choose a decision from a set of alternatives where each voter reports their approved alternatives. Instead of going with the most popular choice in each round, we aim for proportional representation across rounds, using axioms inspired by the multi-winner voting literature. The axioms require that every group of $\alpha\%$ of the voters that agrees in every round (i.e., approves a common alternative), must approve at least $\alpha\%$ of the decisions. A stronger version of the axioms requires that every group of $\alpha\%$ of the voters that agrees in a β fraction of rounds must approve $\beta \cdot \alpha\%$ of the decisions. We show that three attractive voting rules satisfy axioms of this style. One of them (Sequential Phragmén) makes its decisions online, and the other two satisfy strengthened versions of the axioms but make decisions semi-online (Method of Equal Shares) or fully offline (Proportional Approval Voting). We present empirical results for these rules based on synthetic data and U.S. political elections. We also run experiments using the moral machine dataset about ethical dilemmas: We train preference models on user responses from different countries and let the models cast votes. We find that aggregating these votes using our rules leads to a more equal utility distribution across demographics than making decisions using a single global preference model.

JAIR Track: Conference Award Track

JAIR Associate Editor: Hadi Hosseini

JAIR Reference Format:

Nikhil Chandak, Shashwat Goel, and Dominik Peters. 2026. Proportional Aggregation of Preferences for Sequential Decision Making. *Journal of Artificial Intelligence Research* 85, Article 17 (February 2026), 39 pages. DOI: [10.1613/jair.1.18660](https://doi.org/10.1613/jair.1.18660)

1 Introduction

We consider the problem of making a sequence of independent decisions via voting. In each round, we can choose one alternative from a set of several alternatives, based on voters who tell us which alternatives they support (or *approve*). The set of voters stays the same across rounds, though the set of alternatives may change. The “standard” way of making such decisions is to take the alternative with the most supporters in each round. A problem with this method is that non-majority groups of voters may have very little influence on the outcomes. For example, if there is a fixed group of 51% of the voters who all report the same opinion in every round, then 100% of the decisions will be taken in accordance with the wishes of that group, with the other 49% of voters ignored or at most acting as tie-breakers. In many contexts, this is undesirable due to fairness concerns.

An earlier version of this paper was published in the Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI 2024). This journal version of the paper features full proofs, improved presentation, and additional detail in the experiments section.

Authors' Contact Information: Nikhil Chandak, ORCID: [0009-0009-9468-7838](https://orcid.org/0009-0009-9468-7838), nikhil.chandak@tuebingen.mpg.de, Max Planck Institute for Intelligent Systems, Tübingen AI Center, Tübingen, Germany; Shashwat Goel, ORCID: [0009-0004-5193-8135](https://orcid.org/0009-0004-5193-8135), shashwatnow@gmail.com, Max Planck Institute for Intelligent Systems, ELLIS Institute Tübingen, Tübingen, Germany; Dominik Peters, ORCID: [0000-0001-9418-7571](https://orcid.org/0000-0001-9418-7571), dominik.peters@lamsade.dauphine.fr, CNRS, LAMSADE, Université Paris Dauphine - PSL, Paris, France.



This work is licensed under a [Creative Commons Attribution International 4.0 License](https://creativecommons.org/licenses/by/4.0/).

© 2026 Copyright held by the owner/author(s).

DOI: [10.1613/jair.1.18660](https://doi.org/10.1613/jair.1.18660)

Following recent work that studies this model under the name “perpetual voting” (Lackner 2020), we define formal properties of voting rules that capture the intuition that a group of $\alpha\%$ of the voters should be able to control the decisions in $\alpha\%$ of the rounds. Inspired by work in multi-winner approval voting (Lackner and Skowron 2023), we then define a number of voting rules that satisfy these properties. We believe that these rules have promising applications in a variety of domains. Here are some examples:

- *Hiring decisions*: Consider a department that hires new faculty each year, with existing faculty voting over the applicants. The department wishes to hire people representing its spectrum of research interests. For example, if 20% of the department works on one topic and votes for a candidate on that topic, then at least 1 such candidate should be hired every 5 years.
- *Virtual democracy*: In cases where a group of people need to make an extremely large number of decisions, we may wish to automate this process. An approach known as virtual democracy does this by initially learning voters’ preferences over a space of potential alternatives (specified by feature vectors) for example based on pairwise comparisons. Then, each decision is made by letting the models vote on the decision maker’s behalf by predicting their preferences. This approach has led to proof-of-concept systems that automate moral decisions faced by autonomous vehicles (Noothigattu et al. 2018), kidney exchanges (Freedman et al. 2020), collective decision making directly from natural language preferences (Mohsin et al. 2021) and allocation of food donations (Lee et al. 2019). The approach behind these systems has recently been criticized as overweighting the opinions of majorities (Feffer et al. 2023). We show preliminary evidence in our experiments that using proportional voting methods could avoid this issue.
- *Combining the outputs of sequence models*: Modern large language models (LLMs) operate in rounds, in each round suggesting a probability distribution over tokens. We could view these distributions as preferences over tokens, and use proportional methods to mix several LLMs together, potentially allowing for fine-grained customization to user preferences, or aiding in drafting compromise documents. D. Peters (2024) discusses this use case in detail.
- *Policy decisions of coalition governments*: In many countries, the government is formed by coalitions of several parties with different strengths (for example, in Germany 2021–25 it consists of 3 parties who received 26%, 15%, and 12% of the vote). A coalition needs to agree on a program, consisting of decisions on many issues. Our methods could help to design a program where each party’s preferences gets a fair representation in the outcome.
- *The parliamentary approach to moral uncertainty*: An agent may have moral uncertainty, and be unsure which of several moral theories is right. Bostrom (2009) has suggested that an agent could decide on what actions to take given this uncertainty by imagining how the decision would be made by a *parliament* in which each moral theory is represented according to the agent’s credence in the theory. The idea is that the representatives in the imaginary parliament would deliberate, bargain, and vote, and thereby reach good compromise decisions (Greaves and Cotton-Barratt 2023; Newberry and Ord 2021). While the philosophical literature discusses using simple proportional voting methods (Newberry and Ord 2021), this work remains preliminary. Proportional multi-issue voting rules may better reflect the work of a parliament that takes many decisions.

Additional examples of applications mentioned in the literature include a company that donates some resources to a charity each year and lets employees or customers voter over which charity to choose (Elkind, Obratzova, et al. 2025), or meal selection by a group of friends (Elkind, Neoh, et al. 2024), or electing a committee where only part of the committee is replaced in each election (Elkind, Obratzova, et al. 2024).

In many of these applications, the different issues are not completely separable, so our model may not be a precise fit. However, when using online rules, the agents can adjust their reported preferences to take some consistency properties into account.

Table 1. Our results. For each setting (online, semi-online (where the rule only knows the total number of rounds in advance), offline (where the rule knows all the preferences in advance)), we list the best known rule with respect to proportionality axioms. ⁺ = knowing an online rule satisfying Weak EJR would resolve a major open problem in the theory of multi-winner voting. [†] = no semi-online rule can satisfy EJR due to a result of [Elkind, Obratzsova, et al. \(2025, Proposition 7.4\)](#).

Voting Rule	Weak PJR	PJR	Weak EJR	EJR
Sequential Phragmén (Online)	✓	✓	✗ ⁺	✗ [†]
Method of Equal Shares (Semi-online)	✓	✗	✓	✗ [†]
PAV (Offline)	✓	✓	✓	✓

Our Results. We build on the work of [Bultheau et al. \(2021\)](#) who defined two axioms that guarantee group representation. We call them PJR (Proportional Justified Representation) and Weak PJR. These axioms do not require the groups to be pre-defined. Instead, the guarantee applies to *all* subsets of voters who support a common alternative in some of the rounds, and the “weak” versions only apply if they support a common alternative in *all* of the rounds. [Bultheau et al. \(2021\)](#) conjectured that no polynomial-time algorithm can find a Weak PJR outcome. We show that the picture is more positive: an attractive polynomial-time voting rule building on ideas of [Phragmén \(1894\)](#) satisfies PJR ([Section 4.1](#)). This rule is fully online, making decisions round by round.

We then define axioms called Weak EJR (Extended Justified Representation), which provide better group representation guarantees than Weak PJR, and analogously EJR. We show that the existence of a fully online rule that satisfies Weak EJR would positively resolve a major open problem in multi-winner approval voting ([Lackner and Skowron 2023](#)). Still, we can define a simple voting rule (a variant of the “Method of Equal Shares”) that satisfies Weak EJR and that is “semi-online”: it needs to know how many rounds there will be in total, but it can be given preferences in an online fashion ([Section 4.2](#)). However, it does not satisfy PJR or EJR. This suggests the question whether we can satisfy EJR in the offline setting, where all preferences are available in advance. In [Section 4.3](#), we show that Proportional Approval Voting (PAV), an offline rule that maximizes a carefully chosen objective function proposed by [Thiele \(1895\)](#), does satisfy all our axioms, including EJR. While the PAV optimization problem is NP-hard, we prove these results also hold for a polynomial-time local search variant of PAV. Finally, in [Section 5](#) we show that further strengthenings of the PJR/EJR axioms are not satisfiable, and that on inputs where a solution exists, online rules still must fail them.

We close with several experimental results on synthetic ([Section 6.1](#)) and real datasets ([Section 6.2](#)), on which we compare our rules to existing rules from the literature (which do not satisfy our axioms). In [Section 6.3](#), we study the *virtual democracy* application based on the *moral machine* dataset ([Awad et al. 2018](#)). We learn separate preference models that predict the respondents’ judgement for each country and use each country’s model as a ‘voter’. We then sample new decision situations and query the models for their top choices among the alternatives in each situation, aggregating the responses using our voting rules to produce a decision sequence. We compare the decisions produced with training a single universal model on equal number of responses from all countries, and taking its top choices as the decisions. We find that when voters do not already have very similar preferences on the issues, the aggregation approach leads to much fairer outcomes than the decisions made by the combined model.

2 Related Work

Perpetual Voting. Our work is closely related to perpetual voting ([Lackner 2020](#)) which studies online rules, while we also consider offline rules. [Lackner \(2020\)](#) focused on *individual* fairness properties, for example requiring that each voter approves at least 1 decision in every time interval of some bounded length. We are interested in

guarantees for *groups* of voters who agree with each other, such that larger groups receive stronger guarantees. Such guarantees have been considered by [Lackner and Maly \(2023\)](#) and [Bulteau et al. \(2021\)](#); we compare them to ours in [Section 3](#). Very recently, [Phillips et al. \(2025\)](#) proposed new axioms (FJR and EJR+) that strengthen our axioms. [Kozachinskiy et al. \(2025\)](#) study the achievable egalitarian welfare in perpetual voting.

Public Decision Making. [Conitzer et al. \(2017\)](#) study “public decision making” which is an offline model of several decisions, where a voter’s total utility is the sum of the utilities obtained in each round. We study the special case where the utility values in each round are restricted to 0 and 1. [Conitzer et al. \(2017\)](#) focus on fairness notions for individuals (not groups) derived from fair division. [Freeman et al. \(2020\)](#), [Skowron and Górecki \(2022\)](#), and [Brill, Markakis, et al. \(2023\)](#) consider the special case where there are exactly two alternatives (“yes/no”) in each round and utilities are 0/1, which is a special case of our setting where fairness properties are easier to obtain.

Repeated Portioning. In our model, we need to choose one alternative in each round. One could define a fractional version of this model, where in each round, we choose a *mixture* of the available alternatives (formally, a probability distribution or lottery over C_j). In the single-round setting, finding a fair mixture has been considered by several papers ([Aziz, Bogomolnaia, et al. 2020](#); [Bogomolnaia et al. 2005](#); [Brandl et al. 2021](#); [Fain et al. 2016](#)). For a multi-round setting with online decisions, this model is studied by [Banerjee et al. \(2023, Section 5\)](#) who consider a proportional fairness objective and propose methods with good regret bounds.

Multi-winner Approval Voting. For the task of selecting a committee of exactly k out of m candidates given approval votes ([Lackner and Skowron 2023](#)), proportionality has been intensely studied. Many axioms have been proposed that formalize the idea that $\alpha\%$ of the voters should be able to elect at least $\alpha\% \cdot k$ of the committee members ([Aziz, Brill, et al. 2017](#); [Brill and J. Peters 2023](#); [D. Peters and Skowron 2020](#); [Sánchez-Fernández, Elkind, et al. 2017](#)). Voting rules proposed long ago by [Thiele \(1895\)](#) and [Phragmén \(1894\)](#) do very well with respect to these axioms ([Brill, Freeman, et al. 2023](#); [Janson 2018](#)). A more recent rule proposed by [D. Peters and Skowron \(2020\)](#) known as the Method of Equal Shares combines some of the advantages of the Phragmén and Thiele rules. [Skowron, Lackner, et al. \(2017\)](#) discuss adapting multi-winner rules to obtain “proportional rankings” of alternatives, which can be seen as a sequential decision making problem. Many of the rules and axioms we consider in this paper are analogues or generalizations of proposals from multi-winner voting.

Approval-based Apportionment. [Brill, Gözl, et al. \(2022\)](#) considered a model where we need to assign k seats to political parties, based on voters who submit approval sets over parties. This can be seen as a kind of multi-winner election where candidates can be added to the committee several times (because a single party can receive multiple seats). Alternatively, it can be seen as a special case of our model with k rounds, where the sets of alternatives and all voters’ approval sets remain the same across all rounds. Being a special case of our model, some negative results of [Brill, Gözl, et al. \(2022\)](#) carry over to our model, notably that Sequential Phragmén fails Weak EJR and that PAV is NP-hard to compute; we will mention these connections during our discussion.

Combinatorial Voting. A classic literature on voting in combinatorial domains ([Lang and Xia 2016](#)) studies the problem of making decisions on several issues. The main focus is on the representation of complex preferences, and the computational problem of finding good outcomes given such preferences, often in the sense of maximizing utilitarian or egalitarian welfare ([Amanatidis et al. 2015](#)). Some of this work considers (conditional) approval preferences ([Barrot and Lang 2016](#); [Barrot, Lang, and Yokoo 2017](#)).

Additional related work can be found in the recent survey by [Elkind, Obraztsova, et al. \(2024\)](#).

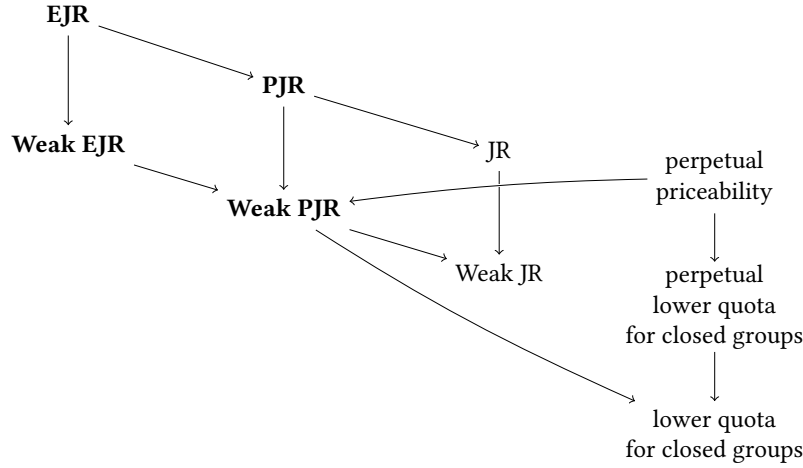


Fig. 1. Implications between axioms. Perpetual priceability and (perpetual) lower quota for closed groups are discussed in Appendix A. We focus on the axioms shown in bold.

3 Problem Formulation

3.1 Model

Let $N = \{1, 2, \dots, n\}$ be the set of voters, with $|N| = n$. There is a set $R = \{1, 2, \dots, T\}$ of $T \geq 1$ rounds, where T is the time horizon. In each round $j \in R$, we are given a non-empty set of alternatives C_j . Each voter $i \in N$ approves some subset A_j^i of C_j (which we do not assume to be non-empty). Thus, we have a sequence of T decision rounds together with a sequence $C = (C_j)_{j \in R}$ of alternative sets and a collection $A = (A_j^i)_{i \in N, j \in R}$ of approval sets. We let (N, R, C, A) denote a decision instance.

A decision sequence $D = (d_1, \dots, d_T) \in C_1 \times \dots \times C_T$ specifies a single decision $d_j \in C_j$ for each round $j \in R$. For a voter $i \in N$, we write $U_D^i = |\{j \in R : d_j \in A_j^i\}|$ for the number of decisions in D that i approves; we treat U_D^i as i 's utility. A decision rule f takes as input a decision instance and returns a decision sequence $D = f(N, R, C, A)$. We call a decision rule *semi-online* if its decision in round $j \in R$ only depends on the time horizon T and on information up to round j , i.e., on N, R, C_1, \dots, C_j and A_1^i, \dots, A_j^i . A rule is *online* if in addition, the decisions are independent of the time horizon T .

3.2 Axioms

We now define several properties (or *axioms*) of decision sequences that formalize the idea of proportional representation. These come in different strengths and Figure 1 shows implication relationships between them. We will focus on the four properties of (Weak) EJR and (Weak) PJR, shown in bold. For completeness, the figure also mentions two properties introduced by Lackner and Maly (2023) (perpetual priceability and lower quota compliance). We discuss these properties, as well as Pareto-efficiency, in Appendix A.

Our first two axioms were introduced by Bulteau et al. (2021).¹ A group $S \subseteq N$ of voters agrees in round $j \in R$ if there is an alternative $c \in C_j$ that all voters in S approve, so $\bigcap_{i \in S} A_j^i \neq \emptyset$. Weak PJR (Weak Proportional Justified Representation, first defined for multi-winner voting by Sánchez-Fernández, Elkind, et al. 2017) requires that a

¹ We use different names for these axioms. Bulteau et al. (2021) say “all periods intersection PJR” for what we call Weak PJR, and say “some periods intersection PJR” for what we call PJR.

Round	1 & 2	3 & 4	5 & 6	7 & 8
Voter 1	{a, b}	{a, b}	{a, b}	{a, b}
Voter 2	{a, c}	{a, c}	{a, c}	{a, c}
Voter 3	{d}	{d}	{d}	{e}
Voter 4	{d}	{d}	{d}	{f}

Fig. 2. Example illustrating our axioms. Voters have the same approval set across the rounds specified in the column header.

group of an α fraction of the voters, if it agrees in every round, must be “happy” in at least $\lfloor \alpha T \rfloor$ rounds, meaning that at least one member of S approves the decision (but this member of S can differ across rounds).

Definition 3.1 (Weak PJR). A decision sequence D satisfies *Weak PJR* if for every $\ell \in \mathbb{N}$ and every group of voters $S \subseteq N$ that agrees in every round and has size $|S| \geq \ell \cdot \frac{n}{T}$, there are at least ℓ rounds $j \in R$ in which the decision d_j of D is approved by some voter in S (i.e., $d_j \in \bigcup_{i \in S} A_i^j$).

Weak PJR only provides guarantees for groups that agree in *all* rounds. PJR also gives guarantees for groups that agree only in some of the rounds, though the groups need to be larger: a group that agrees in (at least) k rounds deserves to be happy with ℓ decisions if the size of the group is at least $\ell \cdot \frac{n}{k}$ (as compared to $\ell \cdot \frac{n}{T}$ for groups that agree in all rounds).² Note that PJR implies Weak PJR (take $k = T$).

Definition 3.2 (PJR). A decision sequence D satisfies *PJR* if for every $k \in \{1, \dots, T\}$, every $\ell \in \mathbb{N}$ and every group of voters $S \subseteq N$ that agrees in k rounds and has size $|S| \geq \ell \cdot \frac{n}{k}$, there are at least ℓ rounds $j \in R$ in which the decision d_j of D is approved by some voter in S (i.e., $d_j \in \bigcup_{i \in S} A_i^j$).

Note that a group’s entitlement ℓ is computed based on the number of rounds in which it agrees, but the ℓ rounds in which the group is satisfied need not come from the set of rounds where it agrees.

An equivalent way of stating the PJR axiom is that an α fraction of the voters who agree in a β fraction of the rounds need to be “happy” with at least an $\alpha \cdot \beta$ fraction of the decisions, in the sense that there are at least $\lfloor \alpha \cdot \beta \cdot T \rfloor$ decisions that are approved by at least one group member. For example, consider a group $S \subseteq N$ of voters with fixed size $\ell \cdot \frac{n}{T}$, and let us ask what PJR guarantees for this group. If S agrees in all rounds, then it says that S should be happy with ℓ decisions. If S agrees in $T/2$ rounds, then it says that S should be happy with $\lfloor \frac{\ell}{2} \rfloor$ decisions. Thus, the guarantee provided to a group is proportional to the number of rounds in which the group agrees. Yet another way to think about this axiom is to pretend that only the rounds in which S agrees exist, and then apply Weak PJR to that situation, while allowing the group to be satisfied also by decisions in rounds where it doesn’t agree.

To illustrate the axioms, consider the instance shown in Figure 2, with $T = 8$ rounds and 4 voters. The group $S = \{1, 2\}$ agrees in all rounds (always approving a). Thus, to satisfy Weak PJR, in at least $\ell = 4$ rounds the outcome needs to be either a , b , or c (because $|S| \geq \ell \cdot \frac{4}{8}$). The group $S' = \{3, 4\}$ agrees in the first 6 rounds (approving d), so with $\ell' = 3$, because $|S'| \geq \ell' \cdot \frac{4}{6}$, PJR requires that in at least 3 rounds, the outcome is either d , e , or f .

A weakness of Weak PJR and PJR is in how they define S being “happy” with a decision (“at least one member of S approves the decision”). This definition can be satisfied by a decision sequence that gives each member of S a utility that is much lower than ℓ (D. Peters and Skowron 2020, Sec. 4.2). In the example of Figure 2,

²Strengthening this axiom to only require the group size to be at least $\ell \cdot \frac{n}{T}$ gives rise to an axiom that may be impossible to satisfy, see Corollary B.2 in the appendix.

the decision sequence (b, b, c, c, d, d, e, f) satisfies PJR, but the first two voters each only approve the decision in 2 rounds, instead of in 4 rounds. Following Aziz, Brill, et al. (2017), we can fix this by defining the axioms EJ_R (Extended Justified Representation) and Weak EJ_R which require that at least one member i of S must approve at least ℓ of the decisions, i.e., must have utility $U_D^i \geq \ell$. In the example, the decision sequence (d, d, d, d, a, a, a, a) satisfies EJ_R.

Definition 3.3 (Weak EJ_R). A decision sequence D satisfies *Weak EJ_R* if for every $\ell \in \mathbb{N}$ and every group of voters $S \subseteq N$ that agrees in all rounds and that has size $|S| \geq \ell \cdot \frac{n}{T}$, there is a voter $i \in S$ who approves at least ℓ decisions in D , i.e., $U_D^i \geq \ell$.

Definition 3.4 (EJ_R). A decision sequence D satisfies *EJ_R* if for every $k \in \{1, \dots, T\}$, every $\ell \in \mathbb{N}$ and every group of voters $S \subseteq N$ that agrees in k rounds and has size $|S| \geq \ell \cdot \frac{n}{k}$, there is a voter $i \in S$ who approves at least ℓ decisions in D , i.e., $U_D^i \geq \ell$.

In analogy to the multiwinner literature, we can also define the notions of JR and Weak JR, which are the special cases of the PJ_R/EJ_R notions with $\ell = 1$. These axioms are also discussed by Bulteau et al. (2021).

Definition 3.5 (Weak JR). A decision sequence D satisfies *Weak JR* if for every $S \subseteq N$ with $|S| \geq \frac{n}{T}$ that agrees in all rounds, there is a voter $i \in S$ with $U_D^i \geq 1$.

Definition 3.6 (JR). A decision sequence D satisfies *JR* if for every $k \in \{1, \dots, T\}$ and every $S \subseteq N$ with $|S| \geq \frac{n}{k}$ that agrees in k rounds, there is a voter $i \in S$ with $U_D^i \geq 1$.

These notions do not guarantee *proportional* representation since larger groups do not obtain higher guarantees, but they do guarantee some minimal representation.

Elkind, Obraztsova, et al. (2025) studied the computational complexity of verifying whether a given decision sequence satisfies a given axiom. They found that this problem is coNP-complete for each of the six axioms, except that verifying Weak JR is tractable for binary issues, i.e., when in each round there are only two alternatives.

A decision rule f satisfies one of these axioms if for all possible inputs, the decision sequence selected by f satisfies it. Note that if f is also online, this means that the proportionality guarantee thereby not only holds for the entire decision sequence, but also for every prefix of it, if that prefix is considered as its own instance.

3.3 Methods

We now define three decision rules that are natural analogues of rules that were first proposed for multi-winner elections: the Sequential Phragmén rule, the Method of Equal Shares, and Proportional Approval Voting. For all the rules, our proportionality guarantees will apply no matter how ties are broken during their execution.

3.3.1 Sequential Phragmén. Phragmén (1894) proposed an approval-based voting method for electing members of the Swedish parliament. Lackner and Maly (2023) adapted this rule to the context of perpetual voting, calling their adaptation “perpetual Phragmén”. We follow their definition. The rule makes decisions round by round. Each decision provides a value (or *load*) of 1, which is distributed to voters who approve the decision. In each round, the rule chooses an alternative such that no voter ends up with too much load, thereby prioritizing voters who do not yet agree with many prior decisions. Formally, each voter i starts with load $x^i = 0$. Sequential Phragmén chooses the alternative for which it can distribute a load of 1 in a way that minimizes the maximum total load assigned to a voter: at each round $j \in R$, we compute the following value for each alternative $c \in C_j$:

$$s_c = \min_{S \subseteq \{i \in N : c \in A^i\}} \frac{\sum_{i \in S} x^i + 1}{|S|}.$$




Fig. 3. Load distribution

It is easy to see that the set S minimizing this quantity will consist of supporters of c with lower current load than the supporters outside S .

The value s_c can be understood using a “water filling” analogy, as shown in Figure 3, where we consider an alternative c approved by 7 voters, and show their current loads as bars. We then fill 1 unit of water on top of the approving voters’ loads. Note that the water never falls on top of the loads of voters 5, 6, and 7 because their load is already quite high; in other words, this process has only assigned load to the set $S = \{1, 2, 3, 4\}$. Then s_c is the “water line”, which is the load of each voter in S after the load of c has been assigned.³

The decision for round $j \in R$ is the alternative c that minimizes s_c , breaking ties arbitrarily. After making the decision, we update loads by setting $x^i = s_c$ for each voter $i \in S$ (where S is the coalition attaining the minimum in the definition of s_c) and leave x^i unchanged for voters not in S . Note that the load of a voter will never decrease.

Clearly, Sequential Phragmén is an online rule. Lackner and Maly (2023) illustrate how the rule works using a detailed example and show that it can be computed in polynomial time (via sorting approvers by load, then scanning through prefixes of that ordering).

3.3.2 Method of Equal Shares (MES). MES is a recently introduced rule for multi-winner voting (D. Peters and Skowron 2020) and also used in practice for participatory budgeting (D. Peters, Pierczyński, et al. 2021). It can be adapted to our setting in a semi-online fashion: the rule needs to know the total number of rounds T in advance, but does not need to know voter preferences of future rounds. Each decision costs $p = \frac{n}{T}$ units, which must be paid by voters that approve the chosen alternative. MES works by subtracting this amount from an initial budget $b_i = 1$ assigned to each voter i .

For $\rho \geq 0$, an alternative $c \in C_j$ is called ρ -affordable if

$$\sum_{i \in N: c \in A_j^i} \min(b_i, \rho) \geq p,$$

so approvers of c can pay p with no one paying more than ρ . In each round, the alternative $d_j \in C_j$ that is ρ -affordable for minimum ρ is chosen, breaking ties arbitrarily. Note that at the minimum ρ , the above inequality holds with equality (since the left-hand side is continuous in ρ). Then for each voter i approving d_j , the remaining budget b_i of i is set to $\max(0, b_i - \rho)$; note that the total amount collected is exactly p .

If in some round, no alternative is affordable for any ρ , MES *terminates prematurely*. Decisions for the remaining rounds can be made arbitrarily (as far as our axioms are concerned), but in practice are done using an appropriate *completion rule* such as Sequential Phragmén. In our experiments, we followed the “ ϵ -completion” strategy introduced by D. Peters, Pierczyński, et al. (2021, Sec. 3.4) which approximates the concept of choosing the lowest affordable alternative, retaining the core idea of MES.

³In the multi-winner setting, we can fix S in the definition of s_c to be the set of *all* voters approving c , because it cannot happen that an approver of c already has more load than s_c , for in that case the rule would have chosen c in an earlier iteration (Brill, Freeman, et al. 2023, Lemma 4.5). This is not true in our setting because c may not have existed in a prior round. Other definitions of Phragmén’s method based on virtual bank accounts (Janson 2018; D. Peters and Skowron 2020) do not easily adapt to our setting because of this phenomenon.

It is also possible to run MES in an offline fashion (“Offline MES”), by letting it choose at each step in which round it would like to make a decision (namely, the round where some alternative is ρ -affordable for minimal ρ , across all alternatives in all rounds where no decision has been taken yet). All our proofs and counterexamples in Section 4.2 below go through for Offline MES, so it also satisfies Weak EJr but fails PJr. This also follows from the results of Masarík et al. (2024), who study Offline MES in a more general setting.

3.3.3 *Proportional Approval Voting (PAV)*. PAV (based on a rule of Thiele 1895) is an offline rule that selects the decision sequence $D \in C_1 \times \dots \times C_T$ that maximizes

$$\text{PAV-score}(D) = \sum_{i \in N} 1 + \frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{U_D^i},$$

where we recall that U_D^i is the number of rounds in which voter i approves the decision of D . (For agents with $U_D^i = 0$, their contribution to the PAV score is 0.) This harmonic objective function is the unique additive objective that leads to a proportional rule (Aziz, Brill, et al. 2017). Finding the optimum decision sequence for PAV is NP-hard (Brill, Gözl, et al. 2022, Thm. 5.1), just like in the multi-winner setting (Aziz, Gaspers, et al. 2015).

However, one can define a polynomial-time local search variant of PAV, that in the multi-winner setting has been found to satisfy the same proportionality axioms as PAV (Aziz, Elkind, et al. 2018).⁴ The variant starts with an arbitrary decision sequence and keeps changing the decision in some round if this increases the PAV-score of the decision sequence by at least n/T^2 .⁵ The process terminates if there is no possible change that leads to a sufficient increase.

Definition 3.7 (Local-Search PAV). Given an instance (N, R, A, C) with n voters and time horizon T , a decision sequence D is a possible output of *Local-Search PAV* if there exists some initial decision sequence D_{init} such that D can be returned by the following algorithm:

```

D ← Dinit;
while there is a round  $j \in R$  and an alternative  $a_j \in C_j$  such that
  PAV-score( $D \setminus \{d_j\} \cup \{a_j\}$ ) ≥ PAV-score( $D$ ) +  $\frac{n}{T^2}$  do
  |  $D \leftarrow D \setminus \{d_j\} \cup \{a_j\}$ 
return D

```

The threshold of n/T^2 is large enough to ensure termination in polynomial time.

PROPOSITION 3.8 (AZIZ, ELKIND, ET AL. 2018). *Local-Search PAV terminates in polynomial time.*

PROOF. Note that a single improving swap can be found and executed in polynomial time, by iterating through all rounds $j \in R$ and all alternatives $a_j \in C_j$. Now, let us assess how many improvements the local search algorithm may perform. Each improvement increases the total PAV-score of the decision sequence by at least n/T^2 . The maximum possible PAV-score of a length- T decision sequence is $n \cdot (1 + 1/2 + \dots + 1/T) = O(n \ln T)$. Thus, there can be at most $O(\frac{n \ln T}{n/T^2}) = O(T^2 \ln T)$ improving swaps. □

4 Satisfying the Axioms

In this section, we establish which of our three rules satisfy which of our four axioms (see Table 1). Bulteau et al. (2021) conjectured that no polynomial-time rule achieves Weak PJr (they gave existence proofs which were based

⁴One could use a sequential version of PAV to get an online rule, which Page et al. (2020) conjectured to satisfy at least Weak JR (a weakening of Weak PJr). However, a counterexample from multiwinner voting (Sánchez-Fernández, Elkind, et al. 2017, Table 2) can be adapted to show that Sequential PAV fails Weak PJr in our model (repeat the profile from their paper for $k = 6$ rounds).

⁵This threshold mirrors the n/k^2 threshold suggested by (Aziz, Elkind, et al. 2018) for multi-winner voting, but see D. Peters (2025, Remark 4.2) for a situation where a somewhat smaller threshold is necessary.

on exponential-time algorithms). Our results show that, in fact, Weak PJR as well as the stronger axioms can be satisfied in polynomial time, by attractive rules, some of which even work online.

4.1 Online: Sequential Phragmén

We begin by analyzing the sequential Phragmén rule.

THEOREM 4.1. *Sequential Phragmén satisfies PjR.*

PROOF. For a contradiction, suppose that $S \subseteq N$ witnesses a violation of PJR: it agrees in rounds $R^* = \{j_1, \dots, j_k\} \subseteq R$ and has size $|S| \geq \ell \cdot \frac{n}{k}$, but there are fewer than ℓ rounds in which at least one member of S approves the decision.

First, we claim that if $d_j \in C_j$ is the alternative chosen in some round $j \in R^*$, then $s_{d_j} \leq \frac{\ell}{|S|}$. Note that the total load $\sum_{i \in S} x^i$ assigned to members of S is at most $\ell - 1$ because each decision incurs a load of 1 and no load is assigned to members of S in rounds where no member approves the decision. Since $j \in R^*$, there is an alternative $c' \in C_j$ that everyone in S approves. Hence,

$$s_{c'} \leq \frac{1 + \sum_{i \in S} x^i}{|S|} \leq \frac{1 + (\ell - 1)}{|S|} = \frac{\ell}{|S|},$$

where the first inequality follows from the definition of $s_{c'}$ as a minimum. Since Sequential Phragmén selected d_j , we have $s_{d_j} \leq s_{c'}$, showing our claim.

Call a round j a *bad round* if $j \in R^*$ and the decision d_j is not approved by any voter in S . Fix some voter $i \in N \setminus S$, and suppose i gets assigned some load during at least one bad round. Consider the point just after the last bad round j where i is assigned some load. At this point we must have $x^i \leq \frac{\ell}{|S|}$ since otherwise $s_{d_j} > \frac{\ell}{|S|}$, contradicting our claim. Thus, at most $\frac{\ell}{|S|}$ load was assigned to i during all bad rounds together. Clearly, this last claim is also true for voters $i \in N \setminus S$ who do not get assigned any load during any bad round.

In a bad round, load is only assigned to voters outside S (since the round is bad). Thus, by summing over all $i \in N \setminus S$, we see that the total load assigned in bad rounds is at most

$$|N \setminus S| \cdot \frac{\ell}{|S|} = \frac{|N|}{|S|} \cdot \ell - \frac{|S|}{|S|} \cdot \ell \leq \frac{k}{\ell} \ell - \ell = k - \ell.$$

However, there are at least $k - (\ell - 1)$ bad rounds, so a total load of at least $k - \ell + 1$ is distributed across bad rounds, a contradiction. \square

In multi-winner voting, Sequential Phragmén is well-known to fail EJR (Brill, Freeman, et al. 2023), so it is unsurprising that it also fails (Weak) EJR in our setting. While it is possible to adapt the example of Brill, Freeman, et al. (2023) to our setting, we use a somewhat simpler example.⁶

Rounds	1–10
Voters 1, 2, 3	{a, b}
Voters 4, 5, 6, 7	{a, c}
Voters 8, 9	{b, c}
Voter 10	{b}

Fig. 4. Instance where Sequential Phragmén fails Weak EJR. The approval sets are the same in all rounds.

⁶This example can also be used to show that Sequential Phragmén fails EJR for the related settings of approval-based apportionment and of multi-winner voting. We thus add a different EJR counterexample to the literature, which may turn out useful in future work.

THEOREM 4.2. *Sequential Phragmén fails Weak EJR.*

PROOF. Consider the instance shown in Figure 4, with $T = 10$ rounds and $n = 10$ voters. Here, the same profile is repeated for all the rounds, so it is also an example where Sequential Phragmén fails EJR in the setting of approval-based apportionment (Brill, Gözl, et al. 2022). In this example, Sequential Phragmén alternates between two alternatives as winner in each round and produces the decision sequence $D = (a, b, a, b, a, b, a, b, a, b)$. The first three voters approve all the decisions of the outcome, so each of them gets a satisfaction of 10 while the rest of the voters approve only half of the decisions from D , so they each get a satisfaction of 5. The coalition $S = \{4, 5, 6, 7, 8, 9\}$ agrees in all rounds on c , so Weak EJR demands that at least one voter in S should approve at least $\ell = 6$ decisions (since $6 = |S| \geq \ell \cdot \frac{n}{T} = 6 \cdot \frac{10}{10}$). However, each voter in S approves the decision of only 5 rounds, which is strictly less than $\ell = 6$, so Weak EJR is violated. Hence, Sequential Phragmén fails Weak EJR. \square

4.2 Semi-online: Method of Equal Shares (MES)

We do not know an online rule that satisfies Weak EJR. In fact, in Appendix B.2, we give a reduction showing that such a rule could be converted to a multi-winner voting rule satisfying Weak EJR and the axiom of *committee monotonicity*, the existence of which is a major open problem (Lackner and Skowron 2023; Sánchez-Fernández and Fisteus 2019).⁷

Can we evade this difficulty with some foresight, by relaxing the online requirement? Knowing the time horizon T is a common assumption found in online learning settings like multi-armed bandits (Barman et al. 2023). Indeed, if we know the total number of rounds T , we can use MES (which is online except that it uses T to determine the price $p = n/T$ of deciding a round). We show that it satisfies Weak EJR.

THEOREM 4.3. *MES satisfies Weak EJR.*

PROOF. Write $p = \frac{n}{T}$ for the cost of deciding a round according to MES. Suppose $S \subseteq N$ witnesses a violation of Weak EJR, where $|S| \geq \ell \cdot p$ but every $i \in S$ approves the decision of at most $\ell - 1$ rounds. We begin by showing that some $i \in S$ has a remaining budget of less than $p/|S|$ when MES terminates.

Suppose MES stops prematurely without making a decision in round $j \in R$. At that point, some voter $i \in S$ must have a remaining budget of strictly less than $p/|S|$, as otherwise S has a combined remaining budget of at least p and could therefore purchase an alternative that S agrees on in round j , contradicting that MES terminated. Otherwise, if MES does not stop prematurely, then it made a decision for all T rounds, in which case all available money is spent (since $T \cdot p = n$), and so all agents have remaining budget 0, hence in particular some $i \in S$ has remaining budget of less than $p/|S|$.

We first handle the case $\ell = 1$. Note that i 's remaining budget is less than $p/|S| \leq p/(\ell \cdot p) = 1/\ell = 1$. Thus, i has spent a strictly positive amount from the initial budget of $b_i = 1$. Since agents only spend money on approved alternatives, i 's utility must be at least 1 and we are done.

Thus, suppose $\ell \geq 2$ (so that we can divide by the quantity $\ell - 1 > 0$). Now, i has spent more than $1 - p/|S|$ by the time that MES has terminated, and has used that money to pay in at most $\ell - 1$ rounds. In those rounds, i therefore paid on average strictly more than

$$\frac{1 - p/|S|}{\ell - 1} \geq \frac{1 - 1/\ell}{\ell - 1} = \frac{1}{\ell}.$$

Hence, there must be a round when $i \in S$ paid strictly more than $1/\ell$ for the decision.

Let $j \in R$ be the *first* round where *some* voter in S paid strictly more than $1/\ell$ for the decision. Suppose this decision is $d_j \in C_j$. Just before paying for d_j , every voter in S had at least $1/\ell$ budget left as they each so far paid

⁷Elkind, Obraztsova, et al. (2025, Proposition 7.4) recently proved that no (semi-)online rule can satisfy EJR, answering a question we asked in earlier versions of the paper.

Rounds	1–10	11–16
Voter 1	{a}	{c}
Voter 2	{a}	{d}
Voter 3	{a}	{e}
Voters 4, 5, 6, 7, 8	{b}	{b}

Fig. 5. Instance where MES does not terminate prematurely yet fails EJR.

at most $1/\ell$ for at most $\ell - 1$ decisions. Recall that S agrees in every round; hence there exists an alternative $c \in C_j$ on which S agrees. But note that c is $1/\ell$ -affordable because

$$|S| \cdot \frac{1}{\ell} \geq \ell \cdot p \cdot \frac{1}{\ell} = p$$

while d_j is not $1/\ell$ -affordable (since some member of S had to pay strictly more than $1/\ell$ for d_j). This is a contradiction to MES choosing the alternative that is ρ -affordable for the lowest ρ . \square

Does MES provide good guarantees for coalitions that do not agree on all rounds? Unfortunately not. We show that MES fails PJR and EJR. This is perhaps surprising since, in other settings, MES usually satisfies at least as many proportionality axioms as Sequential Phragmén. The reason for its failure here is that coalitions may agree only in early rounds where MES greedily maximizes efficiency, and then MES does not satisfy the fairness requirements in subsequent rounds where there may not be enough agreement between voters to support the purchase of any alternative.

We begin by showing that MES does not guarantee EJR.

THEOREM 4.4. *MES fails EJR. It does so even on instances where it does not terminate prematurely, i.e., where it outputs a decision for all rounds.*

PROOF. Consider the instance in Figure 5. We have $T = 16$ rounds and $n = 8$ voters. The group $S_1 = \{1, 2, 3\}$ agrees in the first 10 rounds (jointly approving a) while the group $S_2 = \{4, 5, 6, 7, 8\}$ agrees in all rounds (approving b). The budget of each voter i is $b_i = 1$ unit while the price of each round is $p = \frac{n}{T} = \frac{8}{16} = 0.5$ units.

In round 1, a is ρ -affordable for $\rho = \frac{1}{6}$ whereas b is ρ -affordable for $\rho = \frac{1}{10}$. Thus, b is selected in round 1. The remaining budget of voters in S_2 is each set to $b_i - \rho = 1 - \frac{1}{10} = \frac{9}{10}$. By the same calculation, it can be seen that b is selected in every round until round 10, after which every voter in S_2 has exhausted their budget.

Beginning in round 11, only voters in S_1 have budget remaining. However, the voters 1, 2 and 3 do not agree in these remaining rounds. Further, every alternative apart from b is approved by only one voter, so its supporter will have to pay the full cost of $p = 0.5$ if it is selected. Thus, across the rounds 11–16, each of the 3 voters in S_1 will get to select the outcome of 2 of the rounds (with the exact sequence depending on how ties are broken).

Thus, no matter how ties are broken, the final decision sequence D will have b for the first 10 rounds and c, d, e each appearing twice in the last 6 rounds (in some order). Now notice that S_1 agrees in the first $k = 10$ rounds, so with $\ell = 3$, since $|S_1| \geq \ell \cdot \frac{n}{k} = 3 \cdot \frac{8}{10}$, EJR requires that some voter in S_1 approves decisions of at least $\ell = 3$ rounds. However, each voter in S_1 approves exactly 2 decisions in D , violating EJR. \square

We also have a similar example where MES fails PJR. However, we note that this example (shown in Figure 6) depends on empty approval sets and causes MES to terminate prematurely. We discuss these two caveats further below.

Round	1 & 2 & 3	4 & 5 & 6
Voter 1	{ <i>a</i> }	∅
Voter 2	{ <i>b</i> }	{ <i>b</i> }
Voter 3	{ <i>b</i> }	{ <i>b</i> }

Fig. 6. Instance where MES fails PJR (and JR). Note the empty approval set of voter 1 in rounds 4, 5, 6.

THEOREM 4.5. *MES fails PJR.*

PROOF. Consider the instance shown in Figure 6, with $T = 6$ rounds and $n = 3$ voters. Here, voter 1 approves *a* in the first 3 rounds but does not approve any alternative in the remaining rounds. Voters 2 and 3 approve *b* in all rounds. The budget of each voter i is $b_i = 1$ unit while the price of each round is $p = \frac{n}{T} = \frac{3}{6} = 0.5$ units.

In round 1, *a* is ρ -affordable for $\rho = 0.5$ as only voter 1 approves *a* and thus has to bear its whole cost of $p = 0.5$. However, *b* is ρ -affordable for $\rho = 0.25$ as both voters 2 and 3 approve it, so they can equally divide the cost p . So *b* is selected, and the remaining budget of voters 2 and 3 is set to 0.75. By the same computation, *b* is selected in rounds 2 and 3, after which voters 2 and 3 have a remaining budget of 0.25. In round 4, the only alternative to elect is *b* and voters 2 and 3 have just enough budget to buy it at $\rho = 0.25$, leading to exhaustion of their whole budget. In round 5, voters 2 and 3 do not have any budget to buy any alternative while voter 1 has the budget but does not approve any alternative. Thus, MES terminates prematurely. Any completion rule used from this point can only elect *b* in the remaining rounds as that is the only available alternative. Thus, we obtain the decision sequence $D = (b, b, b, b, b, b)$.

Now, the first voter forms a group $S = \{1\}$ which agrees in the first 3 rounds (as voter 1 approves *a*). With $\ell = 1$, it has size $|S| = 1 \geq \ell \cdot \frac{n}{3} = 1 \cdot \frac{3}{3}$. Thus, PJR requires that in at least $\ell = 1$ rounds, the outcome is *a*. However, MES (with any completion rule) never elects *a*, so MES fails PJR. \square

Remark 4.6. The counterexample of Theorem 4.5 used a requirement of $\ell = 1$ rounds, so it also shows that MES fails JR.

This also implies that MES fails EJ. Elkind, Obraztsova, et al. (2025, Proposition 7.4) show that in fact every semi-online method fails EJ.

As we mentioned, in the example of Figure 6, MES terminates prematurely. Thus, a completion rule would need to be used to make decisions in the remaining rounds. Due to the empty approval sets of Voter 1 in rounds 4, 5, and 6, however, it is clear that no matter how the remaining rounds are decided, the resulting decision sequence will fail PJR.

What about if we insist that all approval sets must be non-empty? In that case we are able to construct examples similar to Figure 6 where MES terminates prematurely (see Figure 7), and all the standard completion methods lead to a decision sequence that fails PJR. However, we are not aware of an example where PJR is failed for *all possible* completions of the MES output.

It turns out that the premature termination is, in a sense, the cause for the PJR failure: We show that if MES does not terminate prematurely (i.e., it outputs decisions for all the T rounds), its output decision sequence D satisfies PJR.

THEOREM 4.7. *If MES outputs a decision sequence D which contains a decision for all rounds (without terminating prematurely), then for all $S \subseteq N$ and $\ell \in \mathbb{N}$ with $|S| \geq \ell \cdot \frac{n}{T}$, there are at least ℓ rounds in which at least one member of S approves the decision of D (no matter in how many rounds S agrees). In particular, D satisfies PJR.*

Round	1–7	8–10
Voters 1–70	{ <i>a</i> }	{ <i>c</i> }
Voters 71–100	{ <i>b</i> }	all disjoint

Fig. 7. Example illustrating how MES fails PJR (and JR) without empty approval sets. The cell “all disjoint” indicates that each of the last 30 voters approves a unique alternative not approved by anyone else in the last 3 rounds. On this example, MES selects *a* in the first 7 rounds and then terminates prematurely. Most completion methods select *c* in the remaining rounds (including Phragmén completion, utilitarian completion, and Add1).

PROOF. Let $\ell \in \mathbb{N}$ and let $S \subseteq N$ be a group of voters with $|S| \geq \ell \cdot \frac{n}{T}$. Recall that MES initially gives each voter a budget of 1 while each round’s decision has a price of $p = \frac{n}{T}$. As MES produced a decision for all rounds, in total $p \cdot T = n$ amount was spent. This implies that every voter’s budget of 1 is used in full. Thus, the voters in S paid a total amount of $|S|$ during the execution of MES. Recall that voters only pay when they approve a round’s decision. Thus, the number of rounds in which the decision is approved by at least 1 voter in S must be at least $|S|/p = |S| \cdot \frac{T}{n} \geq \ell$ rounds, as desired. \square

4.3 Offline: Proportional Approval Voting (PAV)

In some settings, offline voting (where alternatives and approval sets for all rounds are known in advance) is possible, e.g., for voting in combinatorial domains with independent issues. Studying the offline setting can also clarify which axioms are plausible aims for online rules. It turns out that if we make decisions fully offline, there is a rule that satisfies all four of our axioms: PAV, as well as its polynomial-time local search variant. The proof uses a swapping argument: if the output violates EJR, then in at least 1 round, one can change the decision and thereby increase the PAV objective function. This technique is also used in multi-winner voting (Aziz, Brill, et al. 2017). Our theorem was recently generalized by Masařík et al. (2024) to show that PAV satisfies an EJR notion in a general model of voting with matroid constraints (noting that our model can be described as involving constraints that in each round, just a single alternative can be chosen).

THEOREM 4.8. *PAV and Local-Search PAV satisfy EJR.*

PROOF. If D_{PAV} is the decision outcome chosen by PAV, and we run Local-Search PAV with $D_{\text{init}} = D_{\text{PAV}}$, then it immediately terminates. Hence D_{PAV} is a possible output of Local-Search PAV. Thus, to prove the theorem, it is sufficient to prove that every possible output of Local-Search PAV satisfies EJR.

So let D be a decision outcome returned by Local-Search PAV with $|D| = T$. For a contradiction, suppose that $S \subseteq N$ witnesses a violation of EJR of D , with S agreeing on a set $R^* \subseteq R$ of $|R^*| = k$ rounds and with $|S| = s \geq \ell \cdot \frac{n}{k}$ but $U_D^i < \ell$ for all $i \in S$. For any round $r \in R^*$, let $a_r \in C_r$ be an alternative on which S agrees and d_r be the current decision for that round. If we replace d_r by a_r in D , the change in the PAV-score of D is

$$\begin{aligned}
\Delta(a_r, d_r) &= \underbrace{\sum_{\substack{i \in S \\ d_r \notin A_r^i}} \frac{1}{U_D^i + 1} + \sum_{\substack{i \in N \setminus S \\ d_r \notin A_r^i \\ a_r \in A_r^i}} \frac{1}{U_D^i + 1}}_{\text{increase due to addition of } a_r} - \underbrace{\sum_{\substack{i \in N \setminus S \\ d_r \in A_r^i \\ a_r \notin A_r^i}} \frac{1}{U_D^i}}_{\text{due to removal of } d_r} \\
&\geq \sum_{\substack{i \in S \\ d_r \notin A_r^i}} \frac{1}{U_D^i + 1} - \sum_{\substack{i \in N \setminus S \\ d_r \in A_r^i}} \frac{1}{U_D^i} \tag{1}
\end{aligned}$$

Note that the change for a particular round may be negative. Summing this change over R^* , the set of rounds where S agrees, we get

$$\begin{aligned} \sum_{r \in R^*} \Delta(a_r, d_r) &\geq \sum_{r \in R^*} \left(\sum_{\substack{i \in S \\ d_r \notin A_r^i}} \frac{1}{U_D^i + 1} - \sum_{\substack{i \in N \setminus S \\ d_r \in A_r^i}} \frac{1}{U_D^i} \right) && \text{(from (1))} \\ &= \sum_{i \in S} \sum_{\substack{r \in R^* \\ d_r \notin A_r^i}} \frac{1}{U_D^i + 1} - \sum_{i \in N \setminus S} \sum_{\substack{r \in R^* \\ d_r \in A_r^i}} \frac{1}{U_D^i} && \text{(interchanging sums)} \end{aligned}$$

Write $U_{R^*}^i$ for the number of rounds in R^* in which i approves the decision of D . Note that for all $i \in S$, we have $U_{R^*}^i \leq U_D^i \leq \ell - 1$, and so the number of decisions in rounds R^* not approved by i is at least $k - (\ell - 1) = k - \ell + 1$. Hence, we get

$$\begin{aligned} \sum_{r \in R^*} \Delta(a_r, d_r) &\geq \sum_{i \in S} \frac{k - \ell + 1}{U_D^i + 1} - \sum_{i \in N \setminus S} \sum_{\substack{r \in R^* \\ d_r \in A_r^i}} \frac{1}{U_D^i} \\ &\geq \sum_{i \in S} \frac{k - \ell + 1}{U_D^i + 1} - \sum_{i \in N \setminus S} \sum_{\substack{r \in R^* \\ d_r \in A_r^i}} \frac{1}{U_{R^*}^i} && \text{(since } U_{R^*}^i \leq U_D^i \text{ for all } i \in N) \\ &\geq \sum_{i \in S} \frac{k - \ell + 1}{U_D^i + 1} - \sum_{i \in N \setminus S} 1 && \text{(definition of } U_{R^*}^i) \\ &\geq \sum_{i \in S} \frac{k - \ell + 1}{\ell - 1 + 1} - (n - s) && \text{(since } U_D^i \leq \ell - 1 \text{ for all } i \in S) \\ &= \left(\frac{k - \ell + 1}{\ell} \right) \cdot s - (n - s) \\ &= \frac{sk}{\ell} - s + \frac{s}{\ell} - n + s \\ &= s \frac{k}{\ell} + \frac{s}{\ell} - n \geq n + \frac{s}{\ell} - n \geq \frac{n}{k} && \text{(since } s \geq \ell \cdot \frac{n}{k}) \end{aligned}$$

Thus, for an average $r \in R^*$, we have $\Delta(a_r, d_r) \geq \frac{1}{k} \cdot \frac{n}{k} = \frac{n}{k^2} \geq \frac{n}{T^2}$ using that $|R^*| = k \leq T$. Therefore, there exists a round $r \in R^*$ such that $\Delta(a_r, d_r) \geq \frac{n}{T^2}$. Hence D admits a swap that increases its PAV-score by at least $\frac{n}{T^2}$, and thus Local-Search PAV would not terminate and return D . This gives a contradiction and completes the proof. \square

5 Impossibility of Stronger Guarantees

One could try to strengthen the axioms we have discussed in various ways, by attempting to either increase the amount of utility guarantees, or to relax the constraints on group size. However, one can show that many such strengthenings lead to axioms that may be infeasible. For example, if we reduce the requirement on the coalition size by ϵ in the definition of PJR, there are instances where no decision sequence satisfies the strengthening.

THEOREM 5.1. *Let $\epsilon > 0$. Then there exists an instance where no decision sequence D satisfies “ ϵ -Stronger PJR”, defined to require that for every $\ell \in \mathbb{N}$ and every group of voters $S \subseteq N$ that agrees in k rounds and has size*

$|S| \geq (\ell - \epsilon) \cdot \frac{n}{k}$, there are at least ℓ rounds $j \in R$ in which the decision d_j of D is approved by at least one voter in S (i.e., $d_j \in \bigcup_{i \in S} A_j^i$).

The proof of [Theorem 5.1](#) is in [Appendix B.1](#). It constructs a counterexample where in the first k rounds, there are very many coalitions that agree (but cannot all be satisfied simultaneously), while in the remaining rounds, there is no agreement at all (which makes it impossible to satisfy all the justified demands from earlier rounds). In [Appendix B.1](#), we also discuss several other potential strengthenings of our axioms, including what we call “ ℓ -Agreement PJR”, which requires a group of size at least $\ell \cdot n/T$ to only agree in ℓ rounds to deserve to be represented in ℓ -rounds, which is impossible to satisfy due to the same counterexample construction.

[Theorem 5.1](#) is a worst case result, and we could hope for rules that satisfy ϵ -Stronger PJR on inputs where it is possible. However, we show that no (semi-)online rule can do this.

THEOREM 5.2. *Let $\epsilon > 0$. No semi-online decision rule returns an ϵ -Stronger PJR decision sequence whenever one exists.*

Again, the proof can be found in [Appendix B.1](#).

6 Experiments

To understand the performance of our methods empirically, we run our methods on both synthetic and real-world datasets.

Rules. In addition to our proposed rules (Phragmén, MES, PAV), we also consider two online rules proposed by [Lackner \(2020\)](#):

- *Perpetual Quota* is a rule that aims at granting each voter a satisfaction as close as possible to their “quota”. Formally, in round $j \in R$, we compute a weight w_i for each voter $i \in N$, defined as

$$w_i = \max \left(0, \frac{1}{n} \sum_{t=1}^j \max_{c \in A_t^i} |\{i' \in N : c \in A_t^{i'}\}| - U_{j-1}^i \right),$$

where U_{j-1}^i is the number of decisions approved by i in rounds $1, \dots, j-1$ and the max is interpreted as equal to 0 when $A_t^i = \emptyset$. Perpetual Quota chooses an alternative $c \in C_j$ that maximizes the weighted approval score $\sum_{i \in N : c \in A_j^i} w_i$. We refer the reader to [Lackner \(2020\)](#) for the motivation of this rule; intuitively, voters receive higher weights if they are easier to satisfy because they approve more popular alternatives, and they receive lower weights if they have already been satisfied in prior rounds.

- *Perpetual Consensus* is similar to Sequential Phragmén but strictly enforces an equal distribution of the load incurred. Formally, it maintains a weight w_i for each voter $i \in N$, initialized to $w_i = 1$. These weights are allowed to become negative, but in each round the rule only considers voters who currently have positive weight, $N^+ = \{i \in N : w_i > 0\}$. It then selects an alternative $c \in C_j$ that maximizes the weighted approval score among these voters, i.e., maximizes $\sum_{i \in N^+ : c \in A_j^i} w_i$. Let $S^+ = \{i \in N^+ : c \in A_j^i\}$ be the set of voters who contributed to the selected alternative’s score. The rule then updates the weights by setting $w_i \leftarrow w_i + 1 - n/|S^+|$ for all $i \in S^+$ and $w_i \leftarrow w_i + 1$ for other voters.

Both of these rules intuitively aim for proportionality, but do not satisfy any of our axioms as they even fail the weaker axioms proposed by [Lackner and Maly \(2023\)](#). We chose those rules as they performed well in the experiments of [Lackner \(2020\)](#).

Further, we consider two baselines:

- *Approval Voting* which in each round $j \in R$ chooses an alternative with the highest approval score, i.e., an alternative $c \in C_j$ that maximizes $|\{i \in N : c \in A_j^i\}|$.

- *Round Robin* which in each round $j \in R$, identifies the voter $i = j \bmod n$ and chooses an alternative $c \in A_j^i$ approved by that voter (the first one according to a fixed tie-breaking order).

We implemented all the voting rules in Python 3.8, building on existing implementations: For Sequential Phragmén (a.k.a. Perpetual Phragmén), Perpetual Quota, Perpetual Consensus, Approval Voting, and Round Robin (a.k.a. Serial Dictatorship), we used the `perpetual` python package. For PAV, we used a standard integer linear-program (ILP) encoding (see, e.g., [D. Peters and Lackner 2020](#)), solved using Gurobi 9. For MES, our implementation builds upon the `Grzesiek3713/pabutools` package and uses Phragmén completion in case of premature termination.

Metrics. We evaluate our rules on several metrics of voter utility to complement our theoretical guarantees. For comparability of results, we normalize utility and define a single voter i 's utility as the *fraction* of rounds in which the voter approves the decision: U_D^i/T . Based on this, we report three metrics:

- *Average Utility* of the voters (equivalent to utilitarian social welfare).
- *Utility of the 25th Percentile*: We sort the vector of utilities and report its 25th percentile. This is inspired by egalitarian social welfare, which we did not use because the minimum utility was often zero.
- *Gini Coefficient*: This metric quantifies the level of inequality in the voter utilities. A lower value corresponds to a more equal utility distribution (with 0 being obtained in case every voter has the same utility). Formally, the Gini coefficient of a decision sequence D is

$$\text{gini}(D) = \frac{1}{\sum_{i \in N} U_D^i} \sum_{i \in N} \sum_{j \in N} \frac{1}{2n} |U_D^i - U_D^j|.$$

6.1 Synthetic Data

For analysis based on synthetic data, we follow a similar setup to the one used by [Lackner \(2020\)](#) which is based on the popular approach of sampling both voters and alternatives as points in a two-dimensional Euclidean space ([Elkind, Faliszewski, Laslier, et al. 2017](#)). Because we are particularly interested in group representation, we use distributions of voter locations that have several clusters.

More precisely, both voters and alternatives are sampled as points in \mathbb{R}^2 . In each round, a fresh set of alternatives is sampled uniformly in the square $[-1, 1] \times [-1, 1]$. The voter locations stay fixed across rounds, and for each voter $i \in N$, we sample i 's location independently of other voters from a bivariate normal distribution $\mathcal{N}(x_i^*, y_i^*)$ centered at a point (x_i^*, y_i^*) with standard deviation $\sigma = 0.2$ (unless stated otherwise). We present results for four distributions of voter locations, which differ in their choices of points (x_i^*, y_i^*) and the size of the groups:

- **RESTRICTED**: The voters are split into 2 disjoint groups S_1 and S_2 with $\frac{1}{3}$ of voters in S_1 and $\frac{2}{3}$ of voters in S_2 . The locations of voters in S_1 are sampled from $\mathcal{N}(-0.5, -0.5)$, while the locations of voters in S_2 are sampled from $\mathcal{N}(0.5, 0.5)$. We further *restrict* the voter locations to lie in the square $[-1, -1] \times [-1, 1]$, if necessary resampling from the location distributions until this constraint is satisfied.
- **MANY GROUPS**: The voters are split into 4 disjoint groups concentrated in different areas. The first 3 groups each form 20% of the population while the last group forms 40%. For each group, both x - and y -coordinates are drawn independently from $\mathcal{N}(\pm 0.5, \pm 0.5)$.
- **UNBALANCED**: The voters are split into 2 disjoint groups S_1 and S_2 with 20% of voters in S_1 and 80% of voters in S_2 . The locations of voters in S_1 are sampled from $\mathcal{N}(-0.5, -0.5)$, and the locations of voters in S_2 are sampled from $\mathcal{N}(0.5, 0.5)$. For this distribution, we used standard deviation $\sigma = 0.1$.
- **NEAR-BALANCED AND NEARBY**: The voters are split into 2 disjoint groups S_1 and S_2 with 60% of voters in S_1 and 40% of voters in S_2 , with the groups being quite close to each other in comparison to other distributions. All the voters of S_1 were sampled from $\mathcal{N}(-0.25, 0)$ while the voters of S_2 were sampled from $\mathcal{N}(0.25, 0)$.

These distributions are visualized in the blue pictures in the right-most column of [Figure 8](#), where we sampled 2 500 points from these distributions (chosen for visualization purposes; the experiments use fewer voters).

Our experiments are parameterized by a factor $f \geq 1$, which determines how many alternatives are approved by voters. Specifically, voters approve all alternatives whose Euclidean distance is within f times the distance to their closest alternative. In particular, each voter approves at least one alternative. For each of the distributions mentioned above, we consider the following three parameter settings:

- $n = 20$ voters, $T = 50$ rounds and each round having 20 alternatives, with $f = 1.5$. This led to an average approval set size of 2.12.
- $n = 50$ voters, $T = 50$ rounds and each round having 40 alternatives, with $f = 1.5$. This led to an average approval set size of 2.2.
- $n = 50$ voters, $T = 50$ rounds and each round having 40 alternatives, with $f = 3$. This led to a larger approval set size with the average being 7.52.

We repeated each experiment for 1000 trials. We report our results for the various distributions in [Figures 8 to 10](#). Each figure corresponds to a specific distribution and contains multiple subfigures to show the plots for different parameter values. In each subfigure, bar plots are shown for each rule and metric. The length of the bar represents the median value across all the 1000 trials for the rule while the error bars represent the 25th and 75th percentile, with the numeric text after the bar showing the mean value.

Considering each metric, we make the following observations:

- *Average Satisfaction*: In regard to this metric, a clear separation emerges between the voting methods across all the distributions. Approval Voting gives the highest average satisfaction, which is expected (it is easy to see that it always selects the decision sequence maximizing this metric). On the other hand, Round Robin performs the worst among the rules for all parameter settings and all distributions. Among the rules aiming for proportional outcomes, we see the following rank order for almost all parameter settings and distributions: PAV > Sequential Phragmén > MES > Perpetual Quota > Perpetual Consensus. The differences between adjacent rules in this ordering tend to be small. We see some exceptions to this ordering for the UNBALANCED distribution, where Perpetual Quota does better than usual, but is still worse than PAV. The advantage of PAV over other rules is particularly apparent in the MANY GROUPS distributions.
- *Utility of the 25th Percentile*: Here, we find that Approval Voting consistently performed poorly, often resulting in almost no decisions being approved by the bottom 25% of the voters, leading to a score of close to 0 for this metric under the RESTRICTED and the MANY GROUPS distributions. Round Robin performed better than Approval Voting but still much worse than the other rules. Among the fair rules, PAV generally performed the best. However, in the NEAR-BALANCED AND NEARBY distribution, Perpetual Quota outperforms PAV for several parameter settings. Perpetual Consensus performs similarly well for this distribution, but Sequential Phragmén and MES perform notably worse. For the other three distributions, the proportional rules all perform similarly on this metric, with Sequential Phragmén showing the best values for several parameter settings.
- *Gini Coefficient*: Across distributions, Approval Voting consistently produces the outcomes with the most inequality between voters, as measured by the Gini coefficient. That means that some voters are happy with many decisions, and others are happy with few decisions. All other rules produce outcomes with a more equal distribution of utility. Among the proportional rules, there is no consistent ranking between rules in their performance on this metric, and their Gini coefficients are often similar. Under the NEAR-BALANCED AND NEARBY as well as the MANY GROUPS distributions, however, Perpetual Consensus and Perpetual Quota have noticeably lower Gini coefficient. Depending on the parameters and distribution, the Round Robin method performs either slightly worse or comparably to the more sophisticated proportional rules – recalling that, however, the average utility under this method is lower.

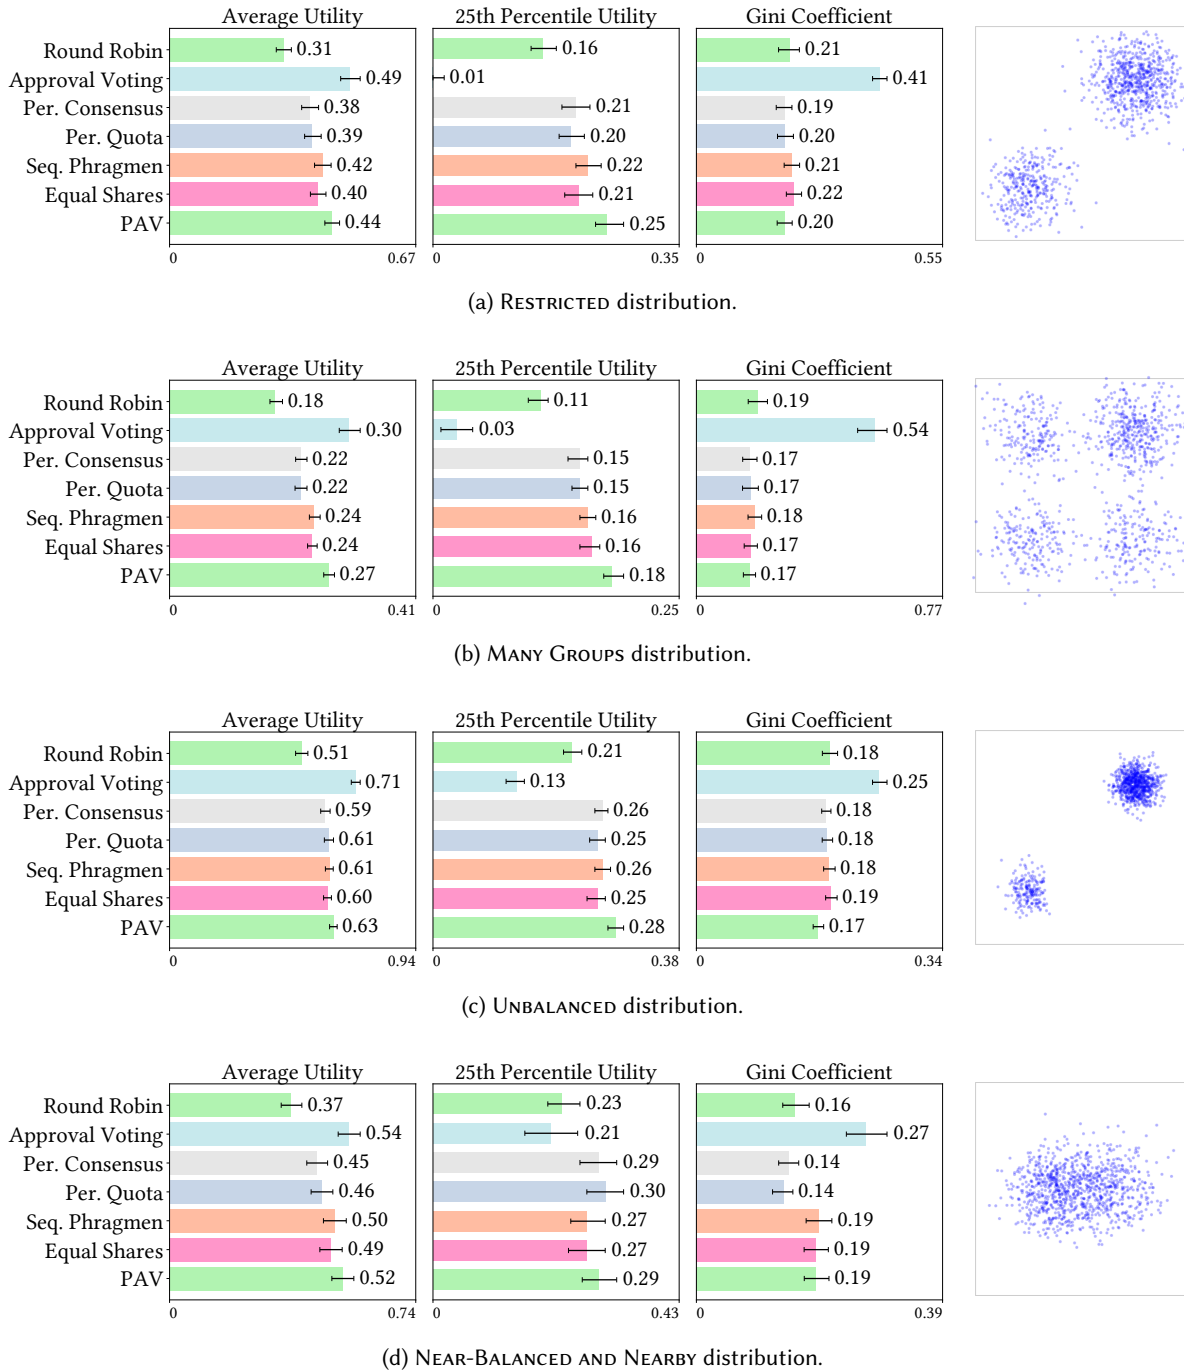


Fig. 8. Performance of the rules for $n = 20$ voters, $T = 50$ rounds, 20 alternatives in each round, with $f = 1.5$. The length of the bar represents the median across trials while the error bars represent the 25th and 75th percentile with the numeric text after the bar showing the mean value. The blue pictures in the right-most column illustrate the distribution of voter locations.

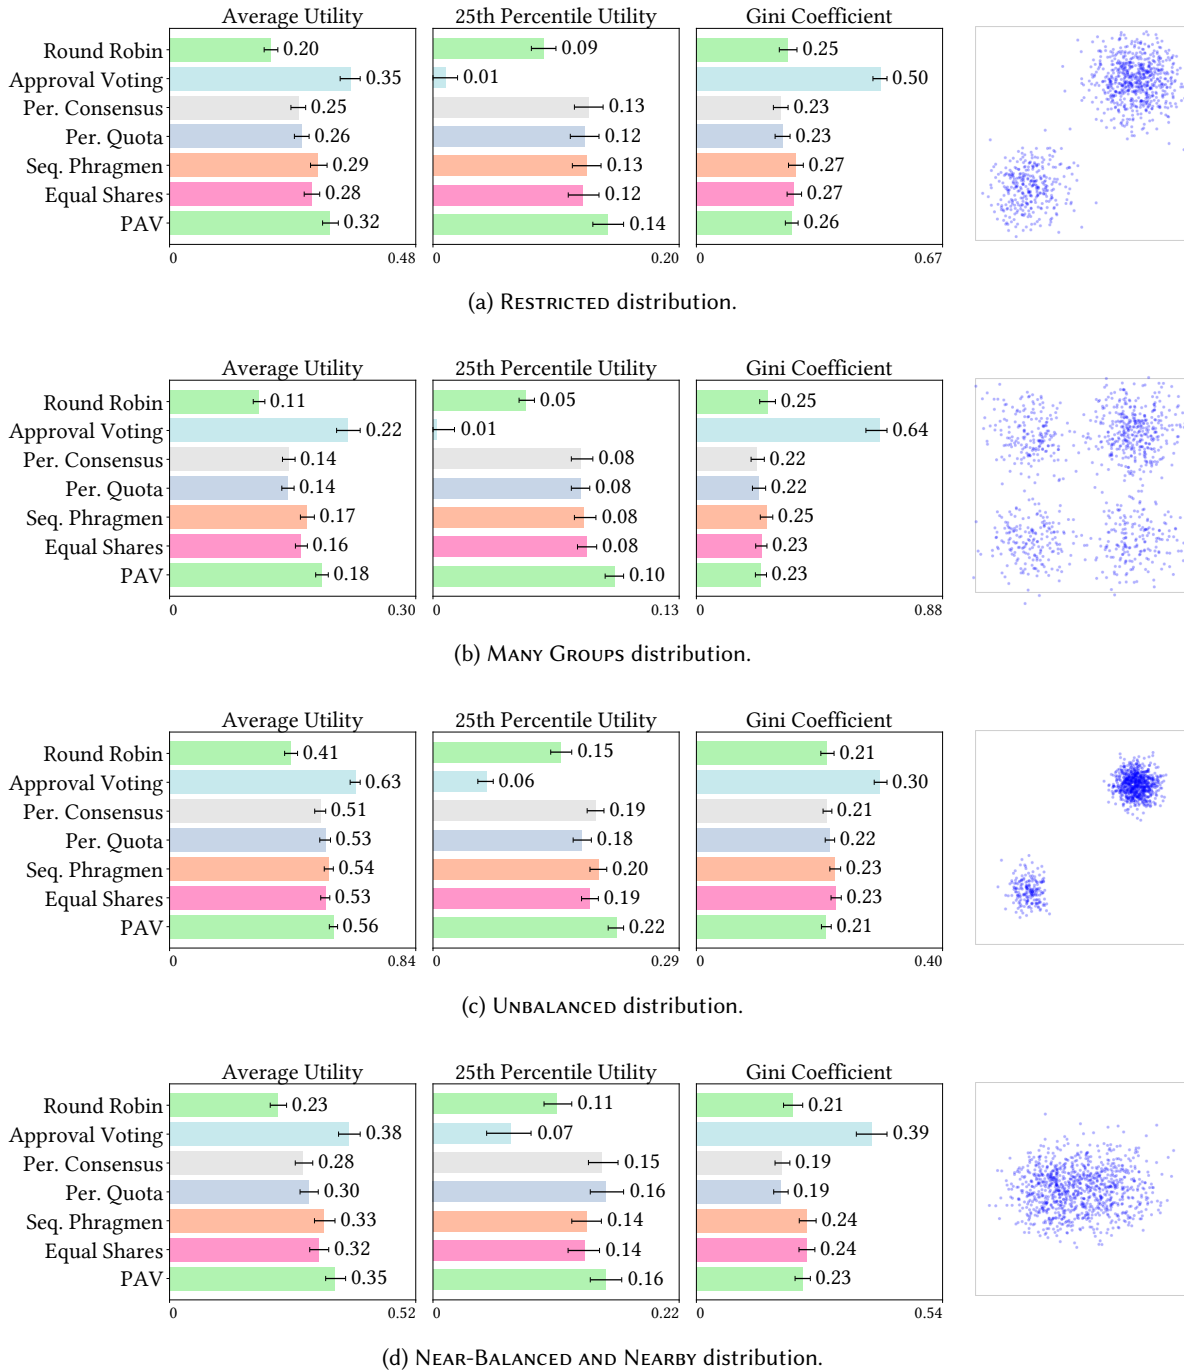




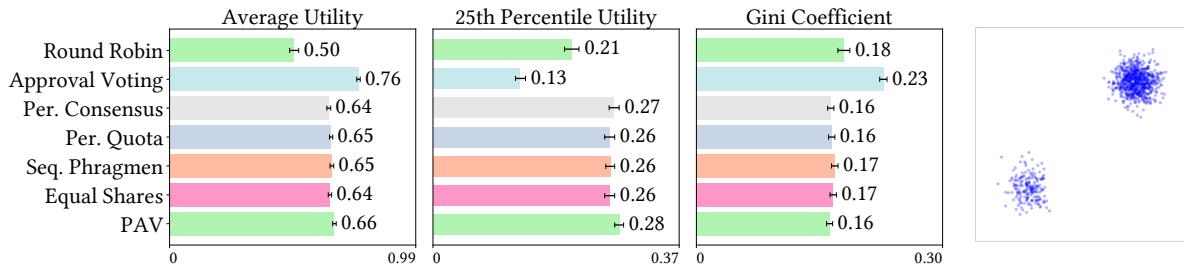
Fig. 9. Performance of the rules for $n = 50$ voters, $T = 50$ rounds, 40 alternatives in each round, with $f = 1.5$. The length of the bar represents the median across trials while the error bars represent the 25th and 75th percentile with the numeric text after the bar showing the mean value. The blue pictures in the right-most column illustrate the distribution of voter locations.




(a) RESTRICTED distribution.



(b) MANY GROUPS distribution.



(c) UNBALANCED distribution.



(d) NEAR-BALANCED AND NEARBY distribution.

Fig. 10. Performance of the rules for $n = 50$ voters, $T = 50$ rounds, 40 alternatives in each round, with $f = 3$. The length of the bar represents the median across trials while the error bars represent the 25th and 75th percentile with the numeric text after the bar showing the mean value. The blue pictures in the right-most column illustrate the distribution of voter locations.

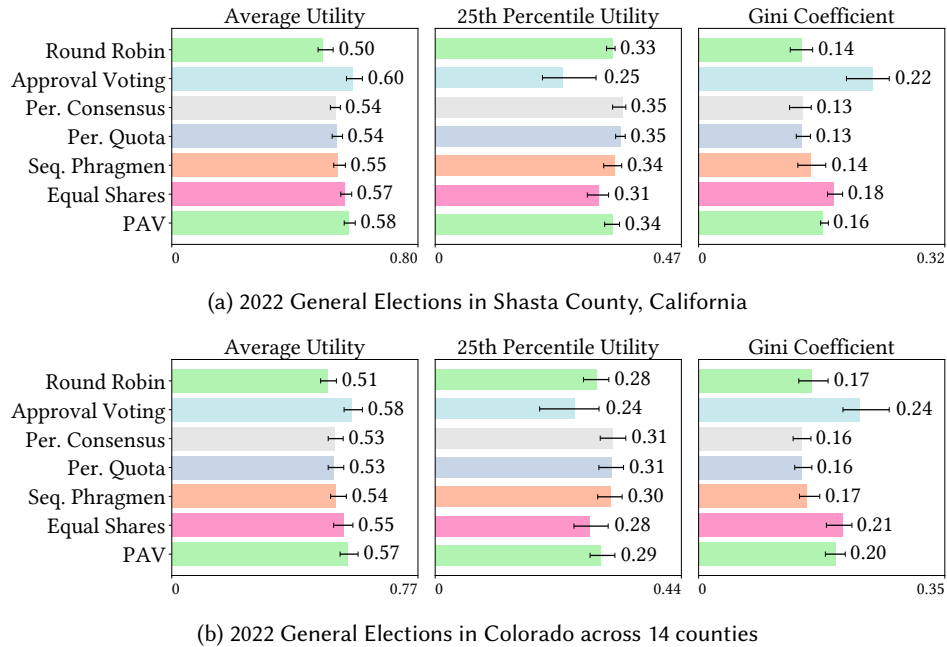


Fig. 11. Performance of the rules on the political dataset. The length of the bar shows the median across all voting districts, the error bars give the 25th and 75th percentile, and the text gives the mean.

6.2 Political Data

In addition to synthetic data, we evaluated the rules on data from U.S. political general elections which we collected from websites of local governments. The resulting dataset is available at <https://osf.io/t6p7s/>. In the elections we collected, voters elect candidates to various federal, state, and local political offices and express opinions on yes/no ballot initiatives. Using public anonymized *Cast Vote Records (CVR)* data we can see each voter's votes on all these issues simultaneously. We collected instances from 15 counties in California and Colorado from 2020 and 2022. In these elections, voters usually only vote for a single alternative for each issue (rather than using approval). However, where several members needed to be elected to a board, voters are usually allowed to vote for several candidates (but no more votes than there are open positions). We reinterpreted such issues to be about the election of a single candidate, and interpreted the votes as forming an approval set. We did this to make the dataset fit better with our formal setup. The voting rule actually used for these elections can be seen as equivalent to Approval Voting (i.e., in each round select the alternative with the highest vote count), one of the baseline rules included in our simulations.

Each county is divided into smaller voting districts, giving rise to a large number of decision instances. For example, in Colorado, we collected 169 instances from 14 counties. Each instance corresponds to all the voters that were issued a particular ballot type (usually, one instance per district of the county, since the candidates and offices differ between districts). We collected 41 instances from Shasta County, California, (one for each voting district) with each instance on average having 32 issues and 1 682 voters. We ignored instances with more than 10 000 voters because the computation of PAV did not finish in reasonable time.

The results are shown in Figure 11 separately for Shasta County, Calif., and the 14 Colorado counties in our dataset. We see that Approval Voting (AV), the method that is actually used to determine winners, gives a higher




Fig. 12. Screenshot of a decision problem from moralmachine.net.

Gini coefficient and lower utility at the 25th percentile compared to all the other rules. Strikingly, PAV leads to a more equal utility distribution while also having high average utility, with the difference to AV being almost negligible. This suggests that proportional rules could be used to produce collections of decisions that are more consensual. We note, however, that much additional analysis would be appropriate, to understand the precise behavior of the rules on political instances.

6.3 Learning Preferences from the Moral Machine

Virtual democracy (Freedman et al. 2020; Mohsin et al. 2021; Noothigattu et al. 2018) is a proposal to automate decision making by learning models of preferences of individual users and using predicted preferences as inputs to a voting rule. This can be particularly useful when preference elicitation costs are high, decisions need to be made in real time, or a very large number of decisions need to be taken. A common approach (Noothigattu et al. 2018; Wang et al. 2019) for virtual democracy is to *average* the learned model parameters across voters to obtain a single model, and effectively treat the predictions of the averaged model as the result of an aggregation method. Feffer et al. (2023) have shown that such averaging may lead to non-proportional outcomes that underweight minority preferences. We hope to alleviate these fairness concerns by querying each individual model and then aggregating their outputs using voting rules. In particular, we will use the proportional aggregation methods that we have studied, and study their effect as compared to non-proportional rules and as compared to the traditional Machine Learning approach of learning a single preference model on a combined dataset.

Following the work of Noothigattu et al. (2018), to empirically test our work on a dataset that has structured features which allow preference learning, we consider virtual democracy applied to the *moral machine* (Awad et al. 2018). This experiment is a modern take on ethical “trolley problems” and involves decisions that a self-driving car might face. Users were asked to express preferences in instances where a self-driving car must either swerve or stay in the lane, with both choices leading to injuring a different group of people. These choices can be seen as alternatives that can be described by a structured feature vector. Several million pairwise comparison responses are available in a public dataset. Note that we do not view ethics for self-driving cars as a promising application of our rules; we just use it as a convenient and evocative data set that allows comparison to previous work.

Noothigattu et al. (2018) learn a model predicting the preferences of each respondent. However, there are only 13 pairwise comparisons per respondent, meaning that such individual models have low accuracy. Instead, we partition the respondents by their country, and learn one model for each country. Kim et al. (2018) found that

grouping respondents from each country leads to higher accuracy on the moral machine dataset (perhaps due to cultural similarities).

We limit ourselves to the 197 countries for which the dataset contains over 100 samples. To learn a preference model for each country, we use the Plackett–Luce (PL) model (Luce 1959; Plackett 1975) which is a *random utility model* appropriate for social choice preference learning (Azari Soufiani et al. 2012). As a baseline, we train a *combined model* on respondents from all countries, using 100 samples from each country for a balanced representation in the training data.

In our experiments, we consider four different parameter settings (a)-(d). We produce 100 decision rounds. In each round, we produce either 50 alternatives ((a)-(b)), 100 alternatives ((c)) or 200 alternatives ((d)), in one of the following ways:


- *Random*. We create a large pool of 240 000 unseen alternatives and sample a random subset of these as the alternatives for each round.
- *High-disagreement alternatives*. We find that on random alternatives, the country models have high agreement (i.e., they tend to approve the same alternatives from the sampled set of alternatives), making the aggregation problem relatively straightforward. Thus, a more interesting case involves finding alternatives with high disagreement. We quantify disagreement on an alternative as high variance in the ranks assigned by each model to the alternative. From the pool of all 240 000 alternatives present in the dataset, we pick the 500 alternatives with the highest variance in ranks. Then in each round we sample alternatives randomly from this high-disagreement set of 500 alternatives.

Finally, we generate approval sets by computing the utilities for each alternative as predicted by the country model, and let the country approve (a) the alternatives with above-average utility, (b)/(d) the top 2 alternatives, (c) the top 1 alternative.


As before, we use our voting rules to compute decision outcomes using the country models as voters. To compare the voting rules to the performance of the combined model, we pick the alternative assigned the highest utility by the combined model as its decision for the round. We present the results for *random alternatives* in Figure 13 and for *high-disagreement alternatives* in Figure 14, in each case having run the experiments for 1000 trials.

For random alternatives (Figure 13), all rules perform very similarly except for Round Robin which performs much worse, presumably because it misses opportunities for compromise. Looking at the results for high-disagreement alternatives in Figure 14, we find it striking how Approval Voting (AV) and the Combined Model attain almost identical values on each metric across parameter settings, and how these are quite different from the values obtained by the proportional rules. For example, for the parameters of the experiment shown in Figure 14(c), AV and the Combined Model choose the same decision in 87% of rounds, while the Combined Model agrees with the five proportional in between only 61% (Perpetual Consensus) and 76% (PAV) of rounds. Notably, the five proportional rules all feature a much smaller Gini coefficient and a higher satisfaction at the 25th percentile. The similar performance between AV and the Combined Model suggests that the Combined Model exhibits a bias towards plurality and majority opinions. This contradicts a possible hope one might have had that, because the Combined Model is trained based on an equal number of samples from each country, it will “merge” their views roughly proportionally.

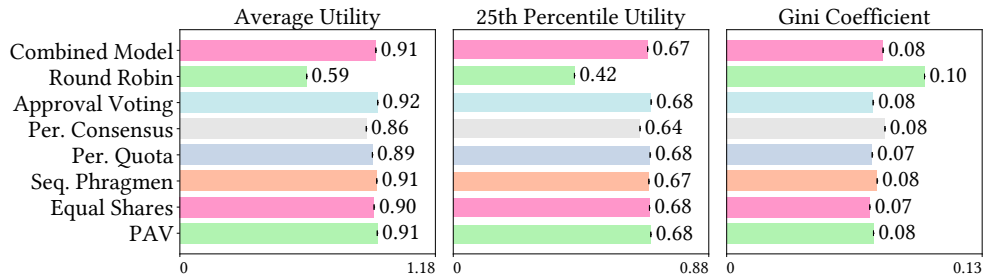
We see the results from this small experiment as a potential starting point for a larger research program that studies how a global preference model trained on preferences of diverse agents makes decisions compared to aggregation rules that explicitly take into account each individual’s preferences. Note that the results do not imply that such models can or should be used in applications, or to automate moral reasoning. To move in such a direction, we would need confidence that the models and methods take into account all relevant considerations and are rigorously tested, and that such automation is a good idea in the first place. We just




(a) 100 rounds with each round having 50 alternatives and each voter approving above-average alternatives.



(b) 100 rounds with each round having 50 alternatives and each voter approving top 2 alternatives.

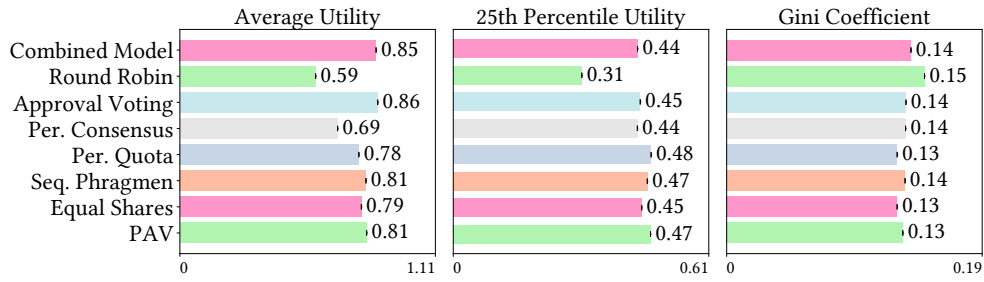


(c) 100 rounds with each round having 100 alternatives and each voter approving top 1 alternatives.




(d) 100 rounds with each round having 200 alternatives and each voter approving top 2 alternatives.


Fig. 13. Performance of the different rules on the Moral Machine Dataset based on *random* alternatives.



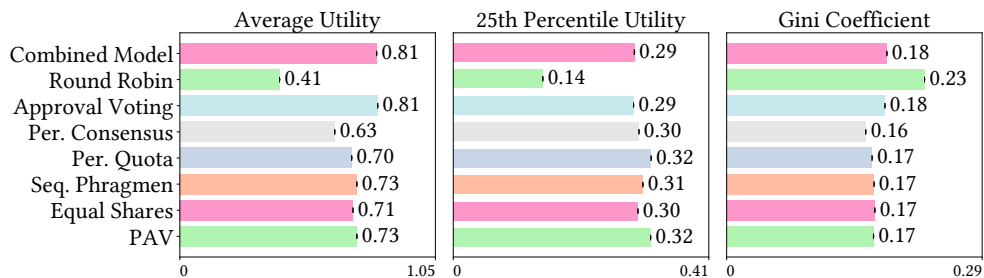
(a) 100 rounds with each round having 50 alternatives and each voter approving above-average alternatives.



(b) 100 rounds with each round having 50 alternatives and each voter approving top 2 alternatives.



(c) 100 rounds with each round having 100 alternatives and each voter approving top 1 alternatives.



(d) 100 rounds with each round having 200 alternatives and each voter approving top 2 alternatives.

Fig. 14. Performance of the different rules on the Moral Machine Dataset, based on *high-disagreement* alternatives.

see these experiments as suggesting that *if* such automation were to be put in place, then taking fairness and proportionality considerations into account would be important.

7 Discussion and Future Work

Extensions. Our model can be extended to make it compatible with more real-world applications. Examples include allowing weighting issues by importance (Page et al. 2020), to allow for dependencies between issues (Brill, Markakis, et al. 2023), or to allow voters to specify utilities or to rank alternatives instead of approvals (D. Peters, Pierczyński, et al. 2021). The latter may be important as some issues may be more critical to certain minority groups.

Open Problems. We leave some theoretical open problems for future work, notably whether an online rule can satisfy Weak EJR – a negative result may be easier to find in this setting than for multi-winner voting. Also open is whether MES can always be completed to satisfy PJR when all approval sets are non-empty. More conceptually, are there stronger versions of (Weak) EJR for this setting that are still satisfiable? Can the concept of proportionality degree (Skowron 2021) be adapted to our setting? What about FJR (Full Justified Representation), EJR+, or the core? Elkind, Obratsova, et al. (2025) recently defined an analog of the Greedy Cohesive Rule (D. Peters, Pierczyński, et al. 2021) for the sequential setting, which would be a natural starting point for studying FJR.

Strategic Issues. D. Peters (2018) proved an impossibility theorem showing that no proportional multi-winner voting rule can be *strategyproof*: voters may be able to get a better outcome by misrepresenting their preferences. For the special case of “approval-based apportionment” (Airiau et al. 2023; Brill, Gözl, et al. 2022), the impossibility still holds (Delemazure et al. 2023; Lackner, Maly, and Nardi 2023). As this is also a special case of our model (when the set of alternatives and voter preferences are the same in each round), it follows that no proportional rule in our setting can be strategyproof. Note that the presentation order of issues can also affect the results of online rules, offering another potential avenue for manipulation.

Acknowledgments

We thank Jérôme Lang for his support and advice throughout this research project, and for proposing key ideas and directions. We thank Jan Maly and Jannik Peters for helpful feedback. 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).

References

- S. Airiau, H. Aziz, I. Caragiannis, J. Kruger, J. Lang, and D. Peters. 2023. “Portioning using ordinal preferences: Fairness and efficiency.” *Artificial Intelligence*, 314, 103809. doi:10.1016/j.artint.2022.103809.
- G. Amanatidis, N. Barrot, J. Lang, E. Markakis, and B. Ries. 2015. “Multiple referenda and multiwinner elections using Hamming distances: Complexity and manipulability.” In: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 715–723. <https://www.ifaamas.org/Proceedings/aamas2015/aamas/p715.pdf>.
- E. Awad, S. Dsouza, R. Kim, J. Schulz, J. Henrich, A. Shariff, J.-F. Bonnefon, and I. Rahwan. 2018. “The moral machine experiment.” *Nature*, 563, 7729, 59–64. doi:10.1038/s41586-018-0637-6.
- H. Azari Soufiani, D. C. Parkes, and L. Xia. 2012. “Random utility theory for social choice.” *Advances in Neural Information Processing Systems*, 25, 126–134. https://papers.nips.cc/paper_files/paper/2012/hash/a512294422de868f8474d22344636f16-Abstract.html.
- H. Aziz, A. Bogomolnaia, and H. Moulin. 2020. “Fair mixing: the case of dichotomous preferences.” *ACM Transactions on Economics and Computation*, 8, 4, 18:1–18:27. doi:10.1145/3417738.
- H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. 2017. “Justified Representation in approval-based committee voting.” *Social Choice and Welfare*, 48, 2, 461–485. doi:10.1007/s00355-016-1019-3.
- H. Aziz, E. Elkind, S. Huang, M. Lackner, L. Sánchez-Fernández, and P. Skowron. 2018. “On the complexity of extended and proportional justified representation.” In: *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, 902–909. doi:10.1609/aaai.v32i1.11478.

- H. Aziz, S. Gaspers, J. Gudmundsson, S. Mackenzie, N. Mattei, and T. Walsh. 2015. “Computational aspects of multi-winner approval voting.” In: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 107–115. <https://www.ifaamas.org/Proceedings/aamas2015/aamas/p107.pdf>.
- S. Banerjee, V. Gkatzelis, S. Hossain, B. Jin, E. Micha, and N. Shah. 2023. “Proportionally fair online allocation of public goods with predictions.” In: *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, 20–28. doi:10.24963/ijcai.2023/3.
- S. Barman, A. Khan, A. Maiti, and A. Sawarni. 2023. “Fairness and welfare quantification for regret in multi-armed bandits.” In: *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, 6762–6769. doi:10.1609/aaai.v37i6.25829.
- N. Barrot and J. Lang. 2016. “Conditional and sequential approval voting on combinatorial domains.” In: *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, 88–94. <https://www.ijcai.org/Proceedings/16/Papers/020.pdf>.
- N. Barrot, J. Lang, and M. Yokoo. 2017. “Manipulation of Hamming-based approval voting for multiple referenda and committee elections.” In: *Proceedings of the 16th Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 597–605. <https://www.ifaamas.org/Proceedings/aamas2017/pdfs/p597.pdf>.
- A. Bogomolnaia, H. Moulin, and R. Stong. 2005. “Collective choice under dichotomous preferences.” *Journal of Economic Theory*, 122, 2, 165–184. doi:10.1016/j.jet.2004.05.005.
- N. Bostrom. 2009. “Moral uncertainty – towards a solution?” In: *Overcoming Bias*. <https://www.overcomingbias.com/p/moral-uncertainty-towards-a-solution.html>.
- F. Brandl, F. Brandt, D. Peters, and C. Stricker. 2021. “Distribution rules under dichotomous preferences: Two out of three ain’t bad.” In: *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, 158–179. doi:10.1145/3465456.3467653.
- M. Brill, R. Freeman, S. Janson, and M. Lackner. 2023. “Phragmén’s voting methods and justified representation.” *Mathematical Programming*. doi:10.1007/s10107-023-01926-8.
- M. Brill, P. Gözl, D. Peters, U. Schmidt-Kraepelin, and K. Wilker. 2022. “Approval-based apportionment.” *Mathematical Programming*. doi:10.1007/s10107-022-01852-1.
- M. Brill, E. Markakis, G. Papanotiropoulos, and J. Peters. 2023. “Proportionality guarantees in elections with interdependent issues.” In: *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, 2537–2545. doi:10.24963/ijcai.2023/282.
- M. Brill and J. Peters. 2023. “Robust and verifiable proportionality axioms for multiwinner voting.” In: *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*, 301. doi:10.1145/3580507.3597785.
- L. Bulteau, N. Hazon, R. Page, A. Rosenfeld, and N. Talmon. 2021. “Justified representation for perpetual voting.” *IEEE Access*, 9, 96598–96612. doi:10.1109/ACCESS.2021.3095087.
- V. Conitzer, R. Freeman, and N. Shah. 2017. “Fair public decision making.” In: *Proceedings of the 2017 ACM Conference on Economics and Computation (EC)*, 629–646. doi:10.1145/3033274.3085125.
- T. Delemazure, T. Demeulemeester, M. Eberl, J. Israel, and P. Lederer. 2023. “Strategyproofness and proportionality in party-approval multiwinner elections.” In: *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, 5591–5599. doi:10.1609/aaai.v37i5.25694.
- V. Do, M. Hervouin, J. Lang, and P. Skowron. 2022. “Online approval committee elections.” In: *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, 251–257. doi:10.24963/ijcai.2022/36.
- E. Elkind, P. Faliszewski, J.-F. Laslier, P. Skowron, A. Slinko, and N. Talmon. 2017. “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)*, 494–501. doi:10.1609/aaai.v31i1.10612.
- E. Elkind, P. Faliszewski, P. Skowron, and A. Slinko. 2017. “Properties of multiwinner voting rules.” *Social Choice and Welfare*, 48, 3, 599–632. doi:10.1007/s00355-017-1026-z.
- E. Elkind, T. Y. Neoh, and N. Teh. 2024. “Temporal elections: Welfare, strategyproofness, and proportionality.” In: *Proceedings of the 27th European Conference on Artificial Intelligence (ECAI)*, 3292–3299. doi:10.3233/FAIA240877.
- E. Elkind, S. Obraztsova, and N. Teh. 2024. “Temporal fairness in multiwinner voting.” In: *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, 22633–22640. doi:10.1609/aaai.v38i20.30273.
- E. Elkind, S. Obraztsova, and N. Teh. 2025. “Verifying proportionality in temporal voting.” In: *Proceedings of the 39th AAAI Conference on Artificial Intelligence (AAAI)*, 13805–13813. arXiv: 2502.05949 (cs.GT). doi:10.1609/aaai.v39i13.33509.
- B. Fain, A. Goel, and K. Munagala. 2016. “The core of the participatory budgeting problem.” In: *Proceedings of the 12th Conference on Web and Internet Economics (WINE)*, 384–399. doi:10.1007/978-3-662-54110-4_27.
- M. Feffer, H. Heidari, and Z. Lipton. 2023. “Moral machine or tyranny of the majority?” In: *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, 5974–5982. doi:10.1609/aaai.v37i5.25739.
- R. Freedman, J. S. Borg, W. Sinnott-Armstrong, J. P. Dickerson, and V. Conitzer. 2020. “Adapting a kidney exchange algorithm to align with human values.” *Artificial Intelligence*, 283, 103261. doi:10.1016/j.artint.2020.103261.
- R. Freeman, A. Kahng, and D. M. Pennock. 2020. “Proportionality in approval-based elections with a variable number of winners.” In: *Proceedings of the 29th International Joint Conferences on Artificial Intelligence (IJCAI)*, 132–138. doi:10.24963/ijcai.2020/19.

- H. Greaves and O. Cotton-Barratt. 2023. “A bargaining-theoretic approach to moral uncertainty.” *Journal of Moral Philosophy*, 21, 1-2, 127–169. doi:10.1163/17455243-20233810.
- S. Janson. 2018. *Phragmén’s and Thiele’s election methods*. <https://arxiv.org/abs/1611.08826>. (2018). arXiv: 1611.08826 (math.HO).
- R. Kim, M. Kleiman-Weiner, A. Abeliuk, E. Awad, S. Dsouza, J. B. Tenenbaum, and I. Rahwan. 2018. “A computational model of commonsense moral decision making.” In: *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, 197–203. https://www.aies-conference.com/2018/contents/papers/main/AIES_2018_paper_130.pdf.
- A. Kozachinskiy, A. Shen, and T. Steifer. 2025. “Optimal bounds for dissatisfaction in perpetual voting.” In: *Proceedings of the 39th AAAI Conference on Artificial Intelligence (AAAI)*, 13977–13984. doi:10.1609/aaai.v39i13.33529.
- M. Lackner. 2020. “Perpetual voting: Fairness in long-term decision making.” In: *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 2103–2110. doi:10.1609/aaai.v34i02.5584.
- M. Lackner and J. Maly. 2023. “Proportional decisions in perpetual voting.” In: *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, 5722–5729. doi:10.1609/aaai.v37i5.25710.
- M. Lackner, J. Maly, and O. Nardi. 2023. “Free-riding in multi-issue decisions.” In: *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2040–2048. <https://www.ifaamas.org/Proceedings/aamas2023/pdfs/p2040.pdf>.
- M. Lackner and P. Skowron. 2023. *Multi-Winner Voting with Approval Preferences*. Springer. doi:10.1007/978-3-031-09016-5.
- J. Lang and L. Xia. 2016. “Voting in Combinatorial Domains.” In: *Handbook of Computational Social Choice*. Ed. by F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia. Cambridge University Press. Chap. 9. doi:10.1017/CBO9781107446984.010.
- M. K. Lee et al. 2019. “WeBuildAI: Participatory framework for algorithmic governance.” *Proceedings of the ACM on Human-Computer Interaction (HCI)*, 3. doi:10.1145/3359283.
- R. D. Luce. 1959. *Individual Choice Behavior: A Theoretical Analysis*. Wiley.
- T. Masařík, G. Pierczyński, and P. Skowron. 2024. “A generalised theory of proportionality in collective decision making.” In: *Proceedings of the 25th ACM Conference on Economics and Computation*, 734–754. doi:10.1145/3670865.3673619.
- F. Mohsin, L. Luo, W. Ma, I. Kang, Z. Zhao, A. Liu, R. Vaish, and L. Xia. 2021. “Making group decisions from natural language-based preferences.” In: *Proceedings of the 8th International Workshop on Computational Social Choice (COMSOC)*. <https://comsoc-community.org/assets/proceedings/comsoc-2021/MohsinEtAl.pdf>.
- T. Newberry and T. Ord. 2021. *The parliamentary approach to moral uncertainty*. FHI Technical Report 2021–2. (2021). <https://ora.ox.ac.uk/objects/uuid:b6b3bc2e-ba48-41d2-af7e-83f07c1fe141/files/svm40xs90j>.
- R. Noothigattu, S. S. Gaikwad, E. Awad, S. Dsouza, I. Rahwan, P. Ravikumar, and A. D. Procaccia. 2018. “A voting-based system for ethical decision making.” In: *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, 1587–1594. doi:10.1609/aaai.v32i1.11512.
- R. Page, E. Shapiro, and N. Talmon. 2020. *Electing the executive branch*. <https://arxiv.org/abs/2009.09734>. (2020). arXiv: 2009.09734 (cs.MA).
- D. Peters. 2024. “Proportional representation for artificial intelligence.” In: *Proceedings of the 27th European Conference on Artificial Intelligence (ECAI) (Frontiers in AI)*, 27–31. doi:10.3233/FAIA240463.
- D. Peters. 2018. “Proportionality and strategyproofness in multiwinner elections.” In: *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 1549–1557. <https://www.ifaamas.org/Proceedings/aamas2018/pdfs/p1549.pdf>.
- D. Peters. 2025. “The core of approval-based committee elections with few candidates.” In: *Proceedings of the 34th International Joint Conference on Artificial Intelligence (IJCAI)*, 4014–4022. doi:10.24963/ijcai.2025/447.
- D. Peters and M. Lackner. 2020. “Preferences single-peaked on a circle.” *Journal of Artificial Intelligence Research*, 68, 463–502. doi:10.1613/jair.1.11732.
- D. Peters, G. Pierczyński, and P. Skowron. 2021. “Proportional participatory budgeting with additive utilities.” In: *Advances in Neural Information Processing Systems*. Vol. 34, 12726–12737. <https://proceedings.neurips.cc/paper/2021/hash/69f8ea31de0c00502b2ae571fbab1f95-Abstract.html>.
- D. Peters and P. Skowron. 2020. “Proportionality and the limits of welfarism.” In: *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*. Full version arXiv:1911.11747, 793–794. arXiv: 1911.11747 (cs.GT). doi:10.1145/3391403.3399465.
- B. Phillips, E. Elkind, N. Teh, and T. Wąs. 2025. *Strengthening proportionality in temporal voting*. (2025). <https://arxiv.org/abs/2505.22513> arXiv: 2505.22513 (cs.GT).
- E. Phragmén. 1894. “Sur une méthode nouvelle pour réaliser, dans les élections, la représentation proportionnelle des partis.” *Öfversigt af Kongliga Vetenskaps-Akademiens Förhandlingar*, 51, 3, 133–137. <https://dominik-peters.de/archive/phragmen1894.pdf>.
- R. L. Plackett. 1975. “The analysis of permutations.” *Journal of the Royal Statistical Society Series C: Applied Statistics*, 24, 2, 193–202. doi:10.2307/2346567.
- L. Sánchez-Fernández, E. Elkind, M. Lackner, N. Fernández, J. A. Fisteus, P. B. Val, and P. Skowron. 2017. “Proportional justified representation.” In: *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 670–676. doi:10.1609/aaai.v31i1.10611.
- L. Sánchez-Fernández and J. A. Fisteus. 2019. “Monotonicity axioms in approval-based multi-winner voting rules.” In: *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 485–493. <https://www.ifaamas.org/Proceedings/aamas2019/pdfs/p485.pdf>.

- P. Skowron. 2021. “Proportionality degree of multiwinner rules.” In: *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, 820–840. doi:10.1145/3465456.3467641.
- P. Skowron and A. Górecki. 2022. “Proportional public decisions.” In: *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, 5191–5198. doi:10.1609/aaai.v36i5.20454.
- P. Skowron, M. Lackner, M. Brill, D. Peters, and E. Elkind. 2017. “Proportional rankings.” In: *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 409–415. doi:10.24963/ijcai.2017/58.
- T. N. Thiele. 1895. “Om Flerfold Valg.” *Oversigt over det Kongelige Danske Videnskabernes Selskabs Fordhandlinger*. <https://dominik-peters.de/archive/thiele1895.pdf>.
- T. Wang, J. Zhao, H. Yu, J. Liu, X. Yang, X. Ren, and S. Shi. 2019. “Privacy-preserving crowd-guided AI decision-making in ethical dilemmas.” In: *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, 1311–1320. doi:10.1145/3357384.3357954.

A More Axioms from the Literature

In this section, we consider other proportionality axioms that have been introduced in prior work by Lackner and Maly (2023), and discuss implication relationships among them (see Figure 1 in the main body). We also discuss Pareto efficiency. Note additionally that a recent paper by Phillips et al. (2025), written after the conference version of this paper, considers several additional proportionality axioms, including the axioms of FJR and EJR+ that strengthen our EJR axiom.

A.1 Quota for Closed Groups

Lackner and Maly (2023) proposed some relatively weak notions of proportionality in the perpetual voting framework. One of their axioms applies only to “simple profiles” (where the set of alternatives and the approval sets are the same in each round, with every voter approving just a single alternative, and the number of rounds being equal to the number of voters). Then “simple proportionality” requires that the number of appearances of alternatives in the output be equal to the number of voters approving them. This axiom is easily seen to be implied by all the other proportionality axioms we consider. Thus, we focus on their axioms based on profiles in which there are “closed groups”: groups of voters that submit identical approval sets in every round and such that any voter outside the group approves a disjoint set of alternatives.

Definition A.1 (Closed Group). A group $S \subseteq N$ is closed if for every round $j \in \{1, \dots, T\}$, it holds that (i) $A_j^j \neq \emptyset$, (ii) $A_j^i = A_j^{i'}$ for all $i, i' \in S$, and (iii) $A_j^i \cap A_j^{i'} = \emptyset$ for all $i \in S$ and $i' \in N \setminus S$.

For example, in Figure 5, the group $\{4, 5, 6, 7, 8\}$ is closed, and it is the unique closed group. In Figure 6, the group $\{2, 3\}$ is closed (but $\{1\}$ is not closed due to having empty approval sets in some rounds).

Lackner and Maly (2023) defined axioms which bound the number of rounds in which the decision should be one of the alternatives approved by a closed group.

Definition A.2 (Perpetual Lower Quota for Closed Groups). We say that a decision sequence $D = (d_1, \dots, d_T)$ satisfies *perpetual lower quota for closed groups*⁸ if, for every k , $1 \leq k \leq T$, writing $D_{:k} = (d_1, \dots, d_k)$ for the prefix of D of length k , it holds for every voter $i \in N$ who is part of a closed group S that $U_{D_{:k}}^i \geq \lfloor k \cdot \frac{|S|}{n} \rfloor$.

Lackner and Maly (2023) also proposed perpetual *upper quota* for closed groups, which is defined in the same way, except that the inequality at the end is “ $U_{D_{:k}}^i \leq \lceil k \cdot \frac{|S|}{n} \rceil$ ”. Because this is a different kind of property (avoiding over-representation rather than guaranteeing representation), and because none of our rules satisfy it, we will not study it in more detail.

Observe that perpetual lower quota for closed groups contains a requirement for all prefixes of the decision sequence D . This makes the axiom “perpetual”. This is natural and desirable if the decision-making process is online. However, in an offline setting, it is less natural, and so we can drop the perpetuality and define a weaker axiom that is more broadly applicable.

Definition A.3 (Lower Quota for Closed Groups). A decision sequence D satisfies *lower quota for closed groups* if for every closed group S , we have for every voter $i \in S$ that $U_D^i \geq \lfloor T \cdot |S|/n \rfloor$.

For an example, refer to Figure 6 where lower quota for closed groups would demand that each voter in $S = \{1\}$ should approve at least $\lfloor 6 \cdot 1/3 \rfloor = 2$ decisions of the outcome, while each voter in $S' = \{2, 3\}$ should approve at least $\lfloor 6 \cdot 2/3 \rfloor = 4$ decisions of the outcome.

We now show that lower quota for closed groups is implied by Weak PJR, and therefore also by the stronger properties PJR and (Weak) EJR.

⁸Lackner and Maly (2023) call this axiom just “lower quota for closed groups”. We added the word “perpetual” to the name since we will define below a weaker version suitable for the offline setting.

THEOREM A.4. *Weak PJR implies lower quota for closed groups.*

PROOF. Let D be a decision sequence that satisfies Weak PJR. Let $S \subseteq N$ be a closed group. Write $\ell = \lfloor T \cdot |S|/n \rfloor$. Then $|S| = (T \cdot |S|/n) \cdot \frac{n}{T} \geq \ell \cdot \frac{n}{T}$, and S agrees in every round since it is a closed group. Thus, Weak PJR implies that there are ℓ rounds where at least one member of S approves the decision of D . But because S is closed, this means that in those ℓ rounds, every member of S approves the decision. Hence D satisfies lower quota for closed groups. \square

Similarly, an *online* rule that satisfies Weak PJR will thereby satisfy *perpetual* lower quota for closed groups.

A.2 Priceability

D. Peters and Skowron (2020) introduced an axiom called *priceability* to ensure proportionality and (approximately) equal influence among voters in the setting of multi-winner voting. The essence of priceability lies in assigning each voter an equal budget of virtual money, which they can only spend on alternatives they approve. Thinking of the virtual money as voting power, the axiom ensures that each voter has approximately equal influence over the outcome. Lackner and Maly (2023) extended priceability to the setting of perpetual voting and showed that Sequential Phragmén (which they call Perpetual Phragmén) satisfies perpetual priceability. We state their axiom below and discuss its relation to our axioms.

Priceability is based on the idea that it costs 1 unit of money to pay for the decision in 1 round. Then a “price system” is used to split this cost between voters that approve this decision.

Definition A.5 (Price System). Suppose we are looking at the first k rounds.⁹ A *price system* is a pair $(B, \{p_j\}_{j \leq k})$ where $B \geq 0$ is the *budget* that each voter starts with, and $p_j : N \times C_j \rightarrow [0, 1]$, for each $j \in \{1, \dots, k\}$, is a *payment function* where $p_j(i, c)$ will indicate how much agent i pays for candidate c in round j . We say that a decision sequence $D = (d_1, \dots, d_k)$ of length k is *supported by the price system* $(B, \{p_j\}_{j \leq k})$ if

- (P1) $p_j(i, c) = 0$ whenever $c \notin A_j^i$, i.e., no voter pays for non-approved alternatives.
- (P2) $\sum_{j=1}^k \sum_{c \in C_j} p_j(i, c) \leq B$, i.e., voters cannot spend more than their budget.
- (P3) $\sum_{i \in N} p_j(i, d_j) = 1$ for $j \in \{1, \dots, k\}$, i.e., each decision of D gets a total payment of 1.
- (P4) $\sum_{i \in N} p_j(i, d) = 0$ for $j \in \{1, \dots, k\}$ and $d \neq d_j$, i.e., decisions not included in D do not receive any payments.

To define priceability for perpetual voting, Lackner and Maly (2023) propose the concept of a *minimal* price system which intuitively requires that no money is wasted at any time step.

Definition A.6 (Minimal Price System, Lackner and Maly 2023). Consider a decision sequence $D = (d_1, \dots, d_k)$ of length k supported by a price system $(B, \{p_j\}_{j \leq k})$ satisfying conditions (P1)-(P4) of Definition A.5. If $k = 0$ (so that $D = ()$), we say that the price system is *minimal* if $B = 0$. If $k \geq 1$, we say that the price system is minimal if it satisfies the following two conditions:

- (P5) there exists a minimal price system $(B^*, \{p_j^*\}_{j \leq k-1})$ that supports (d_1, \dots, d_{k-1})
- (P6) there are no B', d'_k and p'_k such that $B^* \leq B' < B$ and $(B', \{p_j\}_{j \leq k-1} \cup \{p'_k\})$ is a price system supporting $(d_1, \dots, d_{k-1}, d'_k)$.

We can now define perpetual priceability.

Definition A.7 (Perpetual Priceability, Lackner and Maly 2023). A decision sequence D satisfies perpetual priceability if there exists a minimal price system that supports D .

⁹We use a variable k instead of the total number of rounds T since the definition of perpetual priceability concerns decision sequences that are shorter than the full time horizon T .

With the above definition in place, we will now show that perpetual priceability implies Weak PJR. (Similarly, in multi-winner voting, priceability implies multi-winner PJR (D. Peters and Skowron 2020).) The proof below follows the same style as a proof by Lackner and Maly (2023) that perpetual priceability implies perpetual lower quota for closed groups (Definition A.2).

THEOREM A.8. *For decision sequences of length T , Perpetual Priceability implies Weak PJR.*

PROOF. Let $D = (d_1, d_2, \dots, d_T)$ be a decision sequence of length T that satisfies perpetual priceability. Assume for a contradiction that D violates Weak PJR. Thus, there is $\ell \in \mathbb{N}$ and a group $S \subseteq N$ which agrees in all the T rounds, has size $|S| = s \geq \ell \cdot \frac{n}{T}$, but there are fewer than ℓ rounds in which at least one member of S approves the decision in D .

Let $(B, \{p_j\}_{j \leq T})$ be a minimal price system that supports D . Note that $B \geq T/n$, as otherwise it would not be possible to pay for T decisions. Assume first that $B = T/n$. Then no budget is left after round T , and thus every voter has spent their entire budget B . Thus, the total payment of voters in S is $|S| \cdot B = |S| \cdot \frac{T}{n} \geq \ell$. Since each decision is made using 1 unit, in total S must have paid for the decision in at least ℓ rounds. As a voter only pays for an alternative they approve, it follows that in each of those ℓ rounds, at least one voter in S approves the decision. Hence, Weak PJR is satisfied.

So assume $B > \frac{T}{n}$. Let $\epsilon = s \cdot B - \ell$. Note that $\epsilon > \ell \frac{n}{T} \cdot \frac{T}{n} - \ell = 0$. By inductively using the property (P5) of the definition of minimal price systems, for every $1 \leq r \leq T$, we can find a price system $(B_{r-1}, \{p_j^{(r-1)}\}_{j \leq r-1})$ that witnesses the minimality of the price system $(B_r, \{p_j^{(r)}\}_{j \leq r})$, where $(B_T, \{p_j^{(T)}\}_{j \leq T}) = (B, \{p_j\}_{j \leq T})$.

We claim that $B_1 < B$. Observe that in the first round $B_1 \leq 1/s$ must hold, as with a budget of $1/s$ the voters in S can together afford one of the alternatives that they jointly approve; thus $B_1 > 1/s$ would contradict minimality in round 1. Moreover, we can assume that $\ell \geq 1$ as otherwise Weak PJR is trivially satisfied. Thus $s \geq \frac{n}{T}$, and hence $n/s \leq T$. Also, by assumption, $B > T/n$. Putting all of this together, we get

$$B_1 \leq \frac{1}{s} \leq \frac{T}{n} < B.$$

Now let r^* be the largest index r for which $B_r < B$. Recall that at the end, $B_T = B$, and thus $r^* < T$. In addition, by minimality of the price system for the entire sequence, we see that $B_{r^*+1} = B$. We claim that there exists a budget B' with $B_{r^*} \leq B' < B_{r^*+1} = B$, an alternative $a' \in C_{r^*+1}$, and a price p'_{r^*+1} such that $(B', \{p_j^{(r^*)}\}_{j \leq r^*} \cup \{p'_{r^*+1}\})$ is a price system supporting the decision sequence $(d_1, \dots, d_{r^*}, a')$. This would be a contradiction to the minimality of $(B_{r^*+1}, \{p_j^{(r^*+1)}\}_{j \leq r^*+1})$.

Let a' be an alternative approved by everyone in S in round $r^* + 1$. Furthermore, let

$$B' = \max(B - \epsilon/s, B_{r^*}).$$

Observe that $B_{r^*} \leq B' < B$. We now compute the remaining budget of voters in S , after paying for the first r^* rounds according to the price system $\{p_j^{(r^*)}\}_{j \leq r^*}$ with initial budget B' . By the failure of Weak PJR, at most $\ell - 1$ decisions in the entire sequence D are approved by some member of S , and thus the same is true for the sequence restricted to the first r^* rounds. Thus, the members of S have spent at most $\ell - 1$ units in the first r^* rounds. Since they start out with a total budget of sB' , we see that the voters in S have a remaining budget of at least

$$s \cdot B' - (\ell - 1) \geq s \cdot (B - \frac{\epsilon}{s}) - (\ell - 1) = (s \cdot B - \ell) - \epsilon + 1 = 1,$$

by our choice of ϵ . Hence, now we can define p'_{r^*+1} such that $\sum_{i \in S} p'_{r^*+1}(i, a') = 1$ and $p'_{r^*+1}(i', c) = 0$ whenever $i' \notin S$ or $c \neq a'$. Hence, this is a price system that supports $(d_1, \dots, d_{r^*}, a')$ which contradicts the minimality of $(B, \{p_j\}_{j \leq r^*+1})$. \square

Round	1	2	3	4
Voter 1	{a}	{g}	{g}	{g}
Voter 2	{b}	{g}	{g}	{g}
Voter 3	{c}	{g}	{g}	{g}
Voter 4	{d}	{h}	{h}	{h}
Voter 5	{e}	{h}	{h}	{h}
Voter 6	{f}	{h}	{h}	{h}

Fig. 15. Example instance where a rule may produce a perpetually priceable outcome which may not satisfy JR.

As perpetual priceability implies Weak PJR, one may wonder whether it implies stronger axioms. Since Sequential Phragmén satisfies perpetual priceability (Lackner and Maly 2023) but fails Weak EJR (Theorem 4.2), it follows that perpetual priceability does not imply Weak EJR. The next result shows that it does not imply PJR either.

THEOREM A.9. *Perpetual Priceability does not imply PJR.*

PROOF. Consider the decision instance shown in Figure 15 with $T = 4$ rounds and $N = \{v_1, \dots, v_6\}$ with each voter having a single-alternative approval set in each round. Consider the decision sequence $D = (a, h, h, h)$. We show that it satisfies perpetual priceability but fails PJR.

In the first round, everyone has a disjoint approval set, so no matter which alternative is chosen in this round, one voter needs to pay its entire cost. Thus, any price system supporting the decision in the first round must have $B_1 = 1$. Setting $p_1(v_1, a) = 1$ and $p_1(i, c) = 0$ for all $i \neq v_1$ and all $c \neq a$, we obtain a minimal price system (B_1, p_1) supporting $D' = (a)$.

Thus, from the first round, we know that any minimal price system will need to use a per-voter budget of at least 1; we will actually construct a price system for all four rounds with budget exactly 1, which implies minimality. Now, after round 1, voter v_1 has spent all of its money while other voters have their entire budget remaining. In round 2, h can be bought by each of its supporters paying $\frac{1}{3}$ unit as $3 \cdot \frac{1}{3} = p = 1$. Similarly, h can be bought in rounds 3 and 4 as well without requiring to increase the budget of anyone. This defines payment functions p_2, p_3 , and p_4 for those rounds, and setting $B = 1$, we obtain a minimal price system $(B, \{p_j\}_{j \leq 4})$ supporting D . Thus, D satisfies perpetual priceability.

However, D fails PJR: Consider the coalition $S = \{v_2, v_3\}$ which agrees in the last $k = 3$ rounds. For $\ell = 1$, we have $|S| \geq \ell \cdot \frac{n}{k} = \frac{6}{3} = 2$. Thus, PJR requires that at least 1 decision in the outcome must be approved by S . However, no one in S approves either a or h , and thus D fails PJR. Hence, perpetual priceability does not imply PJR. \square

Remark A.10. In the example in Theorem A.9, we used a requirement of $\ell = 1$ rounds, so the proof of Theorem A.9 shows that perpetual priceability does not even imply JR.

Note that perpetual priceability implies PJR for coalitions that agree “initially”, that is they agree in an initial set of rounds $R^* = \{1, \dots, k\}$ (since by definition of perpetual priceability, the decision sequence restricted to the first k decisions is itself perpetually priceable, and therefore this restricted sequence satisfies Weak PJR, which is enough to imply PJR with initial agreement). Note that the example in Theorem A.9 does not feature initial agreement.

Rounds	1–2	3–7
Voter 1	{a}	{c}
Voter 2	{a}	{d}
Voters 3, 4, 5, 6, 7	{b}	{b}

Fig. 16. Example instance where both MES and Sequential Phragmén produce Pareto inefficient outcomes.

In future work, it would be interesting to define a notion of priceability that applies to our setting without an implicit assumption of being online (that is, without necessarily requiring that all prefixes of the decision sequence are themselves priceable).

A.3 Pareto Efficiency

Pareto efficiency forbids that we use our available resources suboptimally. Formally, an outcome is Pareto efficient when it is impossible to change the outcome so as to make some individual better off without making anyone else worse off.

Definition A.11 (Pareto Efficiency). A decision sequence D is *Pareto efficient* if there does not exist another decision sequence $D' \in C_1 \times C_2 \times \dots \times C_T$ such that for all $i \in N$, $U_{D'}^i \geq U_D^i$ and there is some voter $i \in N$ such that $U_{D'}^i > U_D^i$.

Unlike our other axioms, Pareto efficiency is not about proportionality (since a Pareto efficient outcome can be very non-proportional). However, it is interesting to see if any of our rules satisfy Pareto efficiency, and therefore combine an efficiency guarantee with proportionality. Since it maximizes a function of voter utilities, PAV satisfies Pareto efficiency, while also satisfying our strongest proportionality notions. On the other hand, MES and Sequential Phragmén fail Pareto efficiency. This makes PAV a compelling rule for multi-issue collective decision making, when an NP-hard offline method is acceptable. (The outcome of Local-Search PAV need not be Pareto-efficient.)¹⁰

THEOREM A.12. *The decision sequence produced by PAV is Pareto efficient.*

PROOF. Suppose for a contradiction that the decision sequence D produced by PAV is not Pareto efficient. Then there exists another decision sequence say D' such that for all $i \in N$, $U_{D'}^i \geq U_D^i$ and there is some voter $v \in N$ such that $U_{D'}^v > U_D^v$. Thus, because the harmonic series is strictly increasing, the PAV-score of D' is

$$\text{PAV-score}(D') = \sum_{i \in N} \left(1 + \frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{U_{D'}^i} \right) > \text{PAV-score}(D)$$

which contradicts that PAV selects the decision sequence which maximizes the PAV-score. \square

THEOREM A.13. *The decision sequence produced by MES and Sequential Phragmén may not be Pareto efficient.*

PROOF. Figure 16 shows an example where both MES and Sequential Phragmén fail Pareto efficiency, with $T = 7$ rounds and $n = 7$ voters. Here, voters 1 and 2 agree in the first 2 rounds on a , but neither of them agrees with anyone else in other rounds. Meanwhile, voters 3, 4, 5, 6, 7 agree in all rounds on b . Both Sequential Phragmén

¹⁰Here is an example where a local optimum of the PAV score may fail Pareto efficiency. The example also works for multi-winner elections. There are two rounds with identical approval sets in both rounds, and five voters with approval sets $\{a, c, d\}$, $\{a, c\}$, $\{a, d\}$, $\{b, c\}$, $\{b, d\}$. The outcome (a, b) is Pareto dominated by (c, d) because the latter is strictly preferred by the first voters, and the other four voters are indifferent. However, no swap can strictly improve the PAV score of (a, b) .

and MES would produce $D = (b, b, b, b, b, c, d)$ (or breaking ties differently, with d in round 6 and c in round 7) but the decision sequence $D' = (a, a, b, b, b, b, b)$ Pareto dominates D as $U_{D'}^1 = 2 > U_D^1 = 1$ while $U_{D'}^i \geq U_D^i$ for all $i \in N$. Intuitively, it is better to satisfy voters 1 and 2 using the efficient choice a which is available in the first two rounds, rather than using the less efficient alternatives c and d . But due to their (semi-)online nature, neither MES nor Sequential Phragmén can know that an alternative as good as a will not be available in later rounds. \square

In the example in the proof of [Theorem A.13](#), Offline MES can select the same Pareto inefficient decision sequence, so it also fails Pareto efficiency.

B Omitted Proofs from the Main Text

B.1 Impossibility of Stronger Guarantees: More Details

Here we give proofs and additional results related to the discussion in [Section 5](#). The results are based on a counterexample construction formally described in the following result. It has the following structure: fix some coalition size $s \leq n$ of interest. In the first k rounds, for each of the $\binom{n}{s}$ subsets $S \subseteq N$ of size s , there is an alternative a_S with all voters in S approving a_S . Thus, all these coalitions agree in k rounds, and thus the impossible axioms we will consider guarantee all of these coalitions some representation. However, k rounds is not enough to “hit” enough of these coalitions, and in the remaining $T - k$ rounds, voters are in complete disagreement (i.e., the approval sets of any two voters are disjoint) and so there are not enough rounds to make up the deficit accrued in the first k rounds.

THEOREM B.1. *Let $\epsilon > 0$. Choose some $k > \lceil \frac{1-\epsilon}{\epsilon} \rceil$ and some $T > k$. Then for sufficiently large n , there exists an instance with n voters and time horizon T such that for every decision sequence D , there is a coalition $S \subseteq N$ that agrees in k rounds and has size $|S| \geq (1 - \epsilon) \cdot \frac{n}{k}$, yet none of the voters in S approve any of the decisions of D .*

PROOF. Pick $n \in \mathbb{N}$ such that $n \geq \frac{k(T+1)}{\epsilon k + \epsilon - 1}$. Note for later that (because $\epsilon k + \epsilon - 1 > 0$, since $k > \lceil \frac{1-\epsilon}{\epsilon} \rceil$), we have

$$n(\epsilon k + \epsilon - 1) \geq k(T + 1).$$

Adding nk to both sides, and rearranging,

$$nk \geq n \cdot (1 - \epsilon) \cdot k + n(1 - \epsilon) + k(T + 1).$$

Finally, dividing by k , we get

$$n \geq n(1 - \epsilon) + \frac{n(1-\epsilon)}{k} + T + 1. \quad (2)$$

We now construct an instance with n voters and time horizon T . Write $s = \lceil (1 - \epsilon) \cdot \frac{n}{k} \rceil$. In each of the first k rounds, there are $\binom{n}{s}$ many alternatives, one for each coalition $S \subseteq N$ with $|S| = s$, that is approved by exactly the voters in S . In the remaining $T - k$ rounds, there are n alternatives, each approved by a distinct voter (thus each voter has a singleton approval set in those rounds, and no two voters approve the same alternative).

Let D be any decision sequence. We first compute how many voters approve at least 1 decision in D . In the first k rounds, the decisions in D can satisfy at most $k \cdot s$ different voters. In the remaining $T - k$ rounds, at most 1 voter can be satisfied per round, so at most $T - k$ in total. Hence the number of satisfied voters is at most

$$\begin{aligned} k \cdot s + (T - k) &= k \cdot \lceil (1 - \epsilon) \cdot \frac{n}{k} \rceil + (T - k) \\ &\leq n \cdot (1 - \epsilon) + k + (T - k) \\ &= n \cdot (1 - \epsilon) + T. \end{aligned}$$

Thus, the number of voters who approve none of the decisions in D is at least

$$n - (n \cdot (1 - \epsilon) + T) \stackrel{(2)}{\geq} \frac{n(1-\epsilon)}{k} + 1 \geq s.$$

Take any set $S \subseteq N$ of exactly s voters who approve none of the decisions in D . By construction of the instance, they agree in the first k rounds (since they approve the alternative constructed for that coalition), and $|S| = s = \lceil (1 - \epsilon) \cdot \frac{n}{k} \rceil$. Thus, S satisfies the conditions promised in the theorem statement. \square

From [Theorem B.1](#), it follows immediately that for all $\epsilon > 0$, a decision sequence satisfying ϵ -Stronger PJR may not exist (as defined in [Section 5](#)), so [Theorem 5.1](#) from the main body is proven. In fact, this proof establishes the non-existence even for an analogously defined notion of ϵ -Stronger JR.

We can use the same counterexample construction to show that some other related axioms are impossible to satisfy in general. In particular, a natural idea of defining PJR for our setting is by requiring that a group large enough to deserve ℓ decisions only needs to agree in ℓ rounds (instead of in T rounds). This corresponds perhaps most closely to the original definitions of PJR and EJR for multi-winner voting which, in the context of electing k candidates, requires agreement on ℓ candidates not k candidates. It also corresponds to the definition of PJR proposed by [Freeman et al. \(2020\)](#) who study multi-winner voting with a variable number of winners, which is a special case of our model where in each round, only 2 alternatives are available (“add c to the committee”, “do not add c to the committee”). However, this version of the PJR axiom cannot always be satisfied.

COROLLARY B.2 (ℓ -AGREEMENT PJR NEED NOT EXIST). *There exists an instance where no decision sequence D satisfies “ ℓ -Agreement PJR”, defined to require that for every $\ell \in \mathbb{N}$ and every group of voters $S \subseteq N$ that agrees in at least ℓ rounds and that has size $|S| \geq \ell \cdot \frac{n}{T}$, there are at least ℓ rounds j with $d_j \in \bigcup_{i \in S} A_j^i$.*

PROOF. Invoke [Theorem B.1](#) with $\epsilon = 0.5$, $k = 4$, and $T = 40$. The theorem gives an instance where for every decision sequence D , there is a group of voters $S \subseteq N$ which agrees in 4 rounds, has size $|S| \geq (1 - \epsilon) \cdot \frac{n}{k} = 0.125 \cdot n \geq 0.1 \cdot n = 4 \cdot \frac{n}{T}$, but none of the voters in S approve any of the decisions in D , violating ℓ -Agreement PJR. \square

When there are only 2 candidates in every round, Sequential Phragmén, MES, and PAV all satisfy this property ([Freeman et al. 2020](#)).

Following [Do et al. \(2022\)](#), we now consider two relaxations of the ℓ -Agreement PJR axiom which introduce multiplicative approximations by a factor α . We show that for each constant $0 < \alpha < 1$, there are instances where even this relaxed version is not satisfiable.

The first relaxation only requires the group to approve $\lfloor \alpha \cdot \ell \rfloor$ decisions instead of ℓ decisions.

COROLLARY B.3 (α - ℓ -AGREEMENT PJR NEED NOT EXIST). *Let $0 < \alpha < 1$. There exists an instance where no decision sequence D satisfies “ α - ℓ -Agreement PJR”, defined to require that for every $\ell \in \mathbb{N}$ and every group of voters $S \subseteq N$ that agrees in at least ℓ rounds and that has size $|S| \geq \ell \cdot \frac{n}{T}$, there are at least $\lfloor \alpha \cdot \ell \rfloor$ rounds $j \in R$ with $d_j \in \bigcup_{i \in S} A_j^i$.*

PROOF. Invoke [Theorem B.1](#) with $\epsilon = 0.5$, $k = \lceil \frac{1}{\alpha} \rceil$, and $T = \lceil \frac{8}{\alpha^2} \rceil$ to obtain an instance. Assume that D is a decision sequence satisfying α - ℓ -Agreement PJR. The theorem says that there exists a group S which agrees in k rounds and has size $|S| \geq 0.5 \cdot \frac{n}{k}$ and such that none of the decisions in D are approved by any member of S .

Write $\ell = k$. Note that S agrees in ℓ rounds and (using the fact that $\lceil x \rceil \leq 2x$ for all $x > 1$) that

$$|S| \geq 0.5 \cdot \frac{n}{k} = 0.5 \cdot \frac{n}{\lceil 1/\alpha \rceil} \geq 0.5 \cdot \frac{n}{2/\alpha} = \frac{2}{\alpha} \cdot \frac{n}{8/\alpha^2} \geq \ell \cdot \frac{n}{T}.$$

Hence, α - ℓ -Agreement PJR requires that in at least $\lfloor \alpha \cdot \ell \rfloor = \lfloor \alpha \cdot \lceil \frac{1}{\alpha} \rceil \rfloor \geq \lfloor \alpha \cdot \frac{1}{\alpha} \rfloor = 1$ rounds, at least one member of S approves the decision of D , contradiction. \square

The second relaxation requires that the coalition agrees in ℓ/α rounds to be guaranteed to approve ℓ decisions.

COROLLARY B.4 (ℓ/α -AGREEMENT-PJR NEED NOT EXIST). *Let $0 < \alpha < 1$. There exists an instance where no decision sequence D satisfies “ ℓ/α -Agreement-PJR”, defined to require that for every $\ell \in \mathbb{N}$ and every group of voters $S \subseteq N$ that agrees in at least ℓ/α rounds and that has size $|S| \geq \ell \cdot \frac{n}{T}$, there are at least ℓ rounds j with $d_j \in \bigcup_{i \in S} A_j^i$.*

PROOF. Invoke [Theorem B.1](#) with $\epsilon = 0.5$, $k = \lceil \frac{1}{\alpha} \rceil$, and $T = \lceil \frac{4}{\alpha} \rceil$ to obtain an instance.

Assume that D is a decision sequence satisfying ℓ/α -Agreement-PJR. The theorem says that there exists a group S which agrees in k rounds and has size $|S| \geq 0.5 \cdot \frac{n}{k}$ and such that none of the decisions in D are approved by any member of S .

Write $\ell = 1$. Note that S agrees in $k = \lceil \frac{1}{\alpha} \rceil \geq \frac{\ell}{\alpha}$ rounds and that

$$|S| \geq 0.5 \cdot \frac{n}{k} \geq 0.5 \cdot \frac{n}{2/\alpha} = \frac{n}{4/\alpha} \geq \ell \cdot \frac{n}{T}.$$

Hence ℓ/α -Agreement PJR requires that in at least $\ell = 1$ rounds, at least one member of S approves the decision of D , contradiction. \square

We have seen several axioms that cannot always be satisfied. However, we could hope to find rules that satisfy these axioms whenever they are run on an instance where there exists a decision sequence satisfying this axiom. While there are certainly artificial *offline* rules doing this (e.g., the rule that outputs an arbitrary sequence satisfying the axiom should one exists, and otherwise outputs the PAV sequence), it is not clear if there are natural such rules.

Online or semi-online rules, however, provably cannot satisfy such a property, as one can see by adapting the counterexample construction of [Theorem B.1](#): The first $T - 1$ rounds remain exactly the same as in that construction. We observe the decisions made by the online rule in these rounds. In the last round, take $s + 1$ voters who do not approve any of the decisions made thus far, and have each of them approve a distinct alternative (like in the original construction). For the remaining $n - (s + 1)$ voters, introduce an alternative c that they all approve. No matter which alternative the rule chooses in round T , there remain s voters who do not approve any decision, leading to a violation. However, a violation could have been avoided by choosing in round 1 an alternative commonly approved by s of the $s + 1$ voters, and choosing c in round T . Thus, we have arrived at a situation where a violation-free decision sequence exists but an online (or semi-online) rule cannot find it as we alter the input based on its past decisions. This proves [Theorem 5.2](#) from the main body. This kind of construction works for all the axiom variants we have considered in this section.

B.2 Can Online Rules Satisfy Weak EJR?

We have seen that the online rule Sequential Phragmén fails Weak EJR, while the semi-online MES satisfies Weak EJR. Is it possible for an online rule to satisfy EJR or Weak EJR? For EJR, this question has been answered in the negative by [Elkind, Obraztsova, et al. \(2025\)](#), even for semi-online rules. For Weak EJR, we do not know the answer to this question.¹¹ However, we show that any online rule that satisfies Weak EJR for sequential decision making can be translated into a multi-winner voting rule that satisfies EJR and committee monotonicity. As [Lackner and Skowron \(2023, Sec. 7.1\)](#) write, a “main open question is whether there exist [approval-based multi-winner] rules that satisfy EJR and committee monotonicity”. Thus, we leave finding an online rule that satisfies (Weak) EJR as a challenging problem for future work.

In the remainder of this section, we describe the way in which an online decision rule can be used as a multi-winner voting rule. To do so, we must first give basic definitions about the latter topic.

¹¹In the special case where the set of available alternatives and the approval sets are the same across rounds, a known online method satisfies EJR ([Brill, Gözl, et al. 2022, Theorem 4.3](#)). Sequential Phragmén fails Weak EJR even in this special case ([Theorem 4.2](#)).

Definition B.5 (Multi-winner Voting Rule). Given a set of voters N , a set of alternatives C , a profile of approval sets $A = (A^i)_{i \in N}$ (with $A^i \subseteq C$ for all $i \in N$), and a desired committee size k , a *multi-winner voting rule* f outputs a committee $f(N, C, A, k) = W \subseteq C$ with $|W| = k$.

Note that we have defined multi-winner voting rules to be resolute, so they break any ties internally.

Intuitively, a multi-winner voting rule is said to be committee monotonic if it selects the winning committee for size k by first computing the winning committee for size $k - 1$, and then adding a k th alternative. Thus, when an alternative is a member of the winning committee, it can never become a losing alternative when the committee size k is increased.

Definition B.6 (Committee Monotonicity, Elkind, Faliszewski, Skowron, et al. 2017). A multi-winner voting rule f satisfies *committee monotonicity* if for all N, A, C , and $k \geq 1$, we have $f(N, C, A, k - 1) \subseteq f(N, C, A, k)$.

We refer to the EJR axiom originally proposed for the multi-winner voting context by Aziz, Brill, et al. (2017) as Multi-Winner EJR to distinguish it from our axioms for the sequential context.

Definition B.7 (Multi-Winner EJR). Let k be the desired committee size. A committee W satisfies *Multi-Winner EJR* if for every $\ell \in \mathbb{N}$, for every group of voters $S \subseteq N$ of size $|S| \geq \ell \cdot \frac{n}{k}$ with $|\bigcap_{i \in S} A^i| \geq \ell$, there is at least one voter $i \in S$ with $|W \cap A^i| \geq \ell$.

We can now formally state our reduction.

THEOREM B.8. *An online decision rule that satisfies Weak EJR in sequential decision making induces a multi-winner voting rule that simultaneously satisfies Multi-Winner EJR and committee monotonicity.*

PROOF. Let f be an online decision rule that satisfies Weak EJR. We will derive a multi-winner voting rule by describing a procedure that, given a profile of approval sets, selects alternatives in some order one-by-one until all alternatives have been selected. Let us label the alternatives in order of selection as d_1, \dots, d_m . To obtain a committee of size k , we will take $W = \{d_1, \dots, d_k\}$. By construction, such a method must be committee monotonic.

So let N be a set of voters, C be a set of alternatives, and $A = (A^i)_{i \in N}$ be a profile of approval sets over C . We use our online rule f in the following adaptively constructed sequence of $T = |C|$ rounds. In round 1, the set of available alternatives is $C_1 = C$, and the approval sets are $A_1^i = A^i$ for every $i \in N$. We apply f to this round and thereby obtain d_1 , the first selected alternative.

Thereafter, in round $j = 2, \dots, m$, we have already obtained a sequence of decisions d_1, \dots, d_{j-1} which we will assume inductively are pairwise distinct. In round $j \in R$, we take the set of available alternatives to be $C_j = C \setminus \{d_1, \dots, d_{j-1}\}$. Further, we take the approval sets to be $A_j^i = A^i \cap C_j$ for all $i \in N$. Then we apply our rule f to obtain the next decision d_j , which by choice of C_j is distinct from all prior decisions.

To finish, we need to prove that the multi-winner voting rule we have defined satisfies Multi-Winner EJR. Let k be a desired committee size, and let $W = \{d_1, \dots, d_k\}$. Let $\ell \in \mathbb{N}$ and suppose that $S \subseteq N$ is a group of voters with $|S| \geq \ell \cdot \frac{n}{k}$ and $|\bigcap_{i \in S} A^i| \geq \ell$. Assume for a contradiction that every member of S approves at most $\ell - 1$ alternatives in W . Then there must exist an alternative a that all members of S approve, but that is not a winner: $a \in \bigcap_{i \in S} A^i \setminus W$. Now consider the decision instance given by the first k rounds constructed above (thus, having time horizon $T = k$). For that decision instance, note that S forms a group of voters that agrees in every round (because $a \in C_j$ for all $j = 1, \dots, k$). Also $|S| \geq \ell \cdot \frac{n}{T}$. Hence, because f satisfies Weak EJR, there is a voter $i \in S$ who approves at least ℓ of the decisions made by f on this decision instance, a contradiction. \square

Received 19 March 2025; accepted 23 September 2025