



HAL
open science

Proportional Aggregation of Preferences for Sequential Decision Making

Nikhil Chandak, Shashwat Goel, Dominik Peters

► **To cite this version:**

Nikhil Chandak, Shashwat Goel, Dominik Peters. Proportional Aggregation of Preferences for Sequential Decision Making. Proceedings of the 38th Annual AAAI Conference on Artificial Intelligence, AAAI, Feb 2024, Vancouver, Canada. hal-04486833

HAL Id: hal-04486833

<https://hal.science/hal-04486833>

Submitted on 2 Mar 2024

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

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

Proportional Aggregation of Preferences for Sequential Decision Making

Nikhil Chandak¹, Shashwat Goel¹, and Dominik Peters²

¹IIT Hyderabad

²CNRS, LAMSADE, Université Paris Dauphine - PSL

Abstract

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 which of these alternatives they approve. Instead of going with the most popular choice in each round, we aim for proportional representation. We formalize this aim using axioms based on Proportional Justified Representation (PJR), which were proposed in the literature on multi-winner voting and were recently adapted to multi-issue decision making. The axioms require that every group of $\alpha\%$ of the voters, if it agrees in every round (i.e., approves a common alternative), then those voters 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). The first two are polynomial-time computable, and the latter is based on an NP-hard optimization, but it admits a polynomial-time local search algorithm that satisfies the same axiomatic properties. We present empirical results about the performance of these rules based on synthetic data and U.S. political elections. We also run experiments where votes are cast by preference models trained on user responses from the moral machine dataset about ethical dilemmas.

1	Introduction	2	7	Discussion and Future Work	18
2	Related Work	4	8	Bibliography	18
3	Problem Formulation	5	A	More Axioms from the Literature	22
3.1	Model	5	A.1	Quota for Closed Groups	22
3.2	Axioms	5	A.2	Priceability	22
3.3	Methods	6	A.3	Pareto Efficiency	25
4	Satisfying the Axioms	8	B	Omitted Proofs from the Main Text	25
4.1	Online: Sequential Phragmén	8	B.1	MES Fails Strong PJR and Strong EJR	25
4.2	Semi-online: Method of Equal Shares (MES)	9	B.2	Impossibility of Stronger Guarantees: More Details	26
4.3	Offline: Proportional Approval Voting (PAV)	10	B.3	Can Online Rules Satisfy EJR?	28
5	Impossibility of Stronger Guarantees	11	C	Full Results from Experiments	30
6	Experiments	12	C.1	Synthetic Data	30
6.1	Synthetic Data	13	C.2	Learning Preferences from the Moral Machine	34
6.2	Political Data	15			
6.3	Learning Preferences from the Moral Machine	16			

1. Introduction

We consider the problem of making multiple 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 simplest way of making these decisions is to take, in each round, the alternative with the most supporters.

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 will be undesirable, either due to fairness concerns or because the decision maker wishes to hear from each voter to an equal extent.

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 from 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:

- *Hiring decisions*: Consider a department hiring 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 candidate on that topic, then at least 1 such candidate should be hired every 5 years.
- *Uncertainty about the goal*: Suppose we are managing an intelligent system that is repeatedly deciding on an action in order to maximize an objective function that we have specified. But we may not be certain about our desired objective function (for example because we have *moral uncertainty*). Literature in reinforcement learning explores similar problems like *reward uncertainty* [Liang et al., 2022] and in the *cooperative IRL framework* [Hadfield-Menell et al., 2016]. We can decide on a mixture of objectives (20% on objective 1, etc.). Each objective generates recommendations for good actions in each round, and we can make the decision using our voting rules.
- *Virtual democracy*: In cases where a group of people need to make very many similar 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.
- *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 influence is proportional to their vote strength.
- *Combining generative models*: More speculatively, one could use these methods to “mix” different generative models such as large language models that produce a sequence of tokens one by one. Recently introduced models such as chatGPT, Bing Chat, Bard, and others have different styles, personalities, and strengths, and it is common for users to query several models at once. Instead, one could let users specify a desired mixture of these models, and produce the output sequence by letting the models vote on the next token. On a related note, Bakker et al. [2022] have recently used LLMs to generate consensus statements that maximize expected approval for a group of people with diverse opinions. Given that they collect data for the group over several issues, our methods can be used in such cases to produce proportional outcomes.

	Best known voting rule	PJR	Strong PJR	EJR	Strong EJR
Online	Sequential Phragmén	✓	✓	✗*	✗*
Semi-online	Method of Equal Shares	✓	✗	✓	✗
Offline	PAV	✓	✓	✓	✓

Table 1: Our results. For each setting (online, semi-online (where the rule 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 EJR (or Strong EJR) would resolve an open problem in the theory of multi-winner approval voting.

Our results We build on the work of [Bulteau et al. \[2021\]](#) who defined axioms that guarantee group representation. We consider two of them, which we will call PJR (Proportional Justified Representation) and Strong PJR. [Bulteau et al. \[2021\]](#) gave existence proofs showing that for every possible voting input, there exists an output satisfying both axioms, though their proofs do not suggest a natural voting rule or even a polynomial-time algorithm. In fact, [Bulteau et al. \[2021\]](#) conjecture that no polynomial-time algorithm can find a PJR outcome. We show that the picture is more positive: an attractive polynomial-time voting rule building on ideas of [Phragmén \[1894\]](#) satisfies Strong PJR ([Section 4.1](#)). This rule is fully online, making decisions round by round.

We then define an axiom called EJR (Extended Justified Representation), which provides better group representation guarantees than PJR. We define a simple voting rule (a variant of the “Method of Equal Shares”) that satisfies EJR which 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, the Method of Equal Shares does not satisfy Strong PJR. A rule that maximizes a carefully chosen objective function proposed by [Thiele \[1895\]](#) known as Proportional Approval Voting (PAV) does satisfy EJR and Strong PJR, and even satisfies a new axiom that we define called Strong EJR ([Section 4.3](#)). While the PAV optimization problem is NP-hard, we prove that a polynomial-time local search variant of the rule also satisfies Strong EJR. However, PAV is an offline rule and needs to know full preference information for all the rounds in advance. We do not know an online rule satisfying EJR or Strong EJR, and we show (in the appendix) that finding such a rule would resolve a major open problem in the literature on multi-winner approval voting [[Lackner and Skowron, 2023](#)].

We also consider several different ways of defining stronger axioms than (Strong) PJR and EJR, which either give groups better utility guarantees or give guarantees to more groups. However, by giving counterexamples, we show that these strengthenings are not always satisfiable ([Section 5](#)).

We close with several simulations that compare our rules to existing rules from the literature (which do not give group fairness guarantees). These simulations are based on synthetic data (where we randomly generate approvals in such a way that there are a small number of distinct cohesive voter groups) and on U.S. political election data ([Sections 6.1](#) and [6.2](#)). We also run an experiment for the *virtual democracy* application. Following [Noothigattu et al. \[2018\]](#), we use the *moral machine* dataset [Awad et al. \[2018\]](#), containing several million ethical judgments by respondents from around the world about how a self-driving car should behave in trolley-problem-style dilemmas. We partition the responses according to the country of the respondent. For each country, we learn a preference model predicting the judgments of respondents from the country. We treat each country as a ‘voter’. We then sample new decision situations and query the models for the ethically best decisions in that situation, and aggregate the responses using our voting rule. We compare the aggregated response to an alternative approach where a single model is trained based on responses from all countries (with each country contributing the same number of data points for balance). Based on our metrics, we find that the aggregation approach leads to a fairer outcome, in the sense that agreement with the decision is more equally distributed across countries, compared to the decision made by the combined model ([Section 6.3](#)).

2. Related Work

Perpetual Voting Our work is most closely related to perpetual voting [Lackner, 2020] which concerns online sequential decision making based on approval votes. In this paper, we also consider offline rules. Lackner [2020] focussed 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. Lackner and Maly [2023] also consider “upper quota” axioms, which require that groups are not overrepresented in the outcome; we only consider axioms that prevent underrepresentation.

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

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 [Bogomolnaia et al., 2005, Fain et al., 2016, Aziz et al., 2020, Brandl et al., 2021]. 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 et al., 2017, Sánchez-Fernández et al., 2017, Peters and Skowron, 2020, Brill and Peters, 2023]. Voting rules proposed long ago by Thiele [1895] and Phragmén [1894] do very well with respect to these axioms [Janson, 2016, Brill et al., 2023a]. A more recent rule proposed by Peters and Skowron [2020] known as the Method of Equal Shares combines some of the advantages of the Phragmén and Thiele rules. Skowron 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 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 et al. [2022] carry over to our model, notably that Sequential Phragmén fails 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 (where preferences on one issue may be conditional on the decision in other issues), 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].

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 *rounds*, where T is the *time horizon*. In each round $j \in R$, we are given a set of *alternatives* C_j . Each voter $i \in N$ *approves* some subset A_j^i of C_j . Thus, we have a sequence of T decision rounds together with a sequence $C = (C_j)_j$ of alternative sets and a collection $A = (A_j^i)_{i \in N, j}$ of approval sets. We call the quadruple (N, R, C, A) 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 decisions sequence $D = f(N, R, (C_j)_j, (A_j^i)_{i \in N, j})$. We call a decision rule *semi-online* if its decision in round $j \in R$ only depends on the 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 decision in round j does not depend on 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 strongest properties of (Strong) EJR and (Strong) PJR. For completeness, we also mention two other properties introduced by Lackner and Maly [2023]: we discuss *lower quota for closed groups* in Appendix A.1 and *perpetual priceability* in Appendix A.2. In Appendix A.3, we briefly discuss Pareto efficiency.

Our first two main axioms were introduced by Bulteau et al. [2021].¹ Let us say that 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 $\bigcap_{i \in S} A_j^i \neq \emptyset$. PJR (short for Proportional Justified Representation) requires that a group of an α fraction of the voters, if they agree in every round, must be “happy” with at least an α fraction of the decisions. We define this to mean that in $\lfloor \alpha T \rfloor$ many rounds, at least one member of S approves the decision (but this member of S can be different in each round).

Definition 3.1 (PJR). A decision sequence D satisfies *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 at least one voter in S (i.e., $d_j \in \bigcup_{i \in S} A_j^i$).

PJR only provides guarantees for groups that agree in *all* rounds. Strong 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 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). Note that Strong PJR implies PJR (take $k = T$).

Definition 3.2 (Strong PJR). A decision sequence D satisfies *Strong PJR* if for every $\ell \in \mathbb{N}$ and every group of voters $S \subseteq N$ that agrees in k rounds and has size $|S| \geq \ell \cdot \frac{n}{k}$, 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., $d_j \in \bigcup_{i \in S} A_j^i$).

An equivalent way of stating this 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 $\lfloor \alpha \cdot \beta \rfloor$ fraction of the decisions. For example, consider a group $S \subseteq N$ of voters with fixed size $\ell \cdot \frac{n}{T}$, and let us ask what Strong 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 $\frac{T}{2}$ 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.

¹We use different names for these axioms. Bulteau et al. [2021] use the term “all periods intersection PJR” for what we just call PJR, and the term “some periods intersection PJR” for what we call Strong PJR.

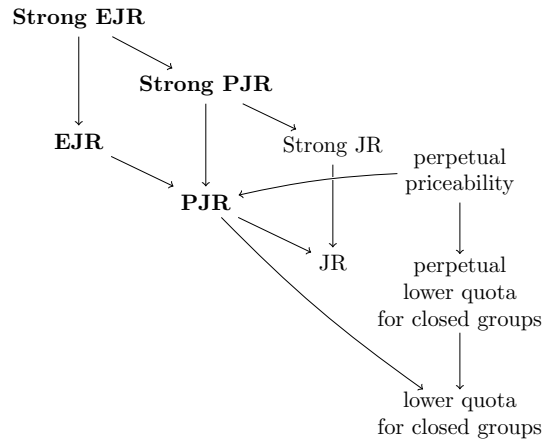


Figure 1: Implications between axioms. We focus on the axioms in bold.

Following the literature on multi-winner voting, we say that a decision sequence satisfies (Strong) *JR* (Justified Representation) if it satisfies (Strong) *PJR* restricted to $\ell = 1$. Several of our counterexamples will apply even to *JR* and to Strong *JR*.

To illustrate the axioms, consider the instance shown in [Figure 2](#), with $T = 8$ rounds and 4 voters. The first two voters form a group $S = \{1, 2\}$ that agrees in all rounds (because they both always approve a). Thus, for an outcome to satisfy *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 $|S'| \geq \ell' \cdot \frac{4}{6}$, Strong *PJR* requires that in at least 3 rounds, the outcome is either d , e , or f .

A weakness of *PJR* and Strong *PJR* is in how they define S being “happy” with a decision (“at least one member of S approves the decision”). This definition could be satisfied by a decision sequence that gives each member of S a utility that is much lower than ℓ [[Peters and Skowron, 2020](#), Sec. 4.2]. In the example of [Figure 2](#), the decision sequence (b, b, c, c, d, d, e, f) satisfies Strong *PJR*, but the first two voters each only approve the decision in 2 rounds, instead of in 4 rounds.

Following [Aziz et al. \[2017\]](#), we can fix this issue by defining the axioms *EJR* (short for Extended Justified Representation) and Strong *EJR*. These axioms 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 Strong *EJR*.

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

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

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.

3.3. Methods

We now define three decision rules that are natural analogues of rules that were first proposed in the context of multi-winner elections.

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. Its main idea is that each decision provides a value of 1, and this value (usually referred to as *load*) will be distributed to the 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 $c \in C_j$:

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

This value 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 blue 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

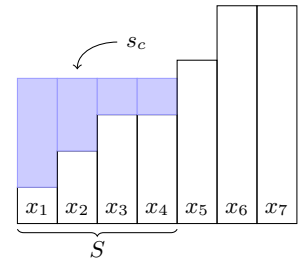


Figure 3: Load distribution

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 }

Figure 2: Example illustrating our axioms.

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.

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.² Clearly, Sequential Phragmén is an online rule that requires no information about future rounds. It can be computed in polynomial time [Lackner and Maly, 2023].

3.3.2. Method of Equal Shares (MES)

MES is a recently introduced voting rule for multi-winner voting [Peters and Skowron, 2020] and participatory budgeting [Peters 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. MES makes decisions round by round. Similar to the concept of load in Sequential Phragmén, each decision costs $p = \frac{n}{T}$ units which must be distributed among voters who approve the alternative. MES works by subtracting this amount from an initial budget assigned to each voter. Formally, let $b^i = 1$ be the initial budget of voter i . For $\rho \geq 0$, we say that an alternative $c \in C_j$ is ρ -affordable if $\sum_{i \in N: c \in A_j^i} \min(b^i, \rho) \geq p$ (so approvers of c can pay the amount p with their remaining budgets even if no one pays 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. Then for each voter i approving d_j , i 's remaining budget b^i is set to $\max(0, b^i - \rho)$. If in some round $j \in R$, no alternative can be bought for any ρ , i.e., $\sum_{i \in N: c \in A_j^i} b^i < p$ for all alternatives $c \in C_j$, MES *terminates prematurely*. For purposes of our theoretical results, it does not matter how the decisions for remaining rounds are made. All these results (positive and negative) hold for every way of completing the partial decision sequence that is output by MES.

For actual implementations of MES, one can choose between several “completion rules”, which are familiar from work on multi-winner voting and participatory budgeting. In our experiments, we followed the “ ε -completion” strategy introduced by [Peters et al., 2021, Sec. 3.4]. In a round $j \in R$ where no alternative is affordable according to the definition of MES, we will choose an alternative $a_j \in C_j$ and make all voters who approve a_j spend all their remaining money to cover as much cost p of the round as possible. The remaining part of the cost is covered by the voters $i \in N$ who don't approve a so that the maximal payment of such a voter, γ , is minimized. We select the alternative $a_j \in C_j$ minimizing γ and update voters' individual budgets.

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 go through for Offline MES, so it also satisfies EJR but fails Strong PJR.

3.3.3. Proportional Approval Voting (PAV)

Thiele [1895] proposed an alternative method to Phragmén's, which is based on optimizing an objective function. We adapt this method to our context in an offline way. PAV 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 . This harmonic objective function is the unique additive objective that leads to a proportional rule [Aziz et al., 2017]. Finding the optimum decision sequence for PAV is NP-hard [Brill et al., 2022, Thm. 5.1], just like in the multi-winner setting [Aziz et al., 2015], but we will show that a polynomial-time local search variant [Aziz et al., 2018] satisfies the same axioms as PAV.³

²In the multi-winner setting, it is not necessary to define s_c using a minimization over S , and we can just take S to be the set of *all* voters approving c . This is because in multi-winner voting, it cannot happen that an approver of c already has more load than s_c , because in that case the rule would have chosen c in an earlier iteration [Brill et al., 2023a, Lemma 4.5]. This argument does not work in our setting because c may not have been available in a prior round. An alternative definition of Phragmén's method in multi-winner voting is based on virtual bank accounts that are continuously filled [Janson, 2016, Peters and Skowron, 2020]. This definition does not easily adapt to our setting due to a similar problem: if in some round no good alternative exists, then the bank accounts will fill up too much, leading to ties in later rounds.

³There also exists a sequential version of PAV that could be applied to our setting to give an online rule. Based on

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 PJR (their existence proof were based on an exponential algorithm). Our results show that in fact, 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 Strong PJR.*

Proof. For a contradiction, suppose that $S \subseteq N$ witnesses a violation of Strong PJR: it agrees in the set of 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. 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 $j \in R^*$ a *bad round* if the decision d_j is not approved by any voter in S . Now let $i \in N \setminus S$ be a voter not in S , and suppose i gets assigned some load during at least one bad round. Consider the point just after the last bad round where i is assigned some load. At this point we must have $x^i \leq \frac{\ell}{|S|}$ because otherwise $s_{d_j} > \frac{\ell}{|S|}$, contradicting our previous claim. Thus, at most $\frac{\ell}{|S|}$ load was assigned to i during bad rounds (i.e., adding up the load additions to i over all bad rounds). 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, we can now bound the total load that gets assigned in bad rounds, by summing over all $i \in N \setminus S$; it 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 et al., 2023a], so it is unsurprising that it also fails EJR in our setting. While it is possible to adapt the example of Brill et al. [2023a] to our setting, we use a somewhat simpler example.⁴

Theorem 4.2. *Sequential Phragmén fails 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 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

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

Figure 4: Instance where Sequential Phragmén fails EJR. Note that the approval sets are the same in all rounds.

simulations, Page et al. [2020] conjecture that this rule would satisfy at least a weak version of PJR. However, a result of Brill et al. [2022, Footnote 4] implies that the sequential version fails this property.

⁴This 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.

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 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 EJR is violated. Hence, Sequential Phragmén fails EJR. \square

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

We do not know an online rule that satisfies EJR, and in [Appendix B.3](#) we show that if there exists such a rule, then there also exists a multi-winner voting rule satisfying EJR and the axiom of *committee monotonicity*. Whether such a multi-winner voting rule exists is a well-known open problem [[Lackner and Skowron, 2023](#)].

However, we can satisfy EJR if we use a little extra information: the total number T of rounds. MES is such a rule: it needs to know T to determine the price $p = n/T$ of deciding a round (or in an equivalent formulation, to decide the initial budget assigned to each voter), but is otherwise online. For example, to make the decision in round 1, MES does not need to know what alternatives will be available in future rounds and what the approval sets will be. We show that MES satisfies EJR.

Theorem 4.3. *MES satisfies 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 EJR, where $|S| \geq \ell \cdot p$ but every $i \in S$ approves the decision of at most $\ell - 1$ rounds.

Suppose MES stops prematurely without making a decision in round $j \in R$. At that point, some voter $i \in S$ has 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 . Otherwise, if it does not stop prematurely, MES makes a decision for all T rounds, in which case all available money is spent (since $T \cdot p = n$), and hence we can again find $i \in S$ with remaining budget 0, i.e., less than $p/|S|$.

Thus, 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 paid strictly more than $1/\ell$ for the decision; let $j \in R$ be the first such round with decision $d_j \in C_j$. Just before paying for d_j , every voter in S had at least $1/\ell$ budget left as they each so far paid at most $1/\ell$ for at most $\ell - 1$ alternatives. But note that the alternative $c \in C_j$ on which S agrees is therefore $1/\ell$ -affordable because

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

while d_j is not $1/\ell$ -affordable (since i 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 Strong PJR and thus also Strong EJR. This is 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 on early rounds where MES greedily maximizes efficiency, and then MES cannot satisfy the fairness requirements in subsequent rounds where there may not be enough agreement between voters.

Theorem 4.4. *MES fails Strong PJR.*

Proof. Consider the instance shown in [Figure 5](#), with $T = 6$ rounds and $n = 3$ voters. Here, voter 1 approves a in 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 . The remaining budget of voters 2 and 3 is set to 0.75. By the same computation, b is selected in rounds

Round	1 & 2 & 3	4 & 5 & 6
Voter 1	{ a }	\emptyset
Voter 2	{ b }	{ b }
Voter 3	{ b }	{ b }

Figure 5: Example illustrating how MES fails Strong PJR (and Strong JR). Note the empty approval set of voter 1 in rounds 4, 5, 6.

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 an individual 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, Strong 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 Strong PJR. \square

Remark 4.5. The proof of [Theorem 4.4](#) used a requirement of $\ell = 1$ rounds, so it also shows that MES fails Strong JR.

In the proof of [Theorem 4.4](#), MES terminates prematurely. In [Appendix B.1](#), we show that on instances where MES makes a decision for every round, it satisfies Strong PJR, but may fail Strong EJR.

4.3. Offline: Proportional Approval Voting (PAV)

If we make decisions fully offline (i.e., we know the approval sets for all rounds in advance), there is a rule that – as we will see – satisfies all four of our axioms: PAV. However, finding the optimum decision sequence for PAV is NP-hard [[Aziz et al., 2015](#), [Brill et al., 2022](#)]. Following a suggestion of [Aziz et al. \[2018\]](#), we can define a polynomial-time local search variant of PAV which turns out to satisfy the same proportionality axioms as PAV. The variant starts with an arbitrary decision sequence D and keeps changing the decision in some round if this increases the PAV-score of the decision sequence by at least n/T^2 . The process terminates if there is no possible change that leads to a sufficient increase.

Definition 4.6 (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( $D \setminus \{d_j\} \cup \{a_j\}$ )  $\geq$  PAV-score( $D$ ) +  $\frac{n}{T^2}$  do
  | D ←  $D \setminus \{d_j\} \cup \{a_j\}$ 
return D

```

The threshold of n/T^2 in the definition of Local-Search PAV is chosen to imply the proportionality properties of PAV while ensuring termination of the local search in polynomial time.

Proposition 4.7 ([Aziz 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. \square

We can now prove that both PAV and Local-Search PAV satisfy Strong EJR. The proof uses a swapping argument: if the output violates Strong EJR, then in at least 1 round, it is possible to change the decision in a way that increases the PAV objective function (by at least n/T^2). This proof technique has also been used to prove that PAV satisfies EJR in multi-winner voting [[Aziz et al., 2017](#)].

Theorem 4.8. *PAV and Local-Search PAV satisfy Strong 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 Strong 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 Strong 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

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

Note that the change for a particular round may be negative. Summing this change over R^* , 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

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

From the pigeonhole principle, it follows that there exists a round $r \in R^*$ such that $\Delta(a_r, d_r) \geq \frac{1}{k} \cdot \frac{n}{k} = \frac{n}{k^2} \geq \frac{n}{T^2}$ (as $|R^*| = k$ and $k \leq T$). 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

In some respects, the axioms we have been considering are weak. PJR and EJR, for example, only provide guarantees to coalitions that agree in *every* round. Their strong versions are more robust, but

their associated guarantees diminish as a coalition has agreement in fewer rounds. Also, readers with a background in multi-winner voting may have expected PJR and EJR to be formulated differently, in a way that a coalition deserving ℓ decisions only needs to agree in ℓ rounds (instead of all T rounds). One might hope that stronger proportionality axioms of this kind could be defined and still be satisfiable.

However, we can show that several strengthenings are in fact impossible to satisfy in general. Our results follow from a single counterexample construction. This counterexample 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.

Formally, the counterexample construction establishes the following result:

Theorem 5.1. *Let $\varepsilon > 0$. Choose some $k > \lceil \frac{1-\varepsilon}{\varepsilon} \rceil$ 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 $|S| \geq (1 - \varepsilon) \cdot \frac{n}{k}$, yet none of the voters in S approve any of the decisions of D .*

The proof appears in [Appendix B.2](#). From [Theorem 5.1](#), it follows immediately that for all $\varepsilon > 0$, it is impossible to satisfy an additive ε -strengthening of Strong PJR (or even Strong JR):

Corollary 5.2 (ε -Strong PJR need not exist). *Let $\varepsilon > 0$. Then there exists an instance where no decision sequence D satisfies “ ε -Strong 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 $|S| \geq (\ell - \varepsilon) \cdot \frac{n}{k}$, 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., $d_j \in \bigcup_{i \in S} A_j^i$).*

Above, we briefly mentioned the idea of defining PJR by requiring that a group large enough to deserve ℓ decisions only need to agree in ℓ rounds. This corresponds perhaps most closely to the original definition of PJR 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”).

Corollary 5.3 (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 5.1](#) with $\varepsilon = 0.5$, $k = 4$, 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 3 rounds, has size $|S| \geq (1 - \varepsilon) \cdot \frac{n}{k} = 0.125 \cdot n \geq 0.1 \cdot n = 4 \cdot \frac{n}{T}$, 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](#)].

[Do et al. \[2022\]](#) propose a way to relax PJR and related concepts in a multiplicative way. In [Appendix B.2](#), we show that such constant-factor relaxations of ℓ -Agreement PJR (defined in [Corollary 5.3](#)) are also not guaranteed to exist. In [Appendix B.2](#), we also explain how to adapt [Theorem 5.1](#) to show that an online rule cannot satisfy any of these strengthened properties *even on instances where they are satisfiable*.

6. Experiments

To understand the performance of our methods empirically, we run our methods on both synthetic and real-world datasets. In addition to our proposed rules, we also consider two rules proposed by [Lackner \[2020\]](#): *Perpetual Quota* (aims at granting each voter a satisfaction as close as possible to their “quota”) and *Perpetual Consensus* (similar to Sequential Phragmén but strictly enforces an equal distribution of the load incurred).⁵ We chose those rules since they performed well in [Lackner’s \[2020\]](#) experiments. Further,

⁵These rules aim for proportionality, but do not satisfy any of our axioms [[Lackner and Maly, 2023](#)].

we also consider two baselines: *Approval Voting* (chooses the alternative with the highest approval score in each round) and *Round Robin* (in each round j , voter $j \bmod T$ chooses an approved alternative).

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’s utility as the *fraction* of rounds in which the voter approves the decision. Using our previous notation, the normalized utility of voter i in a decision sequence D is U_D^i/T . Based on these normalized utilities, we report three metrics:

- *Average Utility*: We report the average utility of the voters, i.e., the fraction of rounds with an approved decision for the average voter. This measures the efficiency of the selected outcome, and it measures the utilitarian social welfare. Approval Voting maximizes this measure by definition.
- *Utility of 25th Percentile*: We sort the vector of utilities and report its 25th percentile. This measure is inspired by the idea of egalitarian welfare (which is the utility of the worst-off agent), but in our experiments, the obtained egalitarian welfare was often very low or even 0. Hence, we report the utility of the 25th percentile which better distinguishes the rules.
- *Gini Coefficient*: This metric quantifies the level of inequality in the voter utilities. A lower value corresponds to a more equal utility distribution (with the lowest possible Gini coefficient of 0 being obtained in case every voter has the same utility). High Gini coefficients are obtained when a few voters are happy with many decisions, and the other voters are happy with few decisions. Note that on its own, a low Gini coefficient need not signify a desirable decision, since it can also be obtained by a decision that gives every voter the same but very low utility.

We implemented all the rules in Python 3.8, based 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.⁶ For PAV, we used a standard integer linear-program (ILP) encoding [Peters and Lackner, 2020], using Gurobi 9. For MES, our implementation builds upon the `pabutools` package.⁷

6.1. Synthetic Data

To compare the performance of the rules, we follow a similar setup to the one used by Lackner [2020]. The data generation is based sampling locations of voters and alternatives in two-dimensional Euclidean space, with voters approving nearby alternatives. 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 a two-dimensional Euclidean space. 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 $\sigma = 0.1$.
- *Balanced and Nearby*: The voters are split into 2 disjoint groups S_1 and S_2 60% of voters in S_1 and 40% of voters in S_2 . In this distribution, the groups are 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 $\mathcal{N}(0.25, 0)$.

⁶<https://github.com/martinlackner/perpetual>

⁷<https://github.com/Grzesiek3713/pabutools>

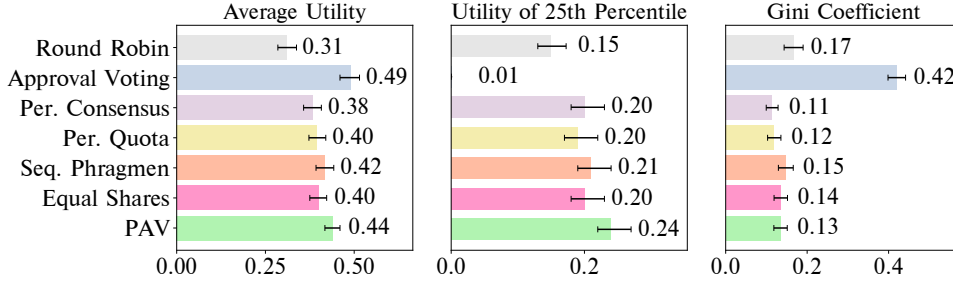


Figure 6: Performance of the different rules on the *Restricted* distribution for $n = 20$ voters, $T = 20$ rounds, 20 alternatives in each round, with $f = 1.5$. The length of the bar represents the median across all the trials while the error bars represent the 25th and 75th percentile with the numeric text after the bar showing the mean value.

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.

For each of the distributions mentioned above, we considered the following 4 cases:

- $n = 20$ voters, $T = 20$ rounds and each round having 20 alternatives, with $f = 1.5$. This led to an average approval set size of 2.11.
- $n = 20$ 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.

Each experiment was repeated for 1000 trials. We report our results for the various distributions in [Appendix C.1](#) (Figures 12, 13, 14 and 15). Each figure contains multiple subfigures to show the plots for different parameter values. In each subfigure, bar plots are shown for each 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. As an example, in [Figure 6](#), we show the results for the *Restricted* distribution and the first bundle of parameter values.

Across the various distributions, we find some prominent patterns for each metric.

- *Overall Satisfaction*: In regard to this metric, a clear separation emerges between the voting methods across all the distributions. Approval Voting demonstrates the highest performance, which is expected (as it naturally maximizes this objective function over all decision sequences). 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.
- *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 Balanced and Nearby distribution, 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.

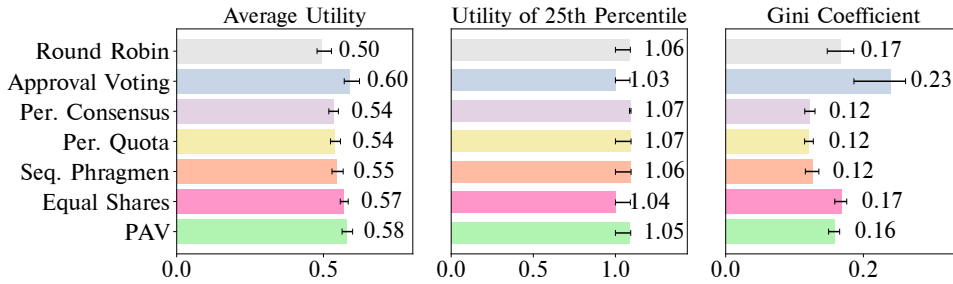
- *Bottom 25%*: Here, we find that Approval Voting consistently performed poorly, often resulting in almost no decisions being approved by the bottom 25% of the voters, leading to a score of close to 0 for this metric under the Restricted and the Many Groups distributions. Round Robin performed better than Approval Voting but still worse than the other rules.

Among those other rules, PAV generally performed the best. However, in the *Balanced and Nearby* distribution, Perpetual Quota occasionally outperformed PAV. Both Sequential Phragmén and MES showed similar performance to Perpetual Quota and Consensus, with Sequential Phragmén having slightly better results sometimes.

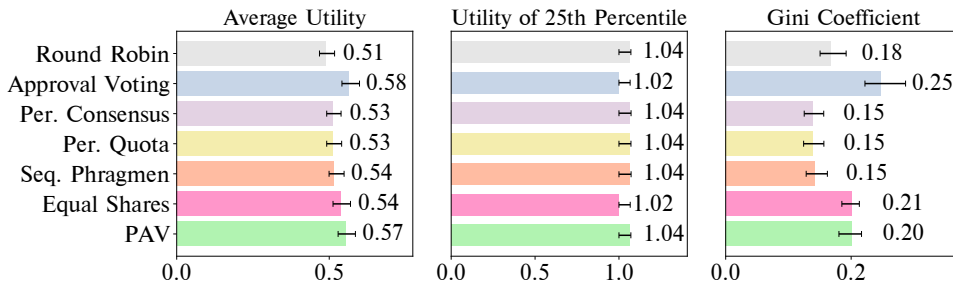
Notably, in the *Balanced and Nearby* distribution, Perpetual Quota and Consensus consistently outperformed MES and Sequential Phragmén. This could indicate that these methods perform particularly well when the voters have similar preferences.

6.2. Political Data

In addition to synthetic data, we evaluated the rules on datasets from U.S. political elections, in particular from the 2022 General Election. In these elections, voters are asked to elect candidates to various federal, state, and local political offices (for example, U.S. president, senators, representatives, state boards, commissioners, judges) as well as to express opinions on yes/no ballot initiatives (for example, constitutional amendments). In response to heightened public interest in election integrity since 2020, several jurisdictions (notably most counties in Colorado) have published (anonymized) *Cast Vote Records (CVR)* data that show, for each voter, their votes on all these issues simultaneously. This makes them interesting for our purposes. Notably, in almost all cases, voters may only vote for a single alternative for each issue (rather than using approval),⁸ but our rules disagree even on such simple instances, so an empirical evaluation is interesting. The current voting rule used for these elections is formally equivalent to what we have called Approval Voting (i.e., in each round select the alternative with the highest vote count), one of the baseline rules included in our simulations.



(a) Results for the 2022 General Elections in California.



(b) Results for the 2022 General Elections in Colorado.

Figure 7: Performance of the different rules on the 2022 General Elections in the state of California and Colorado. The height of the bar represents the median across all the instances; while the error bars represent the 25th and 75th percentile with the numeric text after the bar showing the mean value.

⁸There are some exceptions. In particular, 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.

In the dataset for Colorado, we collected 169 instances from 14 different counties. Each instance corresponds to all the voters that were issued a particular ballot (usually, one instance per district of the county, since the candidates and offices differ between districts). On average, these instances contain 42 issues and 2 255 voters. For California, we collected 41 instances from Shasta county 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.

We show the results of the simulations on this data in [Figure 7](#) with one figure for each of the two states, using the same plot format that we used for results on synthetic data. For each metric, the bar chart shows the median value across the instances from the relevant state.

We can see that Approval Voting (i.e., the method that is actually used to determine winners) has a high Gini coefficient compared to all the other rules. Sequential Phragmén, Perpetual Quota and Consensus perform the best in terms of the Gini coefficient. Remarkably, the average satisfaction of the proportional rules is very close to the optimum average satisfaction achieved by Approval Voting. This is particularly striking for PAV, which achieves an almost optimal overall satisfaction while simultaneously ensuring a more equal utility distribution as measured by Gini coefficient and bottom 25% satisfaction.

6.3. Learning Preferences from the Moral Machine

In the introduction, we discussed the idea of virtual democracy [[Noothigattu et al., 2018](#), [Freedman et al., 2020](#), [Mohsin et al., 2021](#)], which automates decision making by learning models of preferences of individual users (or defined demographic groups of users), and using predicted preferences as inputs to the voting rule. This can be particularly useful when preference elicitation costs are high, or decisions need to be made in real time. Proportional aggregation rules like ours can be useful in this context to ensure that each user (or demographic) gets a fair share in the decision making process.

We are particularly interested in whether proportionally aggregating multiple models representing different demographics can lead to better outcomes than a single learnt model making decisions. We present preliminary empirical evidence in this section. Following the work of [Noothigattu et al. \[2018\]](#), and in the spirit of a prototype application, we consider virtual democracy as applied to the moral machine experiment [[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 are asked to make decisions by pairwise comparisons of 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 groups of people can be seen as alternatives that can be described by a structured feature vector.⁹ Several million responses are available in a public dataset.

[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. [Kim et al. \[2018\]](#) show that learning a single model for respondents from a country leads to reasonable accuracy (perhaps due to cultural similarities).

We limit ourselves to 197 countries for which the dataset contains over 100 samples. We use the Plackett-Luce (PL) model [[Plackett, 1975](#), [Duncan, 1959](#)] which is a *random utility model* appropriate for social choice preference learning [[Azari Soufiani et al., 2012](#)]. Just like in the approach of [Noothigattu et al. \[2018\]](#), we learn for each country a vector w that assigns a weight to each feature. Taking the dot product between w and the feature vector of an alternative gives rise to an estimated latent utility of the agent (i.e., the country) for the alternative. Informally, by training the PL model, we find a vector w which captures the behavior of the respondents of the countries, by (with maximum likelihood) assigning a higher utility to the alternative selected in each pairwise comparison.

To speed up training, for each country we train a model by sampling up to 4000 pairwise comparisons from respondents of that country. We also train a *combined model* trained on respondents from all countries. We will later compare the choices of the combined model with the output of our voting rules which aggregates the predicted responses of the country models. The combined model is trained by sampling 100 samples from each country (to ensure balanced representation of all countries in the training data). The mean accuracy on unseen pairwise comparisons of the 197 country-wise models is 70.06% on samples from the country and for the single combined model it is 71.10% on global samples; this performance is similar to results reported by [Kim et al. \[2018\]](#).

⁹Features include the number of men, women, pregnant women, criminals, children, animals, and other characters, that will be injured in the scenario, as well as features indicating (1) whether these persons are the passengers of the vehicle or pedestrians, (2) whether the pedestrians are currently crossing the road despite a red traffic signal, (3) whether it would be necessary to intervene in the car’s travel path to save the people.

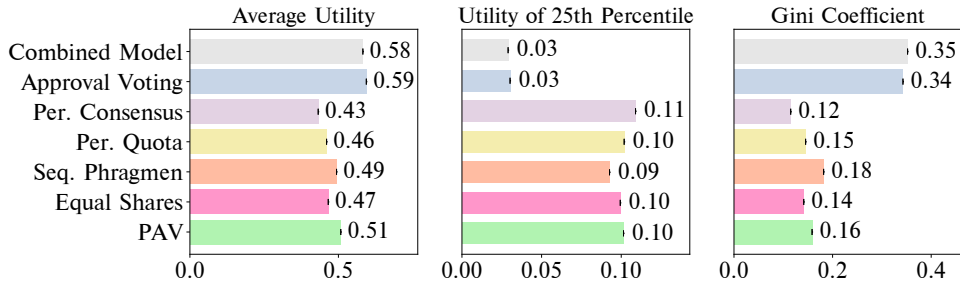


Figure 8: Performance of the different rules on the Moral Machine Dataset, based on high-disagreement alternatives, with parameter setting (d): 100 rounds with each round having 200 alternatives and each voter approving the top 2 alternatives.

For evaluation, we consider four different parameter settings (a)-(d). We produce 100 decision rounds together with alternatives for each round. From this, 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 10 alternatives. In each round, we produce either 50 alternatives ((a)-(c)) or 200 alternatives ((d)), in one of the following ways:

- *Random.* We create a large pool of 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 set of all 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.

We then 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 our results in [Appendix C.2](#), and in [Figure 8](#) we show as an example the results with high-disagreement alternatives and parameter settings (d), chosen to show the largest separations between rules. For parameter settings with large approval sets (especially (a) but also (c)), the conclusions are less clear, which we attribute to high overall agreement among different voters.

Looking at [Figure 8](#), it is striking how Approval Voting and the Combined Model attain almost identical values on each metric, and how these are quite different from the metrics obtained by the proportional rules. Indeed, in the experiment, Approval Voting and the Combined Model choose the same decision in 84% of rounds, but both agree with the 5 proportional rules less frequently. (The Combined Model agrees with proportional rules in 54%–70% of rounds, and Approval Voting agrees in 59%–77% of rounds.) Notably, the 5 proportional rules all feature a much smaller Gini coefficient (indicating a more equal distribution of satisfaction across countries) and a higher satisfaction at the 25th percentile. The similar performance and behavior between Approval Voting 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 identical numbers of samples from each country, it will “merge” their views roughly proportionally.

The conclusions from [Figure 8](#) also hold for high-disagreement alternatives with parameter setting (b), and to a lesser extent for (c), where the performance of the proportional rules is more similar to Approval Voting and the Combined Model. For random (non-high-disagreement) alternatives, we also see the same separation between proportional and non-proportional methods with respect to the Gini coefficient (but not 25th percentile satisfaction), when using parameter settings (b) and (d).

We see the results from this small experiment as a potential starting point for a larger research program that studies how a preference model trained on preferences of diverse agents makes decisions, when compared to aggregation rules that explicitly take into account each individual’s preferences.

7. Discussion and Future Work

We have shown that simple polynomial-time voting rules can provide strong proportionality guarantees in the problem of sequential decision making based on users’ preferences. We hope our work inspires applications of voting rules in sequential decision making.

Open problems We leave some theoretical open problems for future work, notably whether an online rule can satisfy EJR or Strong EJR – a negative result may be easier to find in this setting than for multi-winner voting. Also open is whether a semi-online rule can satisfy Strong EJR. More conceptually, are there are stronger versions of 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 [Peters et al., 2021]) or the core? Can other methods from multi-winner voting also be adapted to our setting, like the Greedy Cohesive Rule [Peters et al., 2021] or the Greedy Justified Candidate Rule [Brill and Peters, 2023]?¹⁰

Extensions Our model can be extended to make it compatible with more real-world applications. Examples include allowing voters to specify utilities or to rank alternatives instead of approvals [Peters et al., 2021], to allow weighting issues by importance [Page et al., 2020], and to allow for dependencies between issues [Brill et al., 2023b]. Also, one could add a budget constraint with alternatives having costs, making the model applicable to participatory budgeting [Peters et al., 2021].

Strategic issues Peters [2018] proved an impossibility theorem showing that no proportional multi-winner voting rule can be *strategyproof*, which means that in the worst case, voters may be able to get a better outcome by misrepresenting their preferences. A special case of multi-winner approval voting is “approval-based apportionment” [Brill et al., 2022, Airiau et al., 2023], and as we have mentioned before, this special case is also a special case of our model (when the set of alternatives and voter preferences are the same in each round). Since the impossibility theorem holds even for this special case [Delemazure et al., 2023, Lackner et al., 2023], it follows that no proportional rule in our setting can be strategyproof. Strategyproofness has also been considered in multi-issue decision making without proportionality objectives [Amanatidis et al., 2015, Barrot et al., 2017].

Acknowledgments

We thank Jérôme Lang for his support and advice throughout this research project, and for proposing useful 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).

8. Bibliography

- Stéphane Airiau, Haris Aziz, Ioannis Caragiannis, Justin Kruger, Jérôme Lang, and Dominik Peters. Portioning using ordinal preferences: Fairness and efficiency. *Artificial Intelligence*, 314:103809, 2023. [→ p. 18]
- Georgios Amanatidis, Nathanaël Barrot, Jérôme Lang, Evangelos Markakis, and Bernard Ries. 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)*, pages 715–723, 2015. [→ p. 4, 18]
- Edmond Awad, Sohan Dsouza, Richard Kim, Jonathan Schulz, Joseph Henrich, Azim Shariff, Jean-François Bonnefon, and Iyad Rahwan. The moral machine experiment. *Nature*, 563(7729):59–64, 2018. [→ p. 3, 16]

¹⁰These greedy rules work by repeatedly searching for a coalition of voters who currently witness a violation of a proportionality axiom such as EJR and then adding candidates to the committee to avoid this violation. The difficulty in adapting this to our setting is that a witnessing coalition could be made happy via decisions in any of their (agreement) rounds. With a naïve choice of which rounds to decide on in each step (such as arbitrary tie-breaking), these rules do not satisfy axioms like Strong PJR.

- Hossein Azari Soufiani, David C. Parkes, and Lirong Xia. Random utility theory for social choice. *Advances in Neural Information Processing Systems*, 25:126–134, 2012. [→ p. 16]
- Haris Aziz, Serge Gaspers, Joachim Gudmundsson, Simon Mackenzie, Nicholas Mattei, and Toby Walsh. Computational aspects of multi-winner approval voting. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 107–115, 2015. [→ p. 7, 10]
- Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 48(2):461–485, 2017. [→ p. 4, 6, 7, 10, 28]
- Haris Aziz, Edith Elkind, Shenwei Huang, Martin Lackner, Luis Sánchez-Fernández, and Piotr Skowron. On the complexity of extended and proportional justified representation. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, pages 902–909, 2018. [→ p. 7, 10]
- Haris Aziz, Anna Bogomolnaia, and Hervé Moulin. Fair mixing: the case of dichotomous preferences. *ACM Transactions on Economics and Computation*, 8(4):18:1–18:27, 2020. [→ p. 4]
- Michiel A. Bakker, Martin J. Chadwick, Hannah Sheahan, Michael Henry Tessler, Lucy Campbell-Gillingham, Jan Balaguer, Nat McAleese, Amelia Glaese, John Aslanides, Matt M. Botvinick, and Christopher Summerfield. Fine-tuning language models to find agreement among humans with diverse preferences. *Advances in Neural Information Processing Systems*, 35:38176–38189, 2022. [→ p. 2]
- Siddhartha Banerjee, Vasilis Gkatzelis, Safwan Hossain, Billy Jin, Evi Micha, and Nisarg Shah. Proportionally fair online allocation of public goods with predictions. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, 2023. [→ p. 4]
- Nathanaël Barrot and Jérôme Lang. Conditional and sequential approval voting on combinatorial domains. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 88–94, 2016. [→ p. 4]
- Nathanaël Barrot, Jérôme Lang, and Makoto Yokoo. 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)*, pages 597–605, 2017. [→ p. 18]
- Anna Bogomolnaia, Hervé Moulin, and R. Stong. Collective choice under dichotomous preferences. *Journal of Economic Theory*, 122(2):165–184, 2005. [→ p. 4]
- Florian Brandl, Felix Brandt, Dominik Peters, and C. Stricker. Distribution rules under dichotomous preferences: Two out of three ain’t bad. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, pages 158–179, 2021. [→ p. 4]
- Markus Brill and Jannik Peters. Robust and verifiable proportionality axioms for multiwinner voting. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*, 2023. [→ p. 4, 18]
- Markus Brill, Paul Gözl, Dominik Peters, Ulrike Schmidt-Kraepelin, and Kai Wilker. Approval-based apportionment. *Mathematical Programming*, 2022. [→ p. 4, 7, 8, 10, 18, 28]
- Markus Brill, Rupert Freeman, Svante Janson, and Martin Lackner. Phragmén’s voting methods and justified representation. *Mathematical Programming*, 2023a. [→ p. 4, 7, 8]
- Markus Brill, Evangelos Markakis, Georgios Papatotiropoulos, and Jannik Peters. Proportionality guarantees in elections with interdependent issues. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, 2023b. [→ p. 4, 18]
- Laurent Bulteau, Noam Hazon, Rutvik Page, Ariel Rosenfeld, and Nimrod Talmon. Justified representation for perpetual voting. *IEEE Access*, 9:96598–96612, 2021. [→ p. 3, 4, 5, 8]
- Vincent Conitzer, Rupert Freeman, and Nisarg Shah. Fair public decision making. In *Proceedings of the 2017 ACM Conference on Economics and Computation (EC)*, pages 629–646, 2017. [→ p. 4]

- Théo Delemazure, Tom Demeulemeester, Manuel Eberl, Jonas Israel, and Patrick Lederer. Strategyproofness and proportionality in party-approval multiwinner elections. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, 2023. [→ p. 18]
- Virginie Do, Matthieu Hervouin, Jérôme Lang, and Piotr Skowron. Online approval committee elections. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 251–257, 2022. [→ p. 12, 27]
- Luce R. Duncan. *Individual choice behavior: A theoretical analysis*. New York: Wiley, 1959. [→ p. 16]
- B. Fain, A. Goel, and K. Munagala. The core of the participatory budgeting problem. In *Proceedings of the 12th Conference on Web and Internet Economics (WINE)*, pages 384–399, 2016. [→ p. 4]
- Michael Feffer, Hoda Heidari, and Zachary Lipton. Moral machine or tyranny of the majority? In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, 2023. [→ p. 2]
- Rachel Freedman, Jana Schaich Borg, Walter Sinnott-Armstrong, John P. Dickerson, and Vincent Conitzer. Adapting a kidney exchange algorithm to align with human values. *Artificial Intelligence*, 283:103261, 2020. [→ p. 2, 16]
- Rupert Freeman, Seyed Majid Zahedi, and Vincent Conitzer. Fair social choice in dynamic settings. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 4580–4587, 2017. [→ p. 4]
- Rupert Freeman, Anson Kahng, and David M. Pennock. Proportionality in approval-based elections with a variable number of winners. In *Proceedings of the 29th International Joint Conferences on Artificial Intelligence (IJCAI)*, pages 132–138, 2020. [→ p. 4, 12]
- Dylan Hadfield-Menell, Stuart J. Russell, Pieter Abbeel, and Anca D. Dragan. Cooperative inverse reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 29, pages 3909–3917, 2016. [→ p. 2]
- Svante Janson. Phragmén’s and Thiele’s election methods. *arXiv:1611.08826*, 2016. [→ p. 4, 7]
- Richard Kim, Max Kleiman-Weiner, Andrés Abeliuk, Edmond Awad, Sohan Dsouza, Joshua B. Tenenbaum, and Iyad Rahwan. A computational model of commonsense moral decision making. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pages 197–203, 2018. [→ p. 16]
- Martin Lackner. Perpetual voting: Fairness in long-term decision making. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, pages 2103–2110, 2020. [→ p. 2, 4, 12, 13]
- Martin Lackner and Jan Maly. Proportional decisions in perpetual voting. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, 2023. [→ p. 4, 5, 6, 7, 12, 22, 23, 24]
- Martin Lackner and Piotr Skowron. *Multi-Winner Voting with Approval Preferences*. Springer, 2023. [→ p. 2, 3, 4, 9, 28]
- Martin Lackner, Jan Maly, and Oliviero Nardi. Free-riding in multi-issue decisions. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2023. [→ p. 18]
- Jérôme Lang and Lirong Xia. Voting in combinatorial domains. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 9. Cambridge University Press, 2016. [→ p. 4]
- Min Kyung Lee, Daniel Kusbit, Anson Kahng, Ji Tae Kim, Xinran Yuan, Allissa Chan, Daniel See, Ritesh Noothigattu, Siheon Lee, Alexandros Psomas, and Ariel D. Procaccia. WeBuildAI: Participatory framework for algorithmic governance. *Proceedings of the ACM on Human-Computer Interaction (HCI)*, 3, 2019. [→ p. 2]
- Xinran Liang, Katherine Shu, Kimin Lee, and Pieter Abbeel. Reward uncertainty for exploration in preference-based reinforcement learning. *arXiv:2205.12401*, 2022. [→ p. 2]

- Farhad Mohsin, Lei Luo, Wufei Ma, Inwon Kang, Zhibing Zhao, Ao Liu, Rohit Vaish, and Lirong Xia. Making group decisions from natural language-based preferences. In *Proceedings of the 8th International Workshop on Computational Social Choice (COMSOC)*, 2021. [→ p. 2, 16]
- Ritesh Noothigattu, Snehal Kumar S. Gaikwad, Edmond Awad, Sohan Dsouza, Iyad Rahwan, Pradeep Ravikumar, and Ariel D. Procaccia. A voting-based system for ethical decision making. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, pages 1587–1594, 2018. [→ p. 2, 3, 16]
- Rutvik Page, Ehud Shapiro, and Nimrod Talmon. Electing the executive branch. *arXiv:2009.09734*, 2020. [→ p. 8, 18]
- Dominik Peters. Proportionality and strategyproofness in multiwinner elections. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1549–1557, 2018. [→ p. 18]
- Dominik Peters and Martin Lackner. Preferences single-peaked on a circle. *Journal of Artificial Intelligence Research*, 68:463–502, 2020. [→ p. 13]
- Dominik Peters and Piotr Skowron. Proportionality and the limits of welfarism. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, pages 793–794, 2020. Full version arXiv:1911.11747. [→ p. 4, 6, 7, 22, 23]
- Dominik Peters, Grzegorz Pierczyński, and Piotr Skowron. Proportional participatory budgeting with additive utilities. In *Advances in Neural Information Processing Systems*, volume 34, pages 12726–12737, 2021. [→ p. 7, 18]
- Edvard Phragmén. Sur une méthode nouvelle pour réaliser, dans les élections, la représentation proportionnelle des partis. *Öfversigt af Kongliga Vetenskaps-Akademiens Förhandlingar*, 51(3):133–137, 1894. [→ p. 3, 4, 6]
- Robin L Plackett. The analysis of permutations. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 24(2):193–202, 1975. [→ p. 16]
- Luis Sánchez-Fernández, Edith Elkind, Martin Lackner, Norberto Fernández, Jesús Fisteus, Pablo Basanta Val, and Piotr Skowron. Proportional justified representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 670–676, 2017. [→ p. 4]
- Piotr Skowron. Proportionality degree of multiwinner rules. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, pages 820–840, 2021. [→ p. 18]
- Piotr Skowron and Adrian Górecki. Proportional public decisions. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5191–5198, 2022. [→ p. 4]
- Piotr Skowron, Martin Lackner, Markus Brill, Dominik Peters, and Edith Elkind. Proportional rankings. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 409–415, 2017. [→ p. 4]
- Thorvald N. Thiele. Om flerfold valg. *Oversigt over det Kongelige Danske Videnskabernes Selskabs Fordhandlingar*, 1895. [→ p. 3, 4, 7]

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. We also discuss Pareto efficiency.

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). This “simple proportionality” 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 $A_j^i = A_j^{i'}$ for all $i, i' \in S$ and that $A_j^i \cap A_j^{i'} = \emptyset$ for all $i \in S$ and $i' \in N \setminus S$.

For example, in [Figure 5](#), the groups $\{1\}$ and $\{2, 3\}$ are closed. [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). A decision sequence $D = (d_1, \dots, d_T)$ satisfies *perpetual lower quota for closed groups*¹¹ 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 5](#) 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 one of our previous axioms.

Theorem A.4. *PJR implies lower quota for closed groups.*

Proof. Let D be a decision sequence that satisfies 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, 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

A.2. Priceability

[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

¹¹[Lackner 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.

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 defined 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). Given a decision sequence $D = (d_1, \dots, d_k)$ of length k ,¹² we say that D is supported by the price system $(B, \{p_j\}_{j \leq k})$ where the real number $B \geq 0$ is the budget that each voter starts with and for each $j \in \{1, \dots, k\}$, $p_j : N \times C_j \rightarrow [0, 1]$ is a payment function, such that the following conditions hold:

- (P1) $p_j(i, c) = 0$ if $c \notin A_j^i$, i.e., no voter pays for an alternative that she does not approve.
- (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 .

With the above definition in place, we will now show that perpetual priceability implies PJR. (Similarly, in multi-winner voting, priceability implies PJR [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 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 PJR. Thus, there is $\ell \in \mathbb{N}$ and a group $S \subseteq N$ which agrees in all the T rounds, has size $|S| = s \geq \ell \frac{T}{n}$, 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 $|S| \cdot \frac{T}{n} \geq \ell$. 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, PJR is satisfied.

So assume $B > \frac{T}{n}$. Because at most $\ell - 1$ decisions in D are approved by at least one member of S , the members of S have spent at most $\ell - 1$ units until after round T . Since they start out with a total budget of sB , we see that the voters in S have a remaining budget of at least

$$sB - (\ell - 1) \geq sB - (s \frac{T}{n} - 1) > s \frac{T}{n} - (s \frac{T}{n} - 1) \geq 1 \quad (2)$$

¹²We use the variable k to denote the length instead of T since the definition of perpetual priceability references decision sequences that are shorter than the full time horizon T .

Hence there is an $\varepsilon > 0$ such that the remaining budget of the voters in S is $1 + \varepsilon$. By the definition of minimal price systems, for every $1 \leq r \leq T$, there is 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 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

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

Now let r^* be the largest index r for which $B_r < B$ (recalling that at the end, $B_T = B$). We claim that there is a B' with $B_{r^*} \leq B' < B_{r^*+1} = B$ such that there is an alternative $a' \in C_{r^*+1}$ and 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

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

Observe that $B_{r^*} \leq B' < B$. 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 $c \neq a'$. This is possible, because the voters in S have at least a budget of 1 in round $r^* + 1$ (from (2)). Hence, this is a price system that supports $(d_1, \dots, d_{r^*}, a')$ which contradicts the minimality of $(B, \{p_j\}_{j \leq T})$. \square

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

Theorem A.9. *Perpetual Priceability does not imply Strong PJR.*

Proof. Consider the decision instance shown in Figure 9 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 Strong 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, everyone receives the budget of 1 unit after round 1 and only 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 brought 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 Strong PJR: Consider the coalition $S = \{v_2, v_3\}$ which agrees in the last $k = 3$ rounds. For $\ell = 1$, we have $|S| \geq \ell \frac{n}{k} = \frac{6}{3} = 2$. Thus, Strong 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 Strong PJR. Hence, perpetual priceability does not imply Strong 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 Strong JR.

Note that perpetual priceability implies Strong 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

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}

Figure 9: Example instance where a rule may produce a perpetually priceable outcome which may not satisfy Strong JR.

sequence satisfies PJR, which is enough to imply Strong PJR with initial agreement). Note that the example in [Theorem A.9](#) does not feature initial agreement.

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.)

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, 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 10](#) 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 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

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

Figure 10: Example instance where both MES and Sequential Phragmén produce Pareto inefficient outcomes.

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. More Details About MES Failing Strong PJR and Strong EJR

In [Theorem 4.4](#), we showed that MES fails Strong PJR. Notably, in our counterexample, MES terminates prematurely. Here we show that, in a sense, this is the cause for the Strong 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 Strong PJR. (However note that the proof of [Theorem 4.4](#) showed an instance where MES fails Strong PJR no matter what completion strategy is used.)

Theorem B.1. *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 $|S| \geq \ell \frac{n}{T}$, 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). Thus, D satisfies Strong PJR.*

Proof. Let $\ell \in \mathbb{N}$ and let $S \subseteq N$ be a group of voters with $|S| \geq \ell \frac{n}{T}$. 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 voters in S must be at least $|S|/p = |S| \cdot \frac{T}{n} \geq \ell$ rounds, as desired. \square

We can ask further if MES satisfies even Strong EJR in cases where it does not terminate prematurely; however this is not the case.

Theorem B.2. *MES can fail Strong EJR even on instances where it outputs a decision for all rounds.*

Proof. Consider the instance shown in [Figure 11](#) which is adapted from the instance shown in [Figure 10](#) where MES (and Sequential Phragmén) do not produce Pareto-efficient outcomes. We have $T = 16$ rounds and $n = 8$ voters. The group of voters $S_1 = \{1, 2, 3\}$ 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 $p = \frac{n}{T} = \frac{8}{16} = \frac{1}{2} = 0.5$ 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, 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 $|S_1| \geq \ell \cdot \frac{n}{k} = 3 \cdot \frac{8}{10}$, Strong 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 Strong EJR. \square

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 }

Figure 11: Example instance where MES does not terminate prematurely yet fails Strong EJR.

B.2. Impossibility of Stronger Guarantees: More Details

Here we give proofs and additional results related to the discussion in [Section 5](#).

Theorem 5.1. *Let $\varepsilon > 0$. Choose some $k > \lceil \frac{1-\varepsilon}{\varepsilon} \rceil$ 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 $|S| \geq (1 - \varepsilon) \cdot \frac{n}{k}$, yet none of the voters in S approve any of the decisions of D .*

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

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

Adding nk to both sides, and rearranging,

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

Finally, dividing by k , we get

$$n \geq n(1 - \varepsilon) + \frac{n(1-\varepsilon)}{k} + T + 1. \quad (3)$$

We now construct an instance with n voters and time horizon T . Write $s = \lceil (1 - \varepsilon) \cdot \frac{n}{k} \rceil$. In 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

$$k \cdot s + (T - k) = k \cdot \lceil (1 - \varepsilon) \cdot \frac{n}{k} \rceil + (T - k) \leq n \cdot (1 - \varepsilon) + k + (T - k) = n \cdot (1 - \varepsilon) + T.$$

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

$$n - (n \cdot (1 - \varepsilon) + T) \stackrel{(3)}{\geq} \frac{n(1-\varepsilon)}{k} + 1 \geq s.$$

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, and $|S| = s = \lceil (1 - \varepsilon) \cdot \frac{n}{k} \rceil$. Thus, S satisfies the conditions promised in the theorem statement. \square

In [Corollary 5.3](#), we showed that a property we called *Multi-Winner PJR* need not exist – this property requires that a group of voters who agree in ℓ rounds and has at least $\ell \frac{n}{T}$ members should approve at least ℓ decisions. Following [Do et al. \[2022\]](#), we now consider two relaxations of this 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 requires a group that the group approves only $\lfloor \alpha \cdot \ell \rfloor$ 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 5.1](#) with $\varepsilon = 0.5$, $k = \lceil \frac{1}{\alpha} \rceil$, 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

$$|S| \geq 0.5 \cdot \frac{n}{k} = 0.5 \cdot \frac{n}{\lceil 1/\alpha \rceil} \geq 0.5 \cdot \frac{n}{2/\alpha} = \frac{1}{4} \cdot \frac{\alpha^2}{\alpha} \cdot n = \frac{2}{\alpha} \cdot \frac{n}{8/\alpha^2} \geq \ell \cdot \frac{n}{T}.$$

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 rules 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 5.1](#) with $\varepsilon = 0.5$, $k = \lceil \frac{1}{\alpha} \rceil$, 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

$$|S| \geq 0.5 \cdot \frac{n}{k} \geq 0.5 \cdot \frac{n}{2/\alpha} = 0.25 \cdot \alpha \cdot n = \frac{n}{4/\alpha} \geq \ell \cdot \frac{n}{T}.$$

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 5.1](#): The first $T - 1$ rounds remain exactly the same as in that construction. We observe the decisions made by the online rule in these rules. 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 the offline rule cannot find it. This kind of construction works for all the axiom variants we have considered in this section.

B.3. Can Online Rules Satisfy EJR?

We have seen that the online rule Sequential Phragmén fails EJR, while the semi-online MES satisfies EJR. Is it possible for an online rule to satisfy EJR? While we do not know the answer to this question,¹³ we show that any online rule that satisfies 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 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.

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 $|W| = k$.

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). A multi-winner voting rule f satisfies *committee monotonicity* if for all N , A , C , and $k \geq 1$, we have $f(N, C, A, k - 1) \subseteq f(N, C, A, k)$.

Our EJR axiom is inspired by the Extended Justified Representation axiom originally proposed for the multi-winner voting context by [Aziz et al. \[2017\]](#).

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.

¹³In 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 et al., 2022](#)]. Sequential Phragmén fails EJR even in this special case.

Theorem B.8. *An online decision rule that satisfies 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 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.

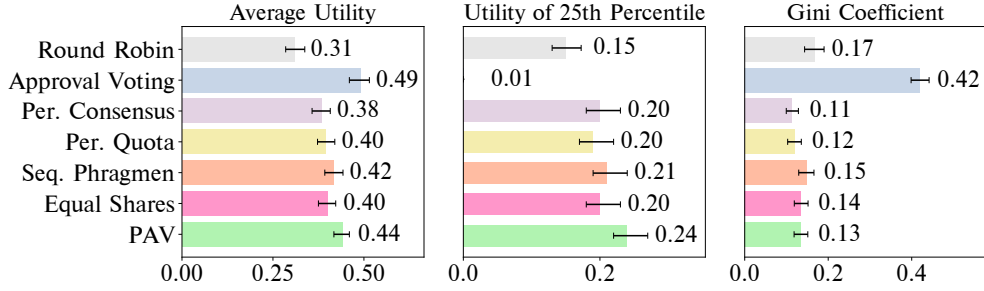
To finish, we need to prove that the multi-winner voting rule we have defined satisfies Multi-Winner EJR. Let k be a desired committee size, and let $W = \{d_1, \dots, d_k\}$. Let $\ell \in \mathbb{N}$ and suppose that $S \subseteq N$ is a group of voters with $|S| \geq \ell \cdot \frac{n}{k}$ and $|\bigcap_{i \in S} A^i| \geq \ell$. 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 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

C. Full Results from Experiments

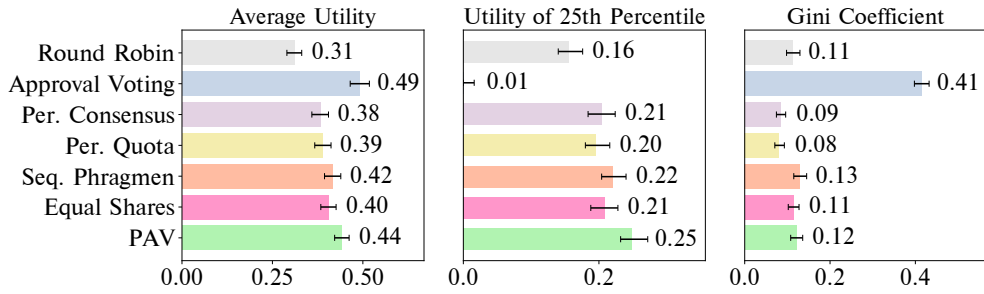
C.1. Synthetic Data

C.1.1. Restricted distribution

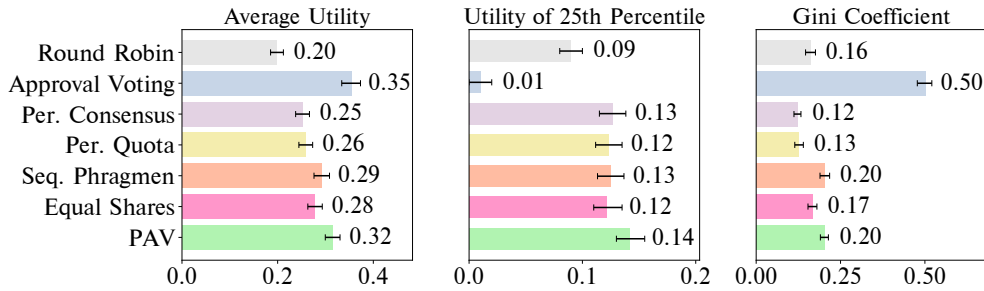
Figure 12: Performance of the different rules on the *Restricted* distribution for various parameter values. The length of the bar represents the median across all the trials while the error bars represent the 25th and 75th percentile with the numeric text after the bar showing the mean value.



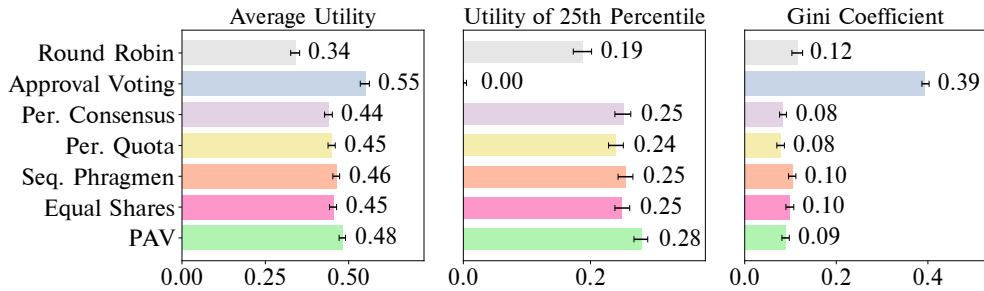
(a) Results for $n = 20$ voters, $T = 20$ rounds, 20 alternatives in each round, with $f = 1.5$.



(b) Results for $n = 20$ voters, $T = 50$ rounds, 20 alternatives in each round, with $f = 1.5$.



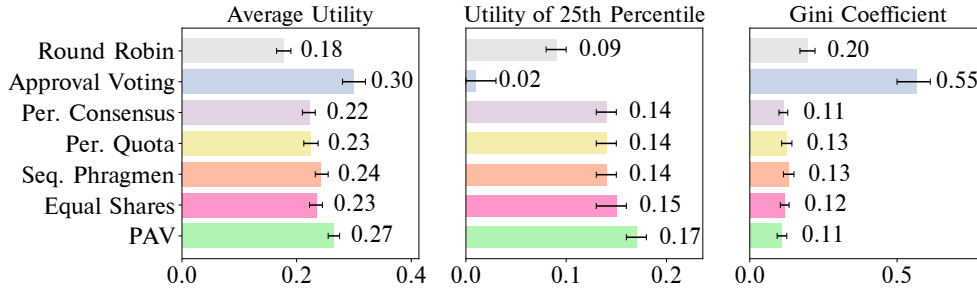
(c) Results for $n = 50$ voters, $T = 50$ rounds, 40 alternatives in each round, with $f = 1.5$.



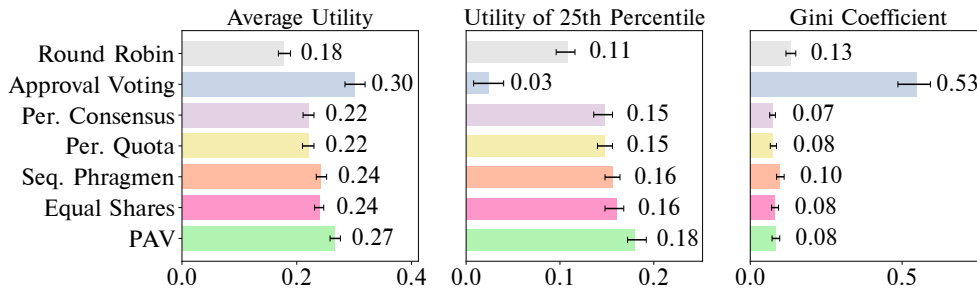
(d) Results for $n = 50$ voters, $T = 50$ rounds, 40 alternatives in each round, with $f = 3$.

C.1.2. Many Groups distribution

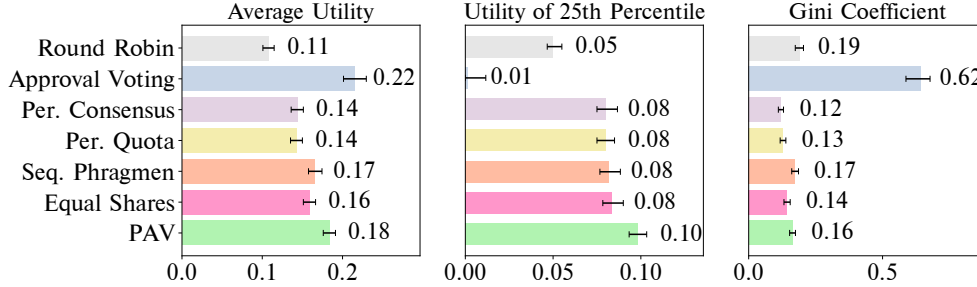
Figure 13: Performance of the different rules on the *Many Groups* distribution for various parameter values. The length of the bar represents the median across all the trials while the error bars represent the 25th and 75th percentile with the numeric text after the bar showing the mean value.



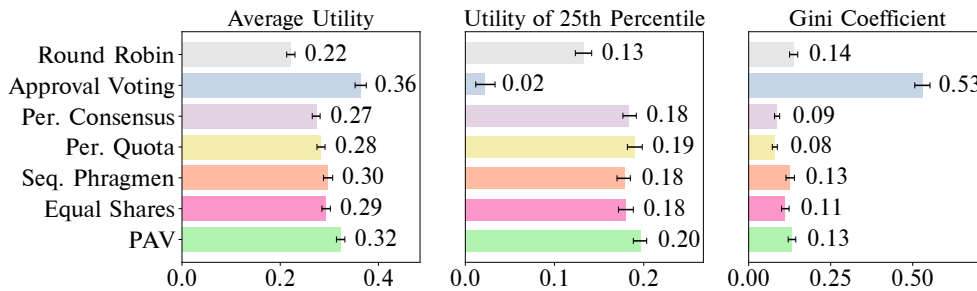
(a) Results for $n = 20$ voters, $T = 20$ rounds, 20 alternatives in each round, with $f = 1.5$.



(b) Results for $n = 20$ voters, $T = 50$ rounds, 20 alternatives in each round, with $f = 1.5$.



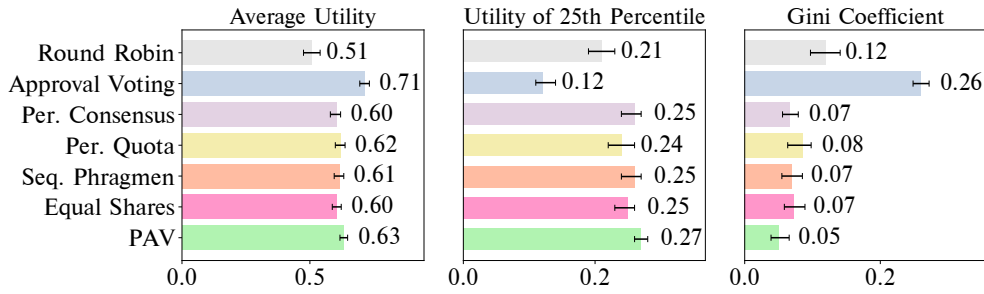
(c) Results for $n = 50$ voters, $T = 50$ rounds, 40 alternatives in each round, with $f = 1.5$.



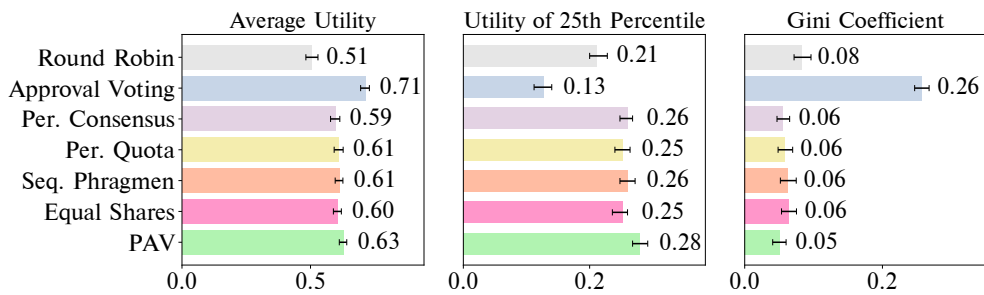
(d) Results for $n = 50$ voters, $T = 50$ rounds, 40 alternatives in each round, with $f = 3$.

C.1.3. Unbalanced distribution

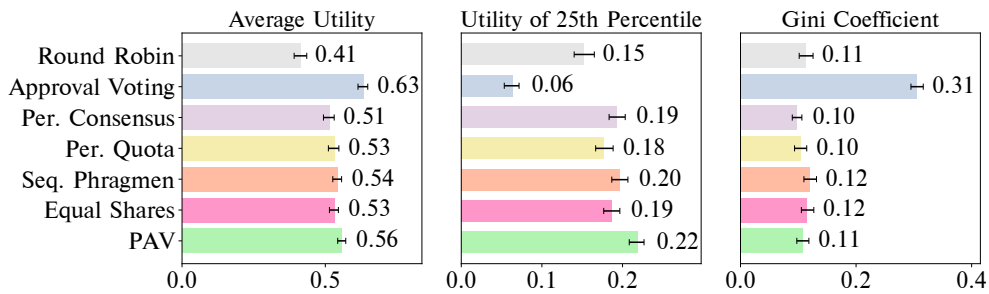
Figure 14: Performance of the different rules on the *Unbalanced* distribution for various parameter values. The length of the bar represents the median across all the trials while the error bars represent the 25th and 75th percentile with the numeric text after the bar showing the mean value.



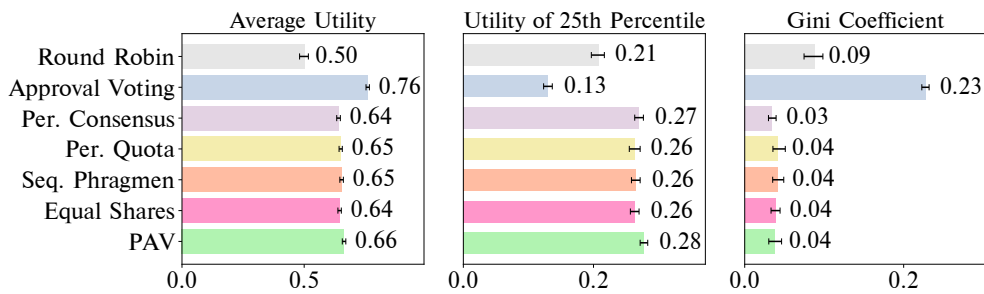
(a) Results for $n = 20$ voters, $T = 20$ rounds, 20 alternatives in each round, with $f = 1.5$.



(b) Results for $n = 20$ voters, $T = 50$ rounds, 20 alternatives in each round, with $f = 1.5$.



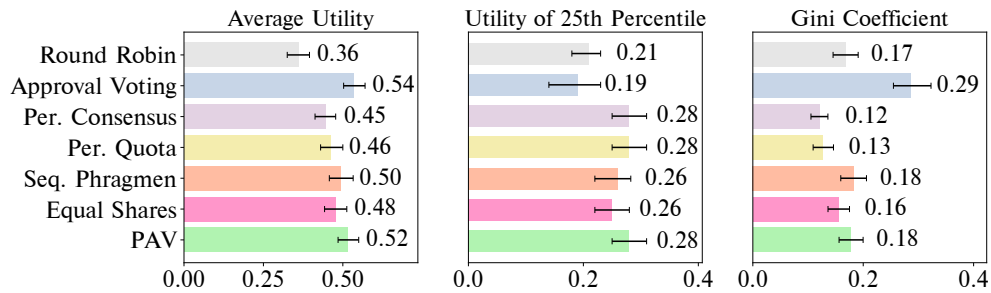
(c) Results for $n = 50$ voters, $T = 50$ rounds, 40 alternatives in each round, with $f = 1.5$.



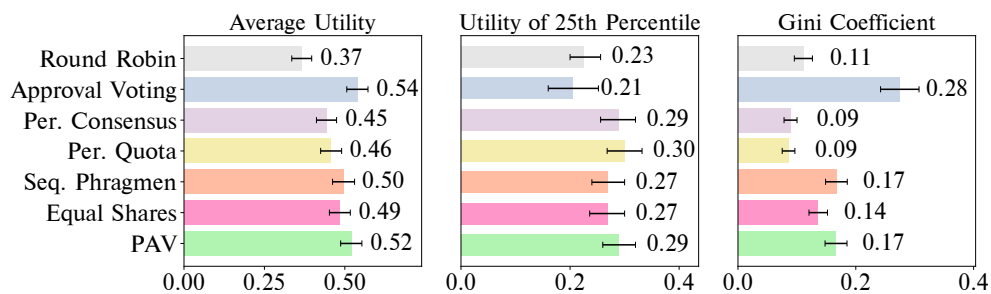
(d) Results for $n = 50$ voters, $T = 50$ rounds, 40 alternatives in each round, with $f = 3$.

C.1.4. Balanced and Nearby distribution

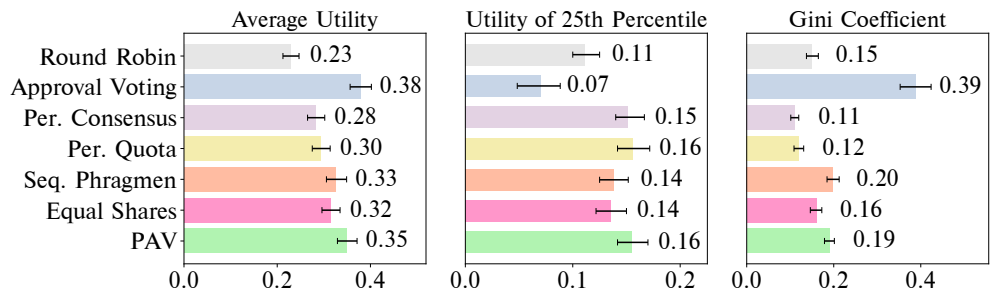
Figure 15: Performance of the different rules on the *Balanced and Nearby* distribution for various parameter values. The length of the bar represents the median across all the trials while the error bars represent the 25th and 75th percentile with the numeric text after the bar showing the mean value.



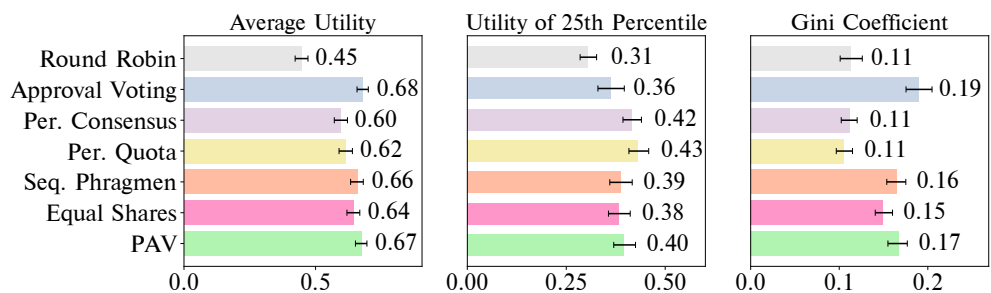
(a) Results for $n = 20$ voters, $T = 20$ rounds, 20 alternatives in each round, with $f = 1.5$.



(b) Results for $n = 20$ voters, $T = 50$ rounds, 20 alternatives in each round, with $f = 1.5$.



(c) Results for $n = 50$ voters, $T = 50$ rounds, 40 alternatives in each round, with $f = 1.5$.

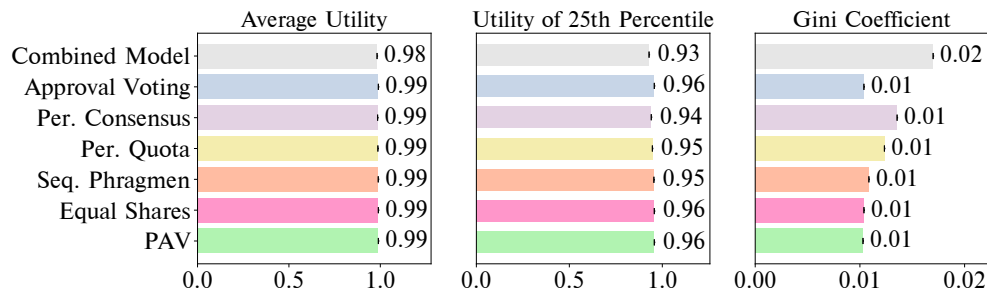


(d) Results for $n = 50$ voters, $T = 50$ rounds, 40 alternatives in each round, with $f = 3$.

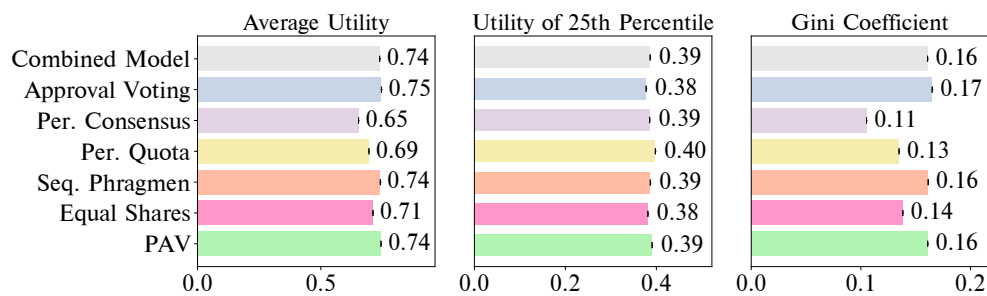
C.2. Learning Preferences from the Moral Machine

C.2.1. Random Alternatives

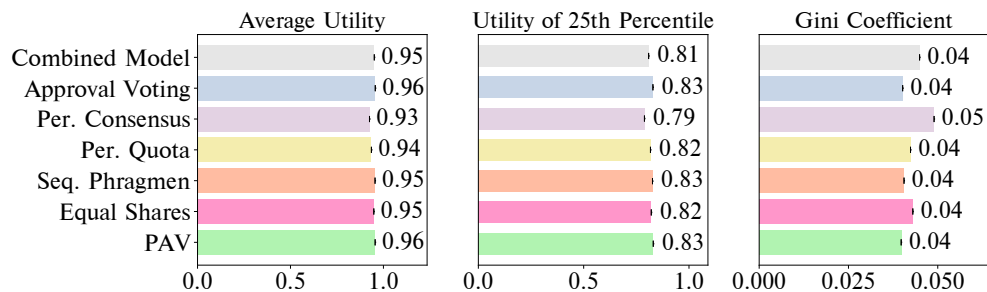
Figure 16: Performance of the different rules on the Moral Machine Dataset based on random alternatives.



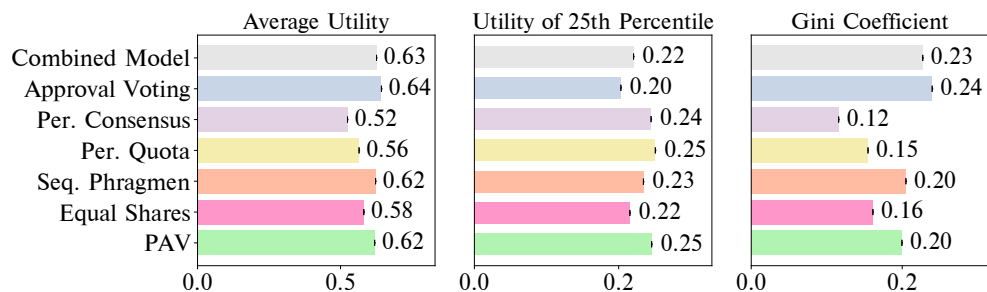
(a) 100 rounds with each round having 50 alternatives and each voter approving alternative having more than average utility.



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



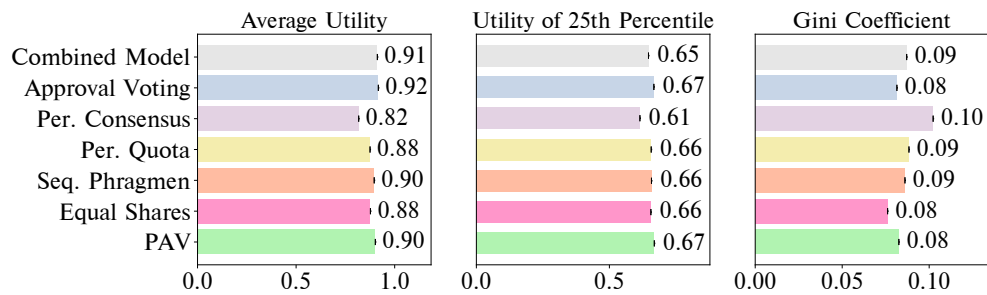
(c) 100 rounds with each round having 50 alternatives and each voter approving top 10 alternatives.



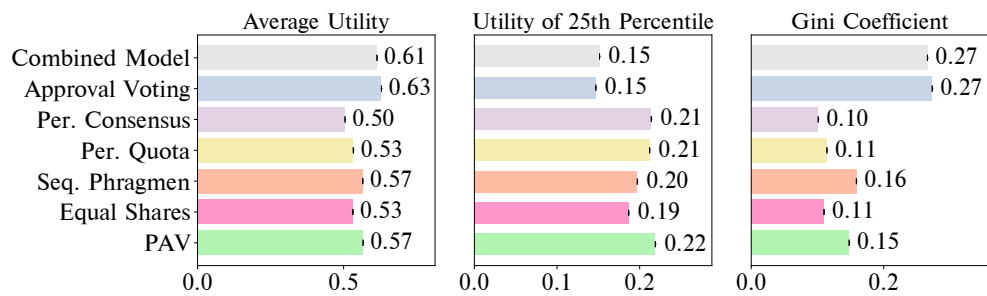
(d) 100 rounds with each round having 200 alternatives and each voter approving top 2 alternatives.

C.2.2. High-disagreement Alternatives

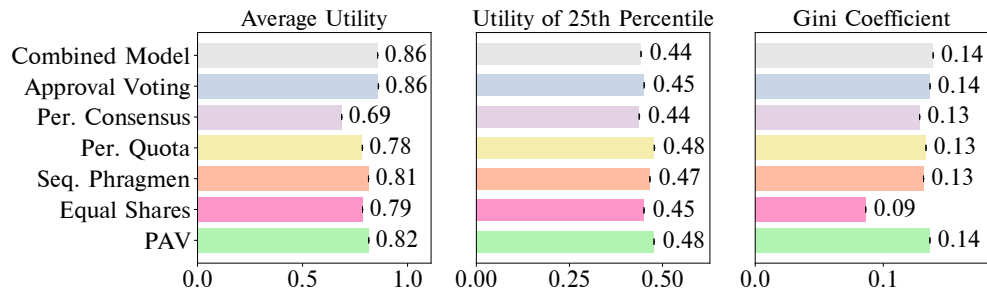
Figure 17: Performance of the different rules on the Moral Machine Dataset, based on high-disagreement alternatives.



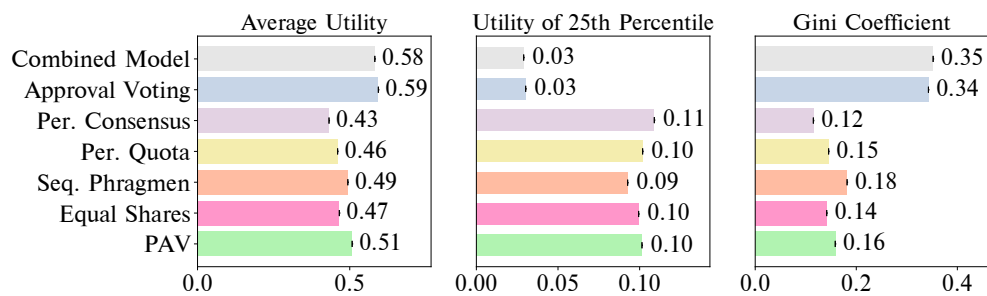
(a) 100 rounds with each round having 50 alternatives and each voter approving alternative having more than average utility.



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



(c) 100 rounds with each round having 50 alternatives and each voter approving top 10 alternatives.



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