

Perpetual Voting: Fairness in Long-Term Decision Making

Martin Lackner

TU Wien

Vienna, Austria

lackner@dbai.tuwien.ac.at

Abstract

In this paper we introduce a new voting formalism to support long-term collective decision making: perpetual voting rules. These are voting rules that take the history of previous decisions into account. Due to this additional information, perpetual voting rules may offer temporal fairness guarantees that cannot be achieved in singular decisions. In particular, such rules may enable minorities to have a fair (proportional) influence on the decision process and thus foster long-term participation of minorities. This paper explores the proposed voting rules via an axiomatic analysis as well as a quantitative evaluation by computer simulations. We identify two perpetual voting rules as particularly recommendable in long-term collective decision making.

1 Introduction

Consider the following simple scenario: A group of five friends have a joint dinner every week. Three of them prefer French food, two of them Indian food. It is apparent that a fair suggestion would consist of going to French restaurants 60% of the time and to Indian restaurants 40% of the time, although there is a majority that prefers French food every single week.

In this simple example it is clear that a majoritarian decision is not desirable, instead a proportional distribution constitutes a fair solution. But how can this form of fairness be achieved when preferences change over time? Moreover, is such fairness possible if the available alternatives change or even if subsequent decisions concern different topics? For example, consider a company, club, or university department where collective decisions are repeatedly made on varying topics. In this paper, we develop voting methods that can deal with such complex scenarios of *perpetual voting*. The central emphasis of this paper is the *long-term* aspect: instead of considering singular decisions to be made, we view the current decision in the context of previously taken decisions, the decision history.

The goal of this paper is to introduce a new form of voting rules—perpetual voting rules—that are suitable in a long-term decision making process and can guide such a pro-

cess in a mathematically principled and fair way. Their main strength is that they can ensure fairness *across* a sequence of decisions, even if the individual decisions are contested and no compromise alternatives are available. Ideally, perpetual voting rules guarantee each voter a fair amount of influence on the decision process. In this way, also minorities are incentivized to participate and we obtain a sustainable (stable) process of collective decision making. The inclusion of minorities can be important both from a social perspective (for the social cohesion of the collective) and from a technical perspective (to be able to base decisions on more representative and detailed preference data). In summary, perpetual voting rules can be beneficial in all repeated collective decision making scenarios where particular emphasis is given to fairness towards minorities and individual agents¹.

The main contribution of our paper is the introduction of perpetual voting as a formalism, together with axioms and metrics to evaluate perpetual voting rules. Our research results in the recommendation of two specific perpetual voting rules: Perpetual Consensus and Perpetual Quota (defined in Section 3). The former achieves the best results in the axiomatic analysis, the latter proves to be particularly strong in the experimental evaluation (while also enjoying good axiomatic properties).

In more detail, the contributions of this paper are as follows. In Section 2 we introduce the perpetual voting formalism based on approval ballots. Then, in Section 3, we suggest eight voting rules suitable for perpetual voting, some of which are inspired by the literature on multiwinner voting (Kilgour 2010; Faliszewski et al. 2017).

We analyze the proposed rules based on three axiomatic properties (Section 4). All three express a certain form of fairness in the perpetual voting setting. The first, *simple proportionality*, requires that perpetual voting rules behave proportionally on certain simple instances related to the apportionment setting (apportionment being the task of assigning seats to parties in a parliamentary assembly (Balinski and Young 1982; Brill, Laslier, and Skowron 2018)). Even though simple proportionality is a rather weak proportional-

¹This assumption does not necessarily hold for political, high-stake decision making, in particular in the presence of extremist opinions that are harmful to society.

ity axiom, it is sufficiently strong to reveal several perpetual voting rules as being non-proportional. The second axiom, *independence of unanimous decisions*, states that decisions with an alternative that everyone agrees with should not impact future decisions. This axiom is founded in the idea that the decision process should not be manipulable by putting “fake” (uncontroversial) decisions on the agenda. The third axiom, *bounded dry spells*, gives voters an individual guarantee that they will be satisfied with at least one decision in every interval of a certain (bounded) length. With this fairness guarantee we can provide voters a reason to participate in the decision process as—at least from time to time—their preference is going to be reflected in the choices.

Our second analysis is based on numerical simulations. To this end, we introduce two performance metrics in Section 5. The first metric, *perpetual lower quota compliance*, can be seen as the likelihood for a voter to be satisfied with at least a proportional fraction of the made choices. The second metric, the *Gini influence coefficient*, measures inequality in the degree of influence that voters have on the decision process. For example, a decision process with a dictator has a high Gini influence coefficient. In our experimental analysis (Section 6) these two metrics yield a numerical evaluation of our proposed methods, and by that refine the picture resulting from the axiomatic analysis.

Furthermore, we include a probabilistic voting rule, Random Serial Dictatorship (RSD), in our experimental setup for comparison (see (Brandt 2017) for a survey on this topic). While probabilistic rules do not explicitly include a mechanism for sequences of decisions, repeated applications do yield fairness properties that can be seen as proportionality *in the long run*. However, in our experiments RSD achieves a worse performance than most perpetual voting rules, which we take as an indication that non-randomized rules have a clear advantage in our envisioned setting. Finally, we discuss research directions and practical issues concerning perpetual voting in Section 7.

Related work

While the specific goal and model of perpetual voting is new, there are related approaches that consider temporal aspects of voting or sequences of decisions. We provide a brief summary:

A particularly relevant formalism is that of Freeman, Zahedi, and Conitzer (2017). They consider the problem of choosing an alternative each round according to reported utility functions. Their goal is to maximize long-term Nash welfare. In its motivation this paper is similar to ours, as it has the same temporal component. However, it differs in the essential aspect that decisions can be based on utility functions, which greatly simplifies the challenge of defining fairness over time (and thus choosing an optimization objective). The connection between perpetual voting and the utility-based methods in this paper can thus be seen as similar to the relation of classic preferential voting rules and collective utility functions, cf. (Moulin 1991).

A preference-based voting rule that incorporates the temporal aspect of perpetual voting is the *storable votes* method (Casella 2005; 2012), which is based on plurality

voting. However, this voting rule does not directly fit into the framework of perpetual voting, as minorities are required to strategically store and spend votes to obtain a fair share of power. In contrast, perpetual voting rules aim to provide such guarantees without the need of strategic voting.

Dynamic voting or dynamic social choice considers voting with changing preference profiles (Tennenholtz 2004; Oren and Lucier 2014; Hemaspaandra, Hemaspaandra, and Rothe 2017; Parkes and Procaccia 2013; Boutilier and Procaccia 2012). These models essentially evolve around a single decision (in contrast to perpetual voting), although the outcome may change over time due to the dynamic setting. In a similar vein, dynamic preferences have also been considered in matching and resource allocation (Dickerson, Procaccia, and Sandholm 2012; Hosseini, Larson, and Cohen 2015; Freeman et al. 2018).

Sequential voting rules (Lang and Xia 2009) are voting rules over combinatorial domains (Lang and Xia 2016) that are based on the assumption that there exists a natural order of issues (issues correspond to decisions in our setting). The key assumption is that preferences concerning an issue are only influenced by issues appearing earlier in the order. Perpetual voting fundamentally differs from sequential voting in several ways: (i) In perpetual voting the number and content of issues/decisions is not known up front. (ii) Perpetual voting rules depend on the history of previous decisions; only due to this dependency we can achieve fairness over time. In sequential voting only the preferences of voters depend on previous choices. (iii) Sequential voting was introduced to tackle a different problem: the representation of preferences in combinatorial domains.

2 Perpetual Voting

We will now introduce our proposed formalism, alongside necessary basic definitions. Let $N = \{1, \dots, n\}$ be a set of voters (agents). Given a set of alternatives C , we assume that each voter $v \in N$ approves some non-empty subset of C . An *approval profile* $A = (A(1), \dots, A(n))$ for C is an n -tuple of subsets of C , i.e., $A(v) \subseteq C$ for $v \in N$. We call the triple (N, A, C) a *decision instance*.

A *k*-*decision sequence* (N, \bar{A}, \bar{C}) is a triple consisting of a set of voters N , a k -tuple of approval profiles $\bar{A} = (A_1, A_2, \dots, A_k)$ and an associated k -tuple of alternatives $\bar{C} = (C_1, \dots, C_k)$. Thus, for each $i \leq k$, the triple (N, A_i, C_i) is a decision instance and can be seen as an individual decision to be made; we refer to (N, A_i, C_i) as *decision in round i*. Furthermore, note that we assume that the set of voters remains the same in all rounds; we discuss how to weaken this assumption in Section 7.

We write $\bar{w} \in \bar{C}$ to denote a k -tuple $\bar{w} = (w_1, \dots, w_k)$ with $w_i \in C_i$ for $i \in \{1, \dots, k\}$; we refer to \bar{w} as a *k*-*choice sequence*. This tuple represents the chosen alternatives in rounds 1 to k . If we combine a k -decision sequence (N, \bar{A}, \bar{C}) and a k -choice sequence $\bar{w} \in \bar{C}$, we speak of a *k*-*decision history* $(N, \bar{A}, \bar{C}, \bar{w})$, which can be seen as the history of past decisions alongside the made choices. We thus know, for any $i \leq k$, that in case of decision instance (N, A_i, C_i) alternative w_i was chosen.

Assume that a group of voters N wants to take a decision and looks back at k decisions already taken. That is, we are presented with a k -decision history $(N, \bar{A}, \bar{C}, \bar{w})$ and a decision instance (N, A_{k+1}, C_{k+1}) . The question is now which alternative in C_{k+1} should be chosen, subject to the preferences in A_{k+1} and under consideration of the decision history. An (*approval-based*) *perpetual voting rule* \mathcal{R} is a function that maps a pair of a k -decision instance (N, A_{k+1}, C_{k+1}) and a decision history $(N, \bar{A}, \bar{C}, \bar{w})$ to an alternative in C_{k+1} . We write $\mathcal{R}(N, \bar{A}, \bar{C})$ to denote the k -choice sequence $\bar{w} \in \bar{C}$ as selected by the perpetual voting rule \mathcal{R} , that is, $\mathcal{R}(N, \bar{A}, \bar{C}) = \bar{w}$ is inductively defined by $w_i = \mathcal{R}(N, (A_1, \dots, A_i), (C_1, \dots, C_i), (w_1, \dots, w_{i-1}))$ for $i \leq k$. Let us now describe some basic statistics of k -decision histories.

Definition 1. Let $(N, \bar{A}, \bar{C}, \bar{w})$ be a k -decision history, and $j \leq k$. The satisfaction of voter $v \in N$ in round j is

$$sat_j(v, \bar{w}) = |\{i \leq j : w_i \in A_i(v)\}|.$$

The support of a voter $v \in N$ in round j is defined as

$$supp_j(v) = \frac{1}{n} \cdot \max_{c \in A_j(v)} |\{u \in V : A_j(u) = c\}|.$$

The quota of voter $v \in N$ in round j is

$$qu_j(v) = \sum_{i \leq j} supp_i(v).$$

Thus, the satisfaction of a voter is the number of past decisions that have satisfied this voter. The support of a voter in round j is the proportion of voters that can collectively agree on some alternative that v approves. The quota of voter v in round j is v 's cumulative support from round 1 to j . Note that although satisfaction, support, and quota clearly depend on (N, \bar{A}, \bar{C}) , we do not explicitly mention that in the notation (and other definitions throughout the paper).

3 Perpetual Voting Rules

In this section we define several approval-based perpetual voting rules. As we assume that perpetual voting rules are resolute, i.e., return exactly one winning alternative, we require a tie-breaking order to resolve ties. Thus, we assume throughout the paper that there exists some arbitrary and fixed order for each set of alternatives that settles ties.

We start by introducing the class of *perpetual weighted approval methods*, which contains many of our proposed rules. These approval-based perpetual voting rules are defined as follows: Each voter has an assigned positive weight, which may change each round; a larger weight corresponds to being assigned a higher importance. Let $\alpha_k(v)$ denote voter v 's weight in round k . Weights are initialized with $\alpha_1(v) = 1$ for all $v \in N$. Given a k -decision history (N, \bar{A}, \bar{C}) and a decision instance (N, A_{k+1}, C_{k+1}) , the rule selects an alternative $w_{k+1} \in C_{k+1}$ with maximum weighted approval score. That is, the score of an alternative c is defined as

$$sc_{k+1}(c) = \sum_{v \in N \text{ with } c \in A_{k+1}(v)} \alpha_{k+1}(v).$$

To compute the weights of voters after round 1, there exist functions f, g such that for all $v \in N$,

$$\alpha_{k+1}(v) = \begin{cases} f(\alpha_k(v)) & \text{if } w_k \notin A_k(v), \\ g(\alpha_k(v)) & \text{if } w_k \in A_k(v). \end{cases}$$

The following five rules are perpetual weighted approval methods. It is thus sufficient to define the weight function α .

AV. The simplest example of a perpetual weighted approval method is approval voting (AV) itself. As it completely ignores the history of past decisions, we mention it mainly for the sake of comparison. AV corresponds to the perpetual weighted approval method with $\alpha_{k+1}(v) = 1$.

Perpetual PAV. This method is inspired by Proportional Approval Voting (Thiele 1895; Faliszewski et al. 2017), or more specifically by its sequential counterpart, and is thus based on the harmonic series. The weight of voters is defined by²

$$\alpha_{k+1}(v) = \frac{1}{sat_k(v, \bar{w}) + 1}.$$

Perpetual Unit-Cost. This rule is based on the idea that satisfied voters “pay” a cost of winning (which is 1), and—after a decision is taken—the weight of *all* voters is increased by 1.

$$\alpha_{k+1}(v) = \begin{cases} \alpha_k(v) + 1 & \text{if } w_k \notin A_k(v), \\ \alpha_k(v) & \text{if } w_k \in A_k(v). \end{cases}$$

Perpetual Reset. This rule is similar to Perpetual Unit-Cost, but the weight of satisfied voters is reset to 1, i.e.,

$$\alpha_{k+1}(v) = \begin{cases} \alpha_k(v) + 1 & \text{if } w_k \notin A_k(v), \\ 1 & \text{if } w_k \in A_k(v). \end{cases}$$

Perpetual Equality. Perpetual Equality is inspired by the Chamberlin–Courant rule (Chamberlin and Courant 1983). Let $s = \min_{v \in N} sat_k(v, \bar{w})$ be the minimum satisfaction after round k . The goal of this rule is to make a choice that satisfies as many voters with a satisfaction of s as possible. In case two choices satisfy the same number of voters with a satisfaction of s , it chooses the alternative that further satisfies as many voters with a satisfaction of $s + 1$ as possible, etc. Formally³,

$$\alpha_{k+1}(v) = n^{-sat_k(v, \bar{w})},$$

i.e., one voter with satisfaction s has a larger weight than $n - 1$ voters with satisfaction $s + 1$.

²To see that Perpetual PAV is indeed a perpetual weighted approval method note that

$$\alpha_{k+1}(v) = \frac{1}{sat_k(v, \bar{w}) + 1} = \begin{cases} \alpha_k(v) & \text{if } w_k \notin A_k(v), \\ \frac{\alpha_k(v)}{\alpha_k(v) + 1} & \text{if } w_k \in A_k(v). \end{cases}$$

³Perpetual Equality is indeed a perpetual weighted approval method as $n^{-sat_k(v, \bar{w})} = \begin{cases} \alpha_k(v) & \text{if } w_k \notin A_k(v), \\ n^{\log_n(\alpha_k(v)) - 1} & \text{if } w_k \in A_k(v). \end{cases}$

The final three rules do not fall into the class of perpetual weighted approval methods.

Perpetual Quota. Perpetual Quota aims at granting each voter a satisfaction as close as possible to their quota. This rule is defined analogously to a perpetual weighted approval method with weights defined as $\alpha_1(v) = qu_1(v)$ and

$$\alpha_{k+1}(v) = \max(0, qu_{k+1}(v) - sat_k(v, \bar{w})).$$

Since $\alpha_{k+1}(v)$ requires knowledge about A_{k+1} (due to $qu_{k+1}(v)$), Perpetual Quota is not a perpetual weighted approval method. Note that the dependency on A_{k+1} implies that the weights of voters cannot be calculated before voters' preferences are known.

Perpetual Consensus. This rule is based on the idea of a “virtual” consensus: in total, the weight of all satisfied voters is decreased by n . In the ideal case of an actual consensus, 1 is subtracted from each voter. Otherwise, more than 1 is subtracted from the satisfied voters, possibly leading to negative weights. After each decision, the weight of all voters is increased by 1. Formally, let

$$N_k^+(c) = \{v \in N : c \in A_k(v) \text{ and } \alpha_k(v) > 0\},$$

and for all $v \in N$, $\alpha_1(v) = 1$ and

$$\alpha_{k+1}(v) = \begin{cases} \alpha_k(v) + 1 & \text{if } v \notin N_k^+(w_k), \\ \alpha_k(v) + 1 - \frac{n}{|N_k^+(w_k)|} & \text{if } v \in N_k^+(w_k). \end{cases}$$

The preferences of voters with non-positive weights are not taken into account. Thus, the score of an alternative c is defined as

$$sc_{k+1}(c) = \sum_{v \in N_{k+1}^+(c)} \alpha_{k+1}(v).$$

Perpetual Nash. This rule can be seen as an adaption of the GREEDY algorithm from Freeman, Zahedi, and Conitzer (2017) to our setting. It maximizes Nash welfare, i.e., the product of voters' utilities. We assume that the voters' utility is their satisfaction if their satisfaction is larger than 0; voters with a satisfaction of 0 have a utility of some small constant ϵ , e.g., $\epsilon = n^{-n}$. Let

$$u_{k+1}(v, c) = \begin{cases} \max(sat_k(v, \bar{w}), \epsilon) & \text{if } c \notin A_{k+1}(v), \\ sat_k(v, \bar{w}) + 1 & \text{if } c \in A_{k+1}(v). \end{cases}$$

be the utility of voter v if c is chosen. The Nash score of an alternative c is $nash_{k+1}(c) = \prod_{v \in N} u_{k+1}(v, c)$. The alternative with maximum Nash score is chosen.

Due to their sequential nature, all aforementioned rules can be computed in polynomial time.

4 Axiomatic Properties

As a first approach to understand the differences between perpetual voting rules, we introduce three axioms that relate to aspects of fairness: simple proportionality, independence of uncontroversial decisions, and bounded dry spells. The first, simple proportionality, is a classic proportionality axiom translated to our setting of perpetual voting.

Definition 2 (Simple proportionality). *We say that a k -decision sequence (N, \bar{A}, \bar{C}) is simple if $A_1 = \dots = A_k$, $C_1 = \dots = C_k$, and $|A_1(v)| = 1$ for all $v \in N$. For any simple n -decision sequence (N, \bar{A}, \bar{C}) with $|N| = n$, we say that $\bar{w} \in \bar{C}$ is proportional if $sat_n(v, \bar{w}) = qu_n(v)$ for every voter $v \in N$. A perpetual voting rule \mathcal{R} satisfies simple proportionality if for any simple $|N|$ -decision sequence (N, \bar{A}, \bar{C}) , it holds that $\mathcal{R}(N, \bar{A}, \bar{C})$ is proportional.*

Although it is a rather weak proportionality requirement (similar to weak proportionality in the apportionment setting (Balinski and Young 1982)), it is sufficiently strong to reveal that some perpetual voting rules are not proportional.

Proposition 1. *AV, Perpetual Equality, Perpetual Reset, and Perpetual Unit-Cost fail simple proportionality.*

Theorem 1. *Perpetual PAV, Perpetual Consensus, Perpetual Nash, and Perpetual Quota satisfy simple proportionality.*

The second axiom concerns the impact of uncontroversial decisions, i.e., decisions where a choice can be made with unanimous agreement. Such uncontroversial decisions should not have an impact on other decisions, as otherwise the inclusion of such decisions could be used to manipulate the outcome of the decision process. Formally, an approval profile A is *uncontroversial due to c* if $\bigcap_{v \in N} A(v) = \{c\}$. Furthermore, given a k -tuple $L = (l_1, \dots, l_k)$ and $i \in \{0, \dots, k\}$, we write $L \oplus_i x$ to denote the $(k+1)$ -tuple $L = (l_1, \dots, l_i, x, l_{i+1}, \dots, l_k)$.

Definition 3 (Independence of uncontroversial decisions). *A perpetual voting rule \mathcal{R} satisfies independence of unanimous decisions if for any k -decision sequence (N, \bar{A}, \bar{C}) , approval profile A for C that is uncontroversial due to c , and $i \in \{0, \dots, k\}$ it holds that*

$$\mathcal{R}(N, \bar{A} \oplus_i A, \bar{C} \oplus_i C) = \mathcal{R}(N, \bar{A}, \bar{C}) \oplus_i c.$$

Proposition 2. *Perpetual PAV, Perpetual Nash, and Perpetual Reset fail independence of unanimous decisions.*

Theorem 2. *AV, Perpetual Equality, Perpetual Quota, Perpetual Unit-Cost, and Perpetual Consensus satisfy independence of unanimous decisions*

The third axiomatic property is called *bounded dry spells*. This property guarantees that every voter agrees with at least some choice in a bounded number of rounds.

Definition 4 (Dry spells). *Given a k -decision history $(N, \bar{A}, \bar{C}, \bar{w})$, we say that a voter $v \in N$ has a dry spell of length ℓ if there exists $t \leq k - \ell$ such that $sat_t(v, \bar{w}) = sat_{t+\ell}(v, \bar{w})$, i.e., voter v is not satisfied by any choice in rounds $t+1, \dots, t+\ell$.*

Let d be function from \mathbb{N} to \mathbb{N} . A perpetual voting rule \mathcal{R} has a dry spell guarantee of d if for any decision sequence (N, \bar{A}, \bar{C}) and $\bar{w} = \mathcal{R}(N, \bar{A}, \bar{C})$, no voter has a dry spell of length $d(|N|)$. A perpetual voting rule \mathcal{R} has bounded dry spells if \mathcal{R} has a dry spell guarantee for some function d .

Equivalently, we can say that a perpetual voting rule has unbounded dry spells if for some fixed N dry spells of arbitrary length can occur.

Proposition 3. *AV, Perpetual PAV, Perpetual Equality, Perpetual Quota, Perpetual Nash, and Perpetual Unit-Cost have unbounded dry spells.*

	Simple Proport.	Indep. of Unanim. Dec.	Bounded Dry Spells
AV	×	✓	×
Per. PAV	✓	×	×
Per. Equality	×	✓	×
Per. Quota	✓	✓	×
Per. Nash	✓	×	×
Per. Reset	×	×	✓
Per. Unit-Cost	×	✓	×
Per. Consensus	✓	✓	✓

Table 1: Overview of the axiomatic analysis

Two of the considered rules have bounded dry spells. The first is Perpetual Reset, for which we can show a guarantee that is linear in n .

Theorem 3. *Perpetual Reset has a dry spell guarantee of $2n-2$. This guarantee is tight, i.e., dry spells of length $2n-3$ may occur.*

For Perpetual Consensus we show a quadratic bound. It remains an open question whether the statement can be strengthened to show a linear dry spell guarantee.

Theorem 4. *Perpetual Consensus has a dry spell guarantee of at least $\frac{n^2+3n}{4}$.*

A summary of our axiomatic analysis can be found in Table 1. Note that Perpetual Consensus is the only rule that satisfies all three properties, and Perpetual Quota the only that satisfies two.

5 Quantitative Properties

We now introduce two metrics that capture certain aspects of desirable behavior of perpetual voting rules. The first one, perpetual lower quota compliance, measures to which degree voters receive at least a “fair share” of favorable choices, and thus have a fair say in the decision process. The second one, the Gini influence coefficient, takes the opposite perspective: it measures whether some voters have a disproportionately large influence on decisions, and thus exert more power than can be justified.

Let us first define perpetual lower quota as a property of choice sequences. The main idea is the following: a voter whose preferences are shared on average by an α fraction of the population should be satisfied on average by an α fraction of the choices.

Definition 5 (Perpetual lower quota). *Let (N, \bar{A}, \bar{C}) be a k -decision sequence. A k -choice sequence $\bar{w} \in \bar{C}$ satisfies perpetual lower quota if for every voter $v \in N$ it holds that*

$$sat_k(v, \bar{w}) \geq \lfloor qu_k(v) \rfloor.$$

As it turns out, this property is too strong to be used as an axiomatic property, i.e., to require that perpetual voting rules produce choice sequences satisfying perpetual lower quota.

Proposition 4. *There are decision sequences for which no choice sequence exists that satisfies perpetual lower quota.*

Proof. Let (N, \bar{A}, \bar{C}) be a k -decision sequence with $N = \{1, \dots, 2^k\}$ and $C_1 = \dots = C_k = \{a, b\}$. We construct \bar{A} as follows: Each voter $v \in N$ corresponds to a length- k string s_v over the alphabet $\{a, b\}$. For each round $j \in [k]$, we define $A_j(v) = \{s_v(j)\}$. Let \bar{w} be an arbitrary k -choice sequence. Then there exists a voter $u \in N$ with $w_j \notin A_j(u)$ for every $j \in [k]$, i.e., $sat_k(u, \bar{w}) = 0$. However, $qu_k(u) = k/2$ and hence perpetual lower quota is not satisfied (by an arbitrarily large margin). \square

As the underlying idea of perpetual lower quota nonetheless appears desirable, we frame it as a quantitative metric that should be guaranteed for as many voters as possible.

Definition 6. *Let $(N, \bar{A}, \bar{C}, \bar{w})$ be a k -decision history. The perpetual lower quota compliance of \bar{w} , $compl(\bar{w})$, is the average proportion of voters in each round that have their perpetual lower quota satisfied, i.e.,*

$$compl(\bar{w}) = \frac{1}{nk} \sum_{i=1}^k |\{v \in N : sat_i(v, \bar{w}) \geq \lfloor qu_i(v) \rfloor\}|.$$

A perpetual lower quota compliance of 1 means that each initial segment of \bar{w} , i.e., (w_1) , (w_1, w_2) , (w_1, w_2, w_3) , etc., satisfies perpetual lower quota.

The second metric is the Gini influence coefficient. It is derived from the Gini coefficient of discrete probability distributions, see e.g. (Moulin 1991). The Gini coefficient is a metric of inequality (often used for income distributions); it is 0 for completely equal distributions and 1 for maximally unequal distributions. We use the Gini coefficient to capture inequality in voters’ influence on the decision process.

We define the influence of a voter on a given choice as 1 divided by the number of voters supporting this choice. For example, if a voter has an influence of 1 on a choice that everyone but her disagrees with; if a choice is supported by all n agents, then their (individual) influence is $\frac{1}{n}$. In the following definition, $\mathbb{I}_{w_j \in A_j(v)}$ is 1 if choice w_j satisfies voter v , and 0 otherwise.

Definition 7. *Let $(N, \bar{A}, \bar{C}, \bar{w})$ be a k -decision history. The influence of voter $v \in N$ on the choice sequence \bar{w} is*

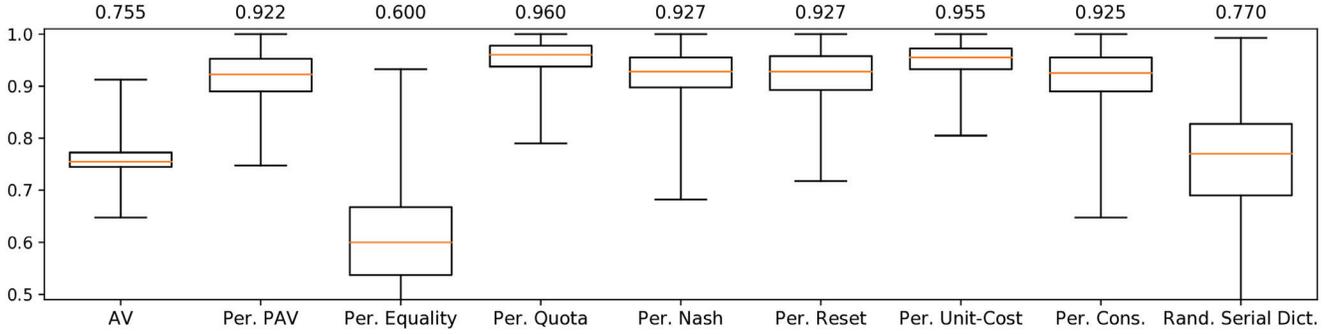
$$infl_k(v, \bar{w}) = \sum_{j \in [k]} \frac{\mathbb{I}_{w_j \in A_j(v)}}{|\{u \in V : A_j(u) = w_j\}|}.$$

Let a be the average influence, i.e., $a = \frac{1}{|N|} \sum_{v \in N} infl_k(v)$. The Gini influence coefficient of \bar{w} , $gini_k(\bar{w})$, is defined as the Gini coefficient of the sequence $(infl_k(v, \bar{w}))_{v \in N}$, i.e.,

$$gini_k(\bar{w}) = \frac{1}{2a|N|^2} \sum_{u \in N} \sum_{v \in N} |infl_k(u, \bar{w}) - infl_k(v, \bar{w})|.$$

We have claimed that perpetual lower quota compliance and the Gini influence coefficient capture two different perspectives. Indeed, the following examples show that there are instances where the Gini influence coefficient is small (that is, an equal distribution of influence), but almost all voters have lower quota violations. Conversely, there are instances where the Gini influence coefficient is large, but all voters have their perpetual lower quota satisfied.

Figure 1: Perpetual lower quota compliance (values on top of the diagram are medians)



Example 1. Let (N, \bar{A}, \bar{C}) be a n -decision sequence with $N = \{1, \dots, n\}$, $C = \{c_1, c_2\}$, and $A_i(v) = \{c_1\}$ if $v = i$ and $A_i(v) = \{c_2\}$ otherwise. The n -choice sequence $\bar{w} = (c_1, \dots, c_1)$ achieves a perfect Gini influence coefficient of 0, since every voter has an influence of 1. The quota of each voter v is $qu_n(v) = (n-1) \cdot \frac{n-1}{n} + \frac{1}{n} \geq n-2$, but $sat_n(v, \bar{w}) = 1$ for all voters v .

Example 2. Let (N, \bar{A}, \bar{C}) be a n -decision sequence with $N = \{1, \dots, n\}$ and $C = \{c_1, \dots, c_n\}$. For $v \in N$, we set $A_1(v) = \dots = A_{n-1}(v) = \{c_v\}$ and $A_n(v) = \{c_1\}$. Consider the n -choice sequence $\bar{w} = (c_1, \dots, c_1)$. Its Gini influence coefficient is $\frac{(n-1)^2}{n^2} \geq 1 - \frac{2}{n}$ (i.e., we have an unequal distribution of influence), but perpetual lower quota is satisfied for all voters: the quota of each voter v is $qu_n(v) = \frac{n-1}{n} + 1$, and $sat_n(v, \bar{w}) \geq 1$.

For simple n -decision sequences, however, these two metrics coincide.

Proposition 5. Let (N, \bar{A}, \bar{C}) be a simple n -decision sequence with $n = |N|$. An n -choice sequence $\bar{w} \in \bar{C}$ satisfies perpetual lower quota if and only if $gini_k(\bar{w}) = 0$.

As we will see in the following experiments, these two metrics indeed highlight different aspects and yield different evaluations of rules.

6 Experiments

The goal of our experiments is to test the proposed rules against the two performance metrics introduced in the previous section. In contrast to the binary statements of an axiomatic analysis, an experimental analysis via numerical simulations gives us an understanding about the average-case behavior. Numerical experiments have to be based on sensibly chosen distributions that model the envisioned application. In our setting these are small- to medium-sized groups with conflicting preferences. For the sake of conciseness, we only present one distribution (other distributions yielded mostly similar results, as described later). We consider a set of 20 voters which decide upon 20 decision instances, i.e., we have 20-decision sequences. For each decision 5 alternatives are available—these differ from round to round. We generate voters and alternatives in a two-dimensional Euclidean space, similar to the setup used by

Elkind et al. (2017). Voters are split in two groups and are placed on the 2d plane by a bivariate normal distribution. For the first group (6 voters) both x - and y -coordinates are independently drawn from $\mathcal{N}(-0.5, 0.2)$; for the second group (14 voters) x - and y -coordinates are from $\mathcal{N}(0.5, 0.2)$. That is, the first, smaller group is centered around $(-0.5, -0.5)$, the second, larger group around $(0.5, 0.5)$.

Alternatives are distributed uniformly in the rectangle $[-1, 1] \times [-1, 1]$. Voters approve all alternatives that have a distance of at most 1.5 times the distance to the closest alternative. This yields approval sets of size 1.8 on average. It is important to note that alternatives change in every round and thus even voters that are close to each other do not necessarily have the same approval sets each round.

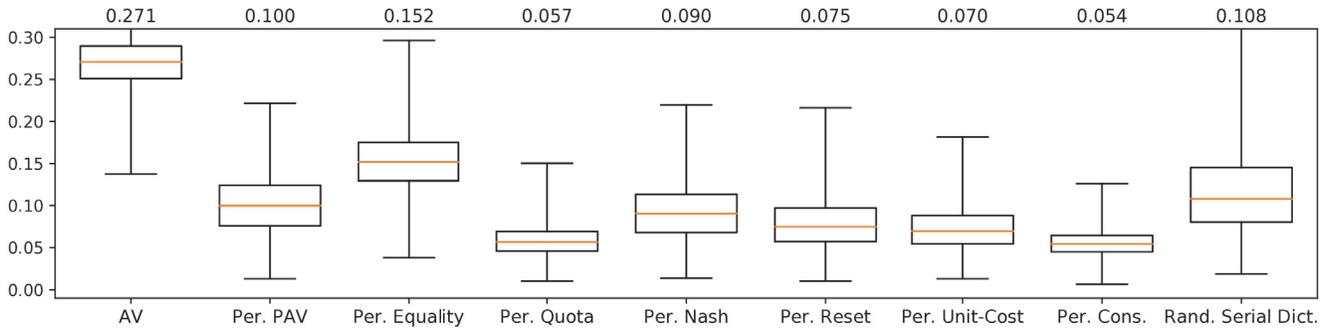
In addition to the perpetual voting rules proposed in Section 3, we also include a probabilistic voting rule: Random Serial Dictatorship (RSD). This rule works as follows. In each round, a permutation of voters is selected uniformly at random. We maintain a set X that starts as the set of all alternatives (in this round). One voter after the other can shrink X further to include only approved alternatives; the set X remains unchanged by voters whose approval set has an empty intersection with X . As soon as X has cardinality 1, this alternative is chosen. If $|X| > 1$ after all voters are considered, one alternative in X is chosen at random.

Our results are based on 10,000 instances. For each instance and each voting rule, we compute the perpetual lower quota compliance and the Gini influence coefficient⁴. The corresponding distributions are visualized in Figures 1 and 2. These box plots show as boxes the first (bottom), second (middle), and third quartile (top), as well as the minimum and maximum values (bottom and top whiskers). Note that for perpetual lower quota compliance values close to 1 are desirable, for the Gini influence coefficient values close to 0. All apparent differences between rules (Figs. 1 and 2) are statistically significant (paired t-test, $p = 0.01$).

Let us first consider perpetual lower quota compliance (Figure 1). We see that three rules perform significantly worse than the others: AV, Perpetual Equality, and RSD. It is unsurprising that AV and Perpetual Equality do not per-

⁴The Python code used for these experiments can be found at <https://github.com/martinlackner/perpetual>.

Figure 2: Gini influence coefficient (values on top of the diagram are medians)



form well, as both—by definition—do not aim to be proportional. For RSD, however, it is noteworthy that perpetual voting rules have significantly better proportionality guarantees. This might be due to the rather short decision sequence (which we believe is a realistic assumption). Among the well-performing rules, Perpetual Quota and Perpetual Unit-Cost stand out as granting the most voters their lower quota.

Now let us look at the Gini influence coefficient (Figure 2). Here, again, AV, Perpetual Equality, and RSD yield unsatisfactory results. But also, perhaps surprisingly, Perpetual PAV and Perpetual Nash do not fare well with regard to this metric. In contrast, Perpetual Consensus and Perpetual Quota achieve the best results.

As mentioned before, we have repeated these experiments also for other distributions (more than two groups, differing variance, larger approval sets, other probability distributions, more voters, longer decision sequences, etc.). The general outcome was that Perpetual Quota performed best for most distributions. The subpar performance of AV, Perpetual Equality, and RSD could also be observed in almost all experiments. Only the exact comparison between the other perpetual rules differed significantly. The main conclusion stands, however, that Perpetual Quota is most recommendable from the perspective of both metrics.

7 Discussion

In this paper we presented the idea of perpetual voting and introduced several perpetual voting rules. In our analysis, two rules stood out as particularly promising: Perpetual Consensus (as the only considered rule that satisfies all three properties studied in Section 4) and Perpetual Quota (with a particularly strong performance in the experiments, Section 6, and satisfying two of three axiomatic properties). Further work is required to explore the design space of perpetual voting rules and obtain a clear understanding of the rules’ properties and behavior. The literature of multiwinner voting can offer inspiration for further perpetual voting rules, e.g., rules could be designed based on the implicit utilities approach by Boutilier et al. (2015) or based on Phragmén’s sequential rule (Brill et al. 2017).

In our numerical simulations, we came to the conclusion that Random Serial Dictatorship (RSD), a probabilistic vot-

ing rule, cannot be recommended in the considered settings. We furthermore note that out of the three axioms studied in Section 4, RSD satisfies only independence of unanimous decisions. RSD does not satisfy simple proportionality due to its probabilistic nature, but it converges to proportional outcomes. Similarly, RSD does not provide a strict guarantee for dry spells. Thus, for settings with *short* decision sequences, we believe that probabilistic voting rules are not a suitable alternative to perpetual voting rules.

We used two notions of proportionality: simple proportionality (Definition 2) and perpetual lower quota (Definition 5). The former is a rather weak axiom, whereas the latter is too strong to be satisfiable in general. It thus of interest to explore the middle ground and pinpoint which degree of proportionality is achievable by perpetual voting rules. Inspiration for this line of work can come from concepts such as extended or proportional justified representation (Aziz et al. 2017; Sánchez-Fernández et al. 2017), both of which have been developed as simpler, seemingly more natural notions of proportionality proved to be unattainable.

Let us end this paper with a discussion of practical challenges pertaining to perpetual voting. First, we have not yet addressed the issue of fluctuating voters. How should weights be adapted if voters abstain some decisions or enter the decision process at a later stage? This question has to be answered individually for each rule and more than one reasonable answer may exist for each rule. It is important to ensure that fairness guarantees also hold for newcomers, but that abstainers do not have an unfair advantage neither.

This leads us immediately to the issue of strategic behavior. In perpetual voting one may encounter the *freerider effect*: if an alternative is guaranteed to be chosen, it is beneficial for voters to misrepresent their preferences (or abstain) so as not to pay the “price of winning”. This is a problem inherent in the idea of guaranteeing fairness over time (cf. Freeman, Zahedi, and Conitzer (2017)). However, its severity may differ between perpetual voting rules. For example, Perpetual Reset is an extreme case where the price of winning is very high, whereas Perpetual Quota has a small price of winning if an alternative with strong support is chosen. It requires further work to understand the impact of the freerider effect on different rules and find strategies to com-

but its negative consequences.

Finally, to find and agree on a compromise is a central feature of human interaction and negotiation. Compromise is also sought in classical voting scenarios, although the possibilities are limited. In a long-term decision process, compromise becomes a powerful concept. For example, if agents assign different importance to individual decisions, compromise can be found by deciding in favor of agents that consider the issue at hand critical, while assigning a higher priority for future decisions to agents that “yielded”. The goal is here to augment perpetual voting rules with methods of compromise and thus increase their ability to aid in a real-world decision process.

Acknowledgements

The author would like to thank Piotr Faliszewski, Marie-Louise Lackner, and Jan Maly for valuable feedback and discussions. This work was supported by the Austrian Science Fund (FWF): P31890.

References

- Aziz, H.; Brill, M.; Conitzer, V.; Elkind, E.; Freeman, R.; and Walsh, T. 2017. Justified representation in approval-based committee voting. *Social Choice and Welfare* 48(2):461–485.
- Balinski, M., and Young, H. P. 1982. *Fair Representation: Meeting the Ideal of One Man, One Vote*. Yale University Press.
- Boutillier, C., and Procaccia, A. D. 2012. A dynamic rationalization of distance rationalizability. In *Proceedings of the 26th Conference on Artificial Intelligence (AAAI-2012)*. AAAI Press.
- Boutillier, C.; Caragiannis, I.; Haber, S.; Lu, T.; Procaccia, A. D.; and Sheffet, O. 2015. Optimal social choice functions: A utilitarian view. *Artificial Intelligence* 227:190–213.
- Brandt, F. 2017. Rolling the dice: Recent results in probabilistic social choice. In Endriss, U., ed., *Trends in Computational Social Choice*. AI Access. chapter 1, 3–26.
- Brill, M.; Freeman, R.; Janson, S.; and Lackner, M. 2017. Phragmén’s voting methods and justified representation. In *Proceedings of the 31st Conference on Artificial Intelligence (AAAI-2017)*, 406–413. AAAI Press.
- Brill, M.; Laslier, J.-F.; and Skowron, P. 2018. Multiwinner approval rules as apportionment methods. *Journal of Theoretical Politics* 30(3):358–382.
- Casella, A. 2005. Storable votes. *Games and Economic Behavior* 51(2):391–419.
- Casella, A. 2012. *Storable Votes: Protecting the Minority Voice*. Oxford University Press.
- Chamberlin, J. R., and Courant, P. N. 1983. Representative deliberations and representative decisions: Proportional representation and the Borda rule. *American Political Science Review* 77(3):718–733.
- Dickerson, J. P.; Procaccia, A. D.; and Sandholm, T. 2012. Dynamic matching via weighted myopia with application to kidney exchange. In *Proceedings of the 26th Conference on Artificial Intelligence (AAAI-2012)*. AAAI Press.
- Elkind, E.; Faliszewski, P.; Laslier, J.-F.; Skowron, P.; Slinko, A.; and Talmon, N. 2017. What do multiwinner voting rules do? An experiment over the two-dimensional euclidean domain. In *Proceedings of the 31st Conference on Artificial Intelligence (AAAI-2017)*, 494–501. AAAI Press.
- Faliszewski, P.; Skowron, P.; Slinko, A.; and Talmon, N. 2017. Multiwinner voting: A new challenge for social choice theory. In Endriss, U., ed., *Trends in Computational Social Choice*. AI Access. chapter 2, 27–47.
- Freeman, R.; Zahedi, S. M.; Conitzer, V.; and Lee, B. C. 2018. Dynamic proportional sharing: A game-theoretic approach. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 2(1):3:1–3:36.
- Freeman, R.; Zahedi, S. M.; and Conitzer, V. 2017. Fair and efficient social choice in dynamic settings. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI-2017)*, 4580–4587. ijcai.org.
- Hemaspaandra, E.; Hemaspaandra, L. A.; and Rothe, J. 2017. The complexity of controlling candidate-sequential elections. *Theoretical Computer Science* 678:14–21.
- Hosseini, H.; Larson, K.; and Cohen, R. 2015. Matching with dynamic ordinal preferences. In *Proceedings of the 29th Conference on Artificial Intelligence (AAAI-2015)*, 936–943. AAAI Press.
- Kilgour, D. M. 2010. Approval balloting for multi-winner elections. In Laslier, J.-F., and Sanver, R., eds., *Handbook on Approval Voting*. Springer. 105–124.
- Lang, J., and Xia, L. 2009. Sequential composition of voting rules in multi-issue domains. *Mathematical Social Sciences* 57(3):304–324.
- Lang, J., and Xia, L. 2016. Voting in combinatorial domains. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*. Cambridge University Press.
- Moulin, H. 1991. *Axioms of Cooperative Decision Making*. Cambridge University Press.
- Oren, J., and Lucier, B. 2014. Online (budgeted) social choice. In *Proceedings of the 28th Conference on Artificial Intelligence (AAAI-2014)*, 1456–1462. AAAI Press.
- Parkes, D. C., and Procaccia, A. D. 2013. Dynamic social choice with evolving preferences. In *Proceedings of the 27th Conference on Artificial Intelligence (AAAI-2013)*. AAAI Press.
- Sánchez-Fernández, L.; Elkind, E.; Lackner, M.; Fernández, N.; Fisteus, J. A.; Val, P. B.; and Skowron, P. 2017. Proportional justified representation. In *Proceedings of the 31st Conference on Artificial Intelligence (AAAI-2017)*, 670–676. AAAI Press.
- Tennenholtz, M. 2004. Transitive voting. In *Proceedings of the 5th ACM Conference on Electronic Commerce (ACM-EC-2004)*, 230–231. ACM.
- Thiele, T. N. 1895. Om flerfoldsvalg. In *Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandlinger*. 415–441.