

What Voting Rules Actually Do: A Data-Driven Analysis of Multi-Winner Voting

Joshua Caiata¹, Ben Armstrong^{*2}, and Kate Larson¹

¹University of Waterloo

²Tulane University

jcaiata@uwaterloo.ca, research@benarmstrong.ca,
kate.larson@uwaterloo.ca

Abstract

Committee-selection problems arise in many contexts and applications, and there has been increasing interest within the social choice research community on identifying which properties are satisfied by different multi-winner voting rules (e.g. [31]). In this work, we propose a data-driven framework to evaluate how frequently voting rules violate axioms across diverse preference distributions in practice, shifting away from the binary perspective of axiom satisfaction given by worst-case analysis. Using this framework, we analyze the relationship between multi-winner voting rules and their axiomatic performance under several preference distributions. We then show that neural networks, acting as voting rules, can outperform traditional rules in minimizing axiom violations. Our results suggest that data-driven approaches to social choice can inform the design of new voting systems and support the continuation of data-driven research in social choice.

1 Introduction

Committee selection is a central problem in social choice theory, wherein voters elect a committee (subset of alternatives) based on their preferences [31, 21]. There are numerous properties or axioms we might wish a multi-winner voting rule to satisfy, however many combinations of properties are known to be impossible to satisfy simultaneously. Traditional research often focuses on ascertaining which axiomatic combinations are possible or are satisfied by a voting rule. Such research is motivated by a desire to identify rules which universally satisfy axioms deemed desirable for a particular setting. However, in many instances, a voting rule may not satisfy an axiom but rarely violates it in practice.

In this paper, we propose a data-driven framework to evaluate and explore voting rules axiomatically. To do so, we establish a measure of axiom violation more fine-grained than binary satisfaction and move away from worst-case analysis and towards an average case evaluation model. Using this framework, we explore the relationship between underlying voter preference distributions and axiomatic properties selected by common multi-winner voting rules using our average-case approach. We then use machine learning to learn bespoke multi-winner voting rules that aim to minimize

^{*}This work was primarily completed while affiliated with the University of Waterloo.

axiom violations as much as possible when electing a committee. These multi-layer perceptrons are designed to select a set of alternatives based on preferences that violate certain axioms less frequently than many well-known rules.

Specifically, we make the following contributions:

1. We initiate a data-driven measure of axiomatic violations for multi-winner voting rules.
2. We explore the relationship between the underlying preference distributions and the axiomatic properties of committees selected by common multi-winner voting rules, illustrating the sensitivity of many of these rules to the underlying voter population.
3. We empirically investigate how different multi-winner voting rules differ from each other both in terms of the committees they select and the frequency with which they violate (sets of) axioms.
4. We demonstrate that it is possible to use machine learning to discover novel multi-winner voting rules that outperform existing rules under our evaluation framework.

In summary, this paper illustrates the power of using data-driven approaches and machine learning to deepen our understanding of social choice and inform the design of new voting systems.

1.1 Related Work

Much prior work has both developed novel axioms to describe desirable properties, or shown that particular axioms are satisfied by certain voting rules. Of relevance to our work, Elkind et al. develop axioms for several types of multi-winner voting rules and establishes satisfaction results [18]. They note, in particular, the difficulty of satisfying Dummett’s condition (where if a large enough group of voters agree on a set of alternatives as their top choices, then those alternatives should be in the winning committee) [15]. A similar approach is taken for rules and axioms based on approval preferences [31, 37].

Our approach builds on data-driven axiomatic analysis [14, 20], drawing upon axioms from the social choice literature. Our results vary greatly based on underlying voter preferences; these are well-studied for the single-winner setting [17] but less is known for the multi-winner setting. We use well-studied distributions shown to approximate human preferences or explore restricted cases [6]. Recent work has highlighted differences in the winners of multi-winner voting rules on generated and real-world preference data, emphasizing how some rules elect committees which are quite different from each other, particularly Minimax Approval Voting, Chamberlin-Courant, and sequential Chamberlin-Courant [22].

Other recent work has explored the possibility of using machine learning with social choice. Prior work has primarily focused on approximating single-winner or probabilistic rules [34, 10, 30]. Existing work on learning new voting rules, primarily in single-winner settings, has studied learning rules under specific axioms [3], reducing susceptibility to manipulation [23], or maximizing utility without axiomatic focus [1]. One paper has explored the possibility of learning multi-winner rules for participatory budgeting with a focus predominantly on measures of social welfare [20]. Our work, instead, is focused on learning multi-winner rules aimed at minimizing axiom violations.

2 Preliminaries: Social Choice Building Blocks

Traditional research in social choice is typically based upon three cornerstones: voting rules, voter preferences, and some measure of outcome quality. Voting rules aggregate voter preferences following

some particular procedure aimed at maximizing some measure of quality. Rules may provide a single winner, a ranking over alternatives, or a set of multiple winners. Voter preferences are the subjects of aggregation; they are often modelled as originating from some fixed distribution or based upon data from real-world voting applications. Preferences might be expressed as rankings over alternatives, approvals, or even as raw scores. Measures of voting rule quality primarily fall into one of two paradigms: *social welfare*, or *axiomatic*.

In this work we consider voting rules which elect multiple winners, using either ordinal or approval-based preferences sampled from a wide variety of distributions, and focus exclusively on a novel measure of the axiomatic properties of voting rules. We now provide the basic notation used through the paper and describe each of these three cornerstones in more detail.

Model and Notation

Let V be a set of n voters and M be a set of m alternatives. Each voter $v_i \in V$ has a preference ranking over M where for $a_i, a_j \in M$, $a_i \succ_v a_j$ means that voter $v \in V$ prefers alternative a_i to a_j . A *preference profile*, $P_{\succ} = (\succ_{v_1}, \dots, \succ_{v_n})$ is a vector specifying the preferences of each voter. In addition to voters' preferences, we are also interested in the alternatives that a voter approves. Let $App(v) \subset M$ be the *approval set* of voter $v \in V$, containing the k most preferred alternatives of v . We let $P_{App} = (App(v_1), \dots, App(v_n))$, and when it is clear from the context we will abuse notation and let P refer to either P_{\succ} or P_{App} . An election, $E = (V, M)$, is defined by its voters and alternatives, along with, implicitly, either preference profiles or approval sets of the voters.

We are interested in *multi-winner* voting rules. Given an election E and its associated P and k , $1 \leq k < m$, $\mathcal{F}_k(E) \subseteq \{C | C \subset M, |C| = k\}$ is a multi-winner voting rule that returns a family of k -sized subsets of M , called the winning *committees*. If \mathcal{F}_k directly uses P_{\succ} we call the voting rule *ordinal*. If \mathcal{F}_k uses P_{App} we say the voting rule is *approval-based*. We are interested in *resolute* multi-winner voting rules and so assume that each rule uses some tie-breaking mechanism so that only one committee is returned. Unless otherwise stated we use lexicographic tie-breaking, sorting on the first non-shared alternative between committees.

2.1 Multi-Winner Voting Rules

We consider multi-winner voting rules from the existing literature which are broadly classified into two categories: ordinal-based and approval-based. While much research focuses exclusively on just one of either ordinal- or approval-based rules we intentionally include rules from both categories in our study. This allows us to highlight behavioural differences among rules which are seldom compared directly.

The ordinal rules which we explore consist of:




Figure 1: An illustrated overview of the principle concepts in this paper. We apply an analytical framework to identify relationships between each of three fundamental cornerstones of social choice.

k -Borda: Each voter assigns $m - 1$ points to their top ranked alternative, $m - 2$ points to their second ranked alternative etc. $\mathcal{F}^{\text{Borda}}$ returns the k alternatives with highest scores.

Single Non-Transferable Vote: Each voter assigns one point to their most preferred alternative. $\mathcal{F}^{\text{SNTV}}$ returns the k alternatives with the highest points.¹

Single Transferable Vote: \mathcal{F}^{STV} recursively adds to the set of winners all alternatives ranked first by voters with weight summing to more than a quota of $\frac{n}{k+1}$. Each voter begins with a weight of 1. When a winner, a_w is found, each voter ranking a_w first has their weight reduced by an amount proportional to the excess weight voting for a_w beyond the quota. If there is not an alternative ranked first by a quota of weighted voters, the alternative ranked first by lowest sum of voter weight is removed from all voter rankings [41].

We also consider the following approval-based rules:

Bloc: Define a score function $sc_{\text{Bloc}}(a) = \sum_{v \in V} \mathbf{1}_{\{a \in \text{App}(v)\}}$ giving one point to each alternative for each approval it receives. $\mathcal{F}^{\text{Bloc}}$ returns the committee C^* containing the k alternatives receiving the most approvals.

$$C^* = \arg \max_{\substack{C \subseteq M, \\ |C|=k}} \sum_{a \in C} sc_{\text{Bloc}}(a)$$

Proportional Approval Voting: Given committee $C \subseteq M, |C| = k$, define

$$sc_{\text{PAV}}(C) = \sum_{v \in V} \sum_{j=1}^{|C \cap \text{App}(v)|} \frac{1}{j}.$$

\mathcal{F}^{PAV} returns $C^* = \arg \max sc_{\text{PAV}}(C)$. Each voter contributes some score to each committee based on the number of alternatives in the committee of which they approve; a voter approving of r alternatives in C adds the r -th harmonic number to the score of the committee.

Chamberlin-Courant (CC): We consider three variations of the Chamberlin-Courant rule. For committee $C \subseteq M, |C| = k$, define

$$sc_{\text{CC}}(C) = \sum_{a \in C} \sum_{v \in V} \mathbf{1}_{\{|C \cup \text{App}(v)| > 0\}}.$$

$sc_{\text{CC}}(C)$ is the number of voters who approve at least one alternative in C .

Approval CC is defined as $\mathcal{F}^{\text{CC}} = C^*$ where $C^* = \arg \max sc_{\text{CC}}(C)$.

Lexicographic CC, $\mathcal{F}^{\text{lex-CC}}$, maximizes sc_{CC} , breaking ties by selecting the committee maximizing the number of voters approving 2 alternatives (then, if ties remain, 3 alternatives, etc.).

Sequential CC, seq-CC, constructs a winning committee by iteratively adding $a \in M$ that increases the sc_{CC} the most at each step.

¹SNTV is an extension of plurality to the multi-winner context.

Monroe: Considers all ways of assigning each voter one alternative in committee W , such that every $a \in W$ is assigned to between $\lfloor \frac{n}{k} \rfloor$ and $\lceil \frac{n}{k} \rceil$ voters. The score of an assignment is the sum of Borda scores of each voter’s assigned alternative. $\mathcal{F}^{\text{Monroe}}$ selects the committee with the highest score.

Greedy Monroe, $\mathcal{F}^{\text{Greedy M.}}$, constructs a winning committee iteratively following a similar scoring process².

Minimax Approval: \mathcal{F}^{MAV} selects the committee that minimizes the maximum Hamming distance between any voter’s approved alternatives and the committee.

Method of Equal Shares: \mathcal{F}^{MES} has two phases. First, each voter has a budget of $\frac{k}{n}$. Proceed for up to k rounds: Adding a voter to the committee has a cost of 1, which can be split between many voters. In each round, consider alternatives A_r which are not in the committee and are approved of by voters that have a remaining budget summing to at least 1. If A_r is empty, go to phase 2. Otherwise, select $a \in A_r$ such that each voter approving of a must spend at most ρ to add them to the committee. Add a to the committee, adjust the remaining budget of each voter, and proceed to the next round. In the second phase, many possible rules can be used to fill any remaining spots on the committee. We use the sequential Phragmen rule. MES satisfies the JR and PJR axioms (but not the Core). See Lackner and Skowron for further details [31].

E Pluribus Hugo: Also called “Single Divisible Vote with Least-Popular Elimination”; \mathcal{F}^{EPH} operates in rounds. In each round, each voter divides a single point evenly between all remaining alternatives of which they approve. Alternatives are ranked in order of total summed points from all voters. The two alternatives with the lowest number of points are compared: the one receiving the fewest approvals overall is eliminated. Rounds of elimination continue until k alternatives remain. \mathcal{F}^{EPH} satisfies the Strong Pareto Efficiency axiom [38].

Random Serial Dictator: \mathcal{F}^{RSD} selects a single voter to serve as “dictator”. The winning committee is exactly the set of alternatives approved of by that voter.

Faliszewski et al. describe three categories for multi-winner voting rules: *individual excellence* (electing alternatives which are individually well-liked), *diversity* (electing alternatives which are different from each other), and *proportionality* (electing a committee which proportionally represents the preferences of voters). The rules we use are generally aligned with one or two of these categories [21, 22]:

Individual Excellence: $\mathcal{F}^{\text{Borda}}$, $\mathcal{F}^{\text{SNTV}}$, $\mathcal{F}^{\text{Bloc}}$, \mathcal{F}^{EPH}

Diversity: $\mathcal{F}^{\text{SNTV}}$, \mathcal{F}^{CC}

Proportionality: \mathcal{F}^{STV} , \mathcal{F}^{PAV} , $\mathcal{F}^{\text{Monroe}}$, \mathcal{F}^{CC} , \mathcal{F}^{MES} , \mathcal{F}^{EPH}

We categorize the variants of each rule in the same way as the original rule. We highlight that these are subjective categorizations and are not mutually exclusive. For example, \mathcal{F}^{CC} has been described as both *diverse* [22] and *proportional* [18]. For this reason we generally group the two categories together. Additionally, \mathcal{F}^{MAV} and \mathcal{F}^{RSD} do not neatly fit into any category (\mathcal{F}^{MAV} considers alternatives as sets, rather than individuals, but does not obviously aim to achieve diversity or proportionality while \mathcal{F}^{RSD} considers only a single voter’s opinion).

²We refer the reader to Lackner and Skowron for a formal definition [31].

2.2 Voting Rule Axioms

Much of the literature on voting rules is axiomatic, describing desirable properties that voting rules may or may not exhibit. We focus exclusively on *intraprofile* axioms – axioms for which we can determine a violation using only the preference profile being given to \mathcal{F} and the resulting committee [39]. Axioms can also be loosely categorized based on the stated priorities of their definitions.

We use the following axioms which describe aspects of *individual excellence*:

Majority Winner If $\lceil \frac{n}{2} \rceil$ or more voters rank alternative x first in their ballot, then x is in the winning committee [24].

Majority Loser If $\lceil \frac{n}{2} \rceil$ or more voters rank alternative x last in their ballot, then x is not in the winning committee.

Condorcet Winner A Condorcet committee C is one in which for all $x \in C$ and for all $y \in M \setminus C$, the majority of voters prefer x to y . A voting rule satisfies the Condorcet Winner axiom if, whenever one exists, it returns a Condorcet committee [26].

Condorcet Loser Let $L \subseteq M$ such that for all $x \in L$, and $y \notin L$, y is preferred to x by a majority of voters. A voting rule satisfies the Condorcet Loser axiom if it never returns L [26].

Strong Pareto Efficiency Committee C dominates C' if every voter approves at least as many alternatives in C as in C' , and at least one voter approves strictly more alternatives in C than in C' . A multi-winner voting rule exhibits Strong Pareto Efficiency if it never returns a dominated committee [31].

Fixed Majority If there exists a set of alternatives C , $|C| = k$ and a set of voters $X \subseteq V$ with $|X| > \frac{n}{2}$ that all rank each alternative in C above each alternative not in C then the winning committee is C [13, 18].

Strong Unanimity If every voter ranks the same k alternatives on top, then those alternatives form the winning committee [18].

And we use the following axioms that prioritize *diversity* or *proportionality*.

Dummett's Condition If there is a group of $\frac{\ell \cdot n}{k}$ voters that all rank the same ℓ alternatives on top, these ℓ alternatives are in the winning committee [15]. Also referred to as “Proportionality for Solid Coalitions” in some literature [9].

Local Stability For $q = \lceil \frac{n}{k} \rceil$, committee C violates local stability if there exists a subset of voters V^* with $|V^*| \geq q$ and an alternative $x \notin C$ such that every voter in V^* prefers x to all members of C [4]. We highlight in particular that a locally stable committee is *not* guaranteed to exist. That is, it is possible for there to exist a preference profile for which any winning committee would violate local stability.

Solid Coalitions If at least $\frac{n}{k}$ voters rank some alternative x first, then x should be in the winning committee [18].

Core A committee C is “in the core” if for each non-empty subset of $\mathcal{V} \subseteq V$ of voters, and each non-empty subset of alternatives $T \subseteq A$ with

$$\frac{|T|}{k} \leq \frac{|V|}{n}$$

there is a $v_i \in V$ such that $|App(v_i) \cap T| \leq |App(v_i) \cap W|$. That is, v_i approves of at least as many alternatives in W as they approve of in T . A voting rule satisfies the Core if it always returns a committee that is in the core [31].

(Extended) Justified Representation For $\ell \geq 1$, a group of voters is ℓ -cohesive if (1) $|V| \geq \ell \cdot \frac{n}{k}$, and (2) $|\bigcap_{i \in V} App(v_i)| \geq \ell$.

A voting rule satisfies Justified Representation (**JR**) if for a winning committee C , it is always the case that every 1-cohesive group of voters \mathcal{V} contains a v_i who approves of at least one member of C ; i.e. $|W \cap App(v_i)| \geq 1$.

A voting rule satisfies Extended Justified Representation (**EJR**) if for a winning committee C , it is always the case that *every* ℓ -cohesive group of voters \mathcal{V} contains a v_i that approves of at least ℓ winners, for $1 \leq \ell \leq k$ [31].

Table 1 illustrates which existing voting rules are known to satisfy each axiom. Green entries indicate known axiomatic satisfaction for a particular axiom and voting rule.

The axioms we use are only one possible set and were chosen to represent a wide variety of desirable properties. While our framework is general and applies to **any** set of intra-profile axioms, we include a study of the relationships between these axioms in Appendix A. While relationships between some of our axioms are already described in the literature, we also identify several novel relationships where the satisfaction of one axiom implies the satisfaction of another.

2.3 Preference Distributions

We use a wide range of standard preference distributions (\mathcal{D}) through our experiments. These are used both for training novel multi-winner rules, as well as testing novel and existing rules. Note that we generate ordinal ballots and, when needed by a particular rule, convert them to an approval-based format where each voter approves of their k top-ranked alternatives. There are many other valid methods of generating approval ballots which we do not explore.

Through our experiments we consider 8 unique families of preference distribution, as well as two additional sets of preferences. We loosely categorize each of these distributions by the amount of structure within the preferences they generate, from Identity preferences where all voters are identical to Impartial Culture where voters are assigned preferences uniformly at random.

Unstructured Distributions

Impartial Culture: Each voter draws its preference ranking from a uniform distribution over all possible rankings [29].

Impartial Anonymous Culture: Given n voters, each n -element multiset of preference rankings is drawn from a uniform distribution [27].

Moderately Structured Distributions

Mallows: Given reference ranking σ and dispersion parameter θ , the probability of drawing ranking r is $P(r|\sigma, \theta) = \frac{1}{Z} \theta^{d(r, \sigma)}$ where Z is a normalization factor and $d(r, \sigma)$ is the Kendall-Tau distance between rankings r and σ [33, 6].

Urn: Parameterized by α . All $m!$ preference orders exist in an “urn.” Voters select orders consecutively. After a preference order is selected, it is returned to the urn along with $\alpha!$ copies of the ranking [16].

Euclidean: Parameterized by a dimension and a topology. Voters and alternatives are placed randomly within the chosen space and preferences correspond to the distance from each voter to each alternative [19].

Highly Structured Distributions

Identity: All voters in V have identical preferences.

Single-Peaked: A global ordering of alternatives exists; each voter has a favourite alternative and prefers alternatives closer to their favourite over those further [5].

Stratification: Parameterized by a weight $w \in (0, 1)$. Alternatives are split into two classes with the first class’s size proportional to w . All voters rank all alternatives in the first class above those in the second class. Within a class voters rank alternatives uniformly at random [6].

Additional Preference Distributions

Mixed: In addition to preference distributions taken on their own, we consider a distribution consisting of profiles sampled from every other distribution with equal probability.

PrefLib: We include all applicable strict, complete preference profiles from PrefLib, an online repository of human preference data [35]. The amount of PrefLib data applicable to our setting is not sufficient for learning from; we use this data exclusively during evaluation.

3 Data-Driven Voting Rule Analysis

Traditionally, axiomatic analysis of voting rules asks whether a rule satisfies some axiom universally, which is a binary question. In the real world and synthetic preference profiles, however, the gap between worst-case analysis and average-case behaviour can be quite large. Rules that are known to explicitly not satisfy an axiom may very rarely violate such an axiom in practice. The key insight, and motivation, for such a shift in analyzing a voting rule is that a fine-grained and empirical lens reveals real and meaningful information on how rules behave across distributions, which are not considered when only looking at axiomatic satisfaction as a binary question.

We introduce two measures that we use to understand the behaviour of voting rules. The data-driven analysis we employ is very well-suited to exploring deep non-binary measures of axiom satisfaction. We capture this by measuring the *rate* at which the outcome of voting rules violates an axiom. We also study the amount of overlap between committees elected by voting rules. This provides us an indication of how similar two rules in practice, regardless of how they function internally.

While both of our data-driven metrics are affected by each of the three cornerstones of our social choice framework, they are more naturally associated with certain components. The axiom violation rate captures the interplay between axiomatic performance and voting rule behaviour³ while the rule difference measure reveals connections between preference distributions and voting rules⁴.

³We argue that preferences are less intrinsically related to axiom violations than our other cornerstones: in the cases where a voting rule satisfies a particular axiom then the underlying preference distribution has no bearing on the violation rate.

⁴This is because axioms have no direct bearing on the output of any given voting rule – although they may indirectly influence the design of voting rules.

3.1 Axiom Violation Rate

We begin by formalizing the axiom violation rate (AVR) as our core metric for empirical axiomatic performance. As we focus exclusively on *intraprofile axioms*, we can determine whether an axiom is violated by a preference profile using only the voting rule \mathcal{F} and the profile itself.

We now define a measure of the frequency that an axiom is violated by a voting rule on given preference profiles. If axiom A is violated by a specific preference profile P and committee c , we say $A(P, c) = 1$. If A is not violated, $A(P, c) = 0$. Then, we define the axiom violation rate (AVR) of \mathcal{F} on a set of profiles \mathbb{P} over axioms \mathbb{A} as:

$$\text{AVR}(\mathcal{F}, \mathbb{P}, \mathbb{A}) = \frac{1}{|\mathbb{A}||\mathbb{P}|} \sum_{A \in \mathbb{A}} \sum_{P \in \mathbb{P}} A(P, \mathcal{F}(P))$$

3.2 Rule Differences

We also measure overlap between elected committees. This tells us (1) whether rules with similar AVR elect similar underlying committees, and (2) the degree of similarity between committees elected by rules with differing AVRs. We say $\cap_P^+(\mathcal{F}^1, \mathcal{F}^2) = \mathcal{F}^1(P) \cap \mathcal{F}^2(P)$ and $\cap_P^-(\mathcal{F}^1, \mathcal{F}^2) = (M \setminus \mathcal{F}^1(P)) \cap (M \setminus \mathcal{F}^2(P))$. We also define a normalization factor $\delta = \frac{m}{m - |m - 2k|}$ which ensures the difference between two rules on a given set of profiles \mathbb{P} has a range from 0 to 1.

$$\text{d}(\mathcal{F}^1, \mathcal{F}^2, \mathbb{P}) = \delta - \frac{\delta}{m|\mathbb{P}|} \sum_{P \in \mathbb{P}} |\cap_P^+(\mathcal{F}^1, \mathcal{F}^2)| + |\cap_P^-(\mathcal{F}^1, \mathcal{F}^2)|$$

4 Learning Rules to Satisfy Axioms

We now describe \mathcal{F}^{NN} , a novel voting rule with a functionality directly based upon goals rooted in our cornerstones of axiom satisfaction and voter preferences. Functionally, \mathcal{F}^{NN} is multi-layer perceptron trained to predict a winning committee based on a given preference profile. This section outlines the procedure we use to train the model. Specific details on the parameters used to train \mathcal{F}^{NN} in our experiments are included in section 5.

4.1 Generating Axiom Violation Data

We generate separate training sets (used by \mathcal{F}^{NN}) and testing sets (used to evaluate all rules). A single example in a dataset is generated as follows:

1. Given some preference distribution, \mathcal{D} (see subsection 2.3) sample a preference profile P from \mathcal{D}
2. Randomly rename each alternative in P .⁵
3. Find, by exhaustive search, the committee $c = \arg \min_c \sum_{A \in \mathbb{A}} A(P, c)$ minimizing the number of axiom violations in \mathbb{A} with ties broken lexicographically.

⁵Renaming alternatives ensures that our data satisfies neutrality by treating alternatives identically. Otherwise, some preference distributions are biased towards specific outcomes. For example, our *identity* distribution only generates the ranking $a_1 \succ a_2 \succ \dots$ which will bias a network to always elect a_1 through a_k .

4.2 Input Data Structure

We transform each preference profile into three separate matrices, preserving aspects of voter preferences while reducing input size to depend only on m , the number of alternatives. This ensures models are applicable for any number of voters n , while being specific to m and committee size k .

Majority Matrix The preference profile is transformed into an $m \times m$ matrix $\mathcal{R}^{\text{majority}}$ with $\mathcal{R}_{ij}^{\text{majority}} = 1$ if a weak majority of voters prefer a_i to a_j and 0 otherwise.

Weighted Preference Matrix The preference profile is transformed into an $m \times m$ matrix $\mathcal{R}^{\text{weighted}}$ with $\mathcal{R}_{ij}^{\text{weighted}} = c$ to indicate that c voters prefer a_i over a_j .

Ranking Matrix The preference profile is transformed into an $m \times m$ matrix $\mathcal{R}^{\text{ranked}}$ where $\mathcal{R}_{ij}^{\text{ranked}} = c$ indicates that c voters rank a_i in position j .

The majority and weighted preference matrices correspond, respectively, to tournaments and weighted tournaments in the literature [8]. We flatten these matrices to 1-dimensional arrays. $\mathcal{R}^{\text{weighted}}$ and $\mathcal{R}^{\text{ranked}}$ are stored in a normalized form so that the model output is agnostic of the exact number of voters present. We concatenate each of these arrays to form the input to our models. To simplify notation, we continue to write $\mathcal{F}(P)$ when discussing the output of a trained model; however, in place of P , the function receives only the above transformations of P .

4.3 Model Loss

During training, committees of input examples are encoded as k -hot vectors: e.g. $[0, 1, 1, 0, 1]$ indicates that alternatives 2, 3, and 5 are in the winning committee while alternatives 1 and 4 are not in the winning committee. Our model uses the L1 loss function⁶, defined as:

$$L_1(\mathbf{y}, \hat{\mathbf{y}}) = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i|$$

where y is the committee from the training example and \hat{y} is the model’s predicted output. As the model does not output integer values, the output must be processed to form a committee: the highest k values in the length m list output by the model are set to 1 while the other values are set to 0.

5 Experimental Results and Analysis

We now describe the experiments we have run, provide their results and discuss their implications. From both existing and learned rules we are able to extract novel conclusions about each of our cornerstones and the learnability of novel rules.

5.1 Experimental Training Parameters

In our experiments we learn two different rule configurations. These configurations differ primarily in which axioms are evaluated in the loss function.

Learning Configuration 1: All Axioms Our initial learned rule $\mathcal{F}^{\text{NN-all}}$ is trained on data generated from evaluating violations for *all* axioms we have described. We train $\mathcal{F}^{\text{NN-all}}$ on all test distributions described below.

⁶We tested all applicable loss functions in Pytorch and found that L1 Loss provided the best results [36].




Figure 2: Axiom violation rate averaged over all axioms for each number of voters and winners on Mixed preferences.

Learning Configuration 2: Root Axioms We also train $\mathcal{F}^{\text{NN-root}}$ on a second set of axioms: those axioms which, if satisfied, imply the satisfaction of all other axioms we use. This includes: Majority Loser, Condorcet Winner, Dummett’s Condition, Local Stability, Strong Pareto, and the Core. See Appendix A for details of these axiom relationships. We train $\mathcal{F}^{\text{NN-root}}$ on all test distributions described below, except as noted.

5.1.1 Preference Distributions

Across both configurations most parameter values are held constant. We train (as applicable) and test all rules (learned and existing) on all combinations of $m = \{5, 6, 7\}$ alternatives and $k = \{1, 2, \dots, m - 1\}$ winners. In all cases we use profiles with $n = 50$ voters. All rules are evaluated on 25,000 profiles sampled from each of the following test distributions with learned rules being tested on the distribution upon which they were trained.

- IC, IAC, Identity take no parameter
- Mallow’s with θ sampled uniformly at random as described by Boehmer et al.[7].
- Urn with α sampled from a Gamma distribution [6].
- Single-peaked distributions as described by Conitzer [11] and Walsh [43].
- Stratification with $w = 0.5$.
- 8 Euclidean distributions with each combination of: 3 or 10 dimensions, a Ball or Cube topology, and Uniform or Gaussian placement of voters. Note: Due to the similarity of results across Euclidean distributions from $\mathcal{F}^{\text{NN-all}}$, we trained $\mathcal{F}^{\text{NN-root}}$ only on one Euclidean distribution, the 3-dimensional Gaussian Ball.
- 1 distribution containing an even mixture of the 16 other distributions. We test mixed distributions on $\mathcal{F}^{\text{NN-all}}$ and pre-existing rules, but not $\mathcal{F}^{\text{NN-root}}$.

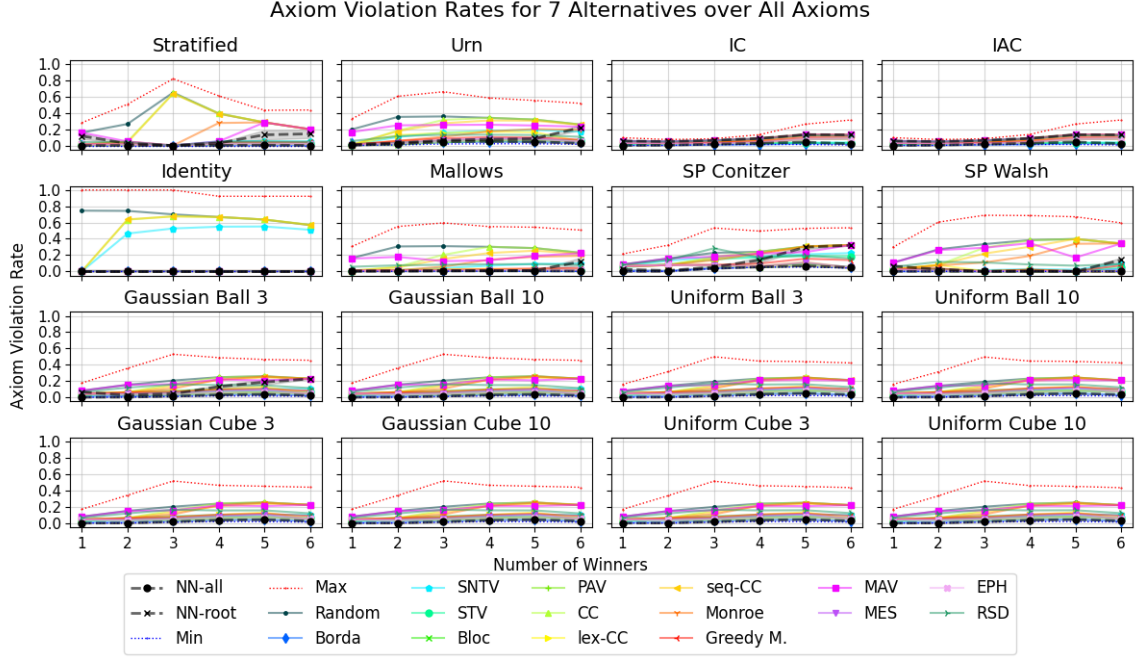


Figure 3: Axiom violation rates for each rule under each individual preference distribution for $m = 7$. In all cases, $\mathcal{F}^{\text{NN-all}}$ has AVR lower than, or similar to, other rules.

5.1.2 Learning Parameters

For each configuration, we train 20 neural networks using PyTorch [36], with 5 hidden layers of 256 nodes. Networks are trained for up to 50 epochs using the Adam optimizer with a learning rate of 0.0001, stopping early if no improvement of 0.0005 occurs over 10 epochs. We generate separate training and testing data sets of 25,000 examples each for all 255 unique combinations of m , k , and \mathcal{D} . Each profile contains 50 voters⁷. We do not filter profiles to ensure no overlap between training and test sets but find that there is minimal overlap between these sets in almost all cases.⁸

5.2 Understanding Our Cornerstones

We first examine what we can observe by considering each pair of our three cornerstones: axioms, voting rules, and preferences. By examining multiple aspects of our area of focus we are able to form novel connections between each topic.

⁷In preliminary experiments we have trained networks using data containing a number of voters sampled from a normal distribution truncated between 25 and 75 voters. We found no difference to performance and use profiles with 50 voters for simplicity.

⁸In the identity distribution all profiles are identical. On profiles sampled from the Urn distribution for 5, 6, and 7 alternatives roughly 3.5%, 2%, and 0.5% of profiles overlapped on a training distribution of 25,000 profiles. For Mallows preferences with 5 alternatives, up to 0.2% of profiles overlap. In all other cases training distributions have zero overlap.

| Method | Mean | Maj W | Maj L | Cond W | Cond L | Pareto | F Maj | Unanimity | Dummett's | JR | EJR | Core | S. Coalitions | Stability |
|-----------|-------------|----------|-------|-------------|----------|-------------|----------|-----------|-----------|----------|----------|----------|---------------|-------------|
| NN-all | .017 | .000 | .000 | .015 | .000 | .004 | .000 | .000 | .061 | .001 | .001 | .001 | .046 | .092 |
| NN-root | .038 | .001 | .030 | .200 | .010 | .045 | .024 | .000 | .056 | .000 | .000 | .002 | .044 | .082 |
| Min | .009 | 0 | .001 | .036 | 0 | .001 | .001 | 0 | .012 | 0 | .000 | .000 | .005 | .059 |
| Max | .440 | .125 | .340 | .919 | .635 | .620 | .175 | .076 | .555 | .234 | .354 | .381 | .521 | .787 |
| Borda | .021 | .001 | .004 | .125 | 0 | .004 | .011 | 0 | .044 | .000 | .000 | .000 | .031 | .056 |
| EPH | .040 | .000 | .001 | .270 | .002 | .000 | .001 | 0 | .082 | .000 | .000 | .000 | .063 | .096 |
| SNTV | .099 | 0 | .098 | .619 | .007 | .227 | .106 | .049 | .062 | .001 | .054 | .058 | 0 | .012 |
| Bloc | .039 | .000 | .001 | .254 | .002 | 0 | 0 | 0 | .080 | .000 | .000 | .000 | .061 | .106 |
| STV | .048 | 0 | .037 | .442 | .002 | .118 | .029 | 0 | 0 | .000 | .000 | .001 | 0 | .001 |
| PAV | .043 | .001 | .001 | .308 | .002 | 0 | .004 | 0 | .088 | 0 | 0 | 0 | .068 | .091 |
| MES | .049 | .001 | .002 | .351 | .002 | .001 | .008 | 0 | .096 | 0 | 0 | 0 | .075 | .095 |
| CC | .195 | .036 | .146 | .756 | .031 | .344 | .141 | .062 | .308 | 0 | .084 | .091 | .232 | .301 |
| seq-CC | .183 | .032 | .139 | .740 | .025 | .297 | .140 | .061 | .292 | 0 | .078 | .081 | .216 | .278 |
| lex-CC | .061 | .005 | .007 | .440 | .002 | 0 | .024 | 0 | .117 | 0 | .000 | .000 | .091 | .112 |
| Monroe | .130 | .007 | .078 | .649 | .026 | .234 | .060 | 0 | .214 | 0 | .002 | .006 | .180 | .231 |
| Greedy M. | .063 | .002 | .019 | .448 | .003 | .012 | .023 | 0 | .112 | 0 | 0 | 0 | .089 | .118 |
| MAV | .157 | .022 | .110 | .750 | .044 | .279 | .084 | 0 | .219 | .015 | .022 | .022 | .179 | .300 |
| RSD | .105 | .008 | .056 | .594 | .016 | 0 | .036 | 0 | .148 | .030 | .032 | .033 | .120 | .299 |
| Random | .237 | .063 | .171 | .845 | .057 | .406 | .160 | .071 | .326 | .049 | .125 | .134 | .252 | .419 |

Table 1: Axiom violation rates for 7 alternatives averaged over all distributions and numbers of winners. Voting rules and axioms are separated to indicate to which category of rule/axiom they belong. Root axioms are underlined. Bold values indicate the best result of a column, italic values have been rounded to zero. Shaded green indicates that previous work has shown the rule satisfies this axiom [18, 31]. Due to differences in tie-breaking with previous work, some edge cases do not match prior theoretical results.

5.2.1 Axioms and Voting Rules

In Table 1 we show the violation rate of each voting rule on each individual axiom, as well as an average violation rate over all axioms. This allows us to directly explore the relationship between voting rules and axiomatic performance. After dividing rules and axioms into their informal categories of excellence-based or diverse/proportional some consistent patterns emerge.

Excellence based rules generally have quite low axiom violation rates. Being good at excellence-based axioms leads to strong performance on diverse/proportional axioms. The axioms that are formally satisfied by some proportionality-based rules (i.e., JR, EJR) are almost never violated in practice by most excellence-based rules. Other proportionality-based axioms, such as Solid Coalitions, are violated more frequently by all diverse/proportional rules except \mathcal{F}^{STV} .

Observation: Alternatives which are liked individually are also likely to be members of committees which provide strong proportional properties. Identifying the alternatives which are liked individually may be easier to do well at than identifying strong sets of alternatives.

Observation: Certain axioms – Condorcet Winner, Dummett’s Condition, Local Stability – are significantly harder to avoid violating than others. This pattern is consistent across all rules and is also reflected in the violation rate of the best committees on these axioms.

Observation: Learning a rule to minimize axiom violation rate results in strong performance across *all* axioms. Curiously, the learned rules violate the Condorcet Winner condition at a rate lower than they would were they to perfectly minimize violations across all axioms.

Observation: Learning basic axioms can inform us about more complex axioms. $\mathcal{F}^{\text{NN-root}}$, the rule trained to minimize violations on only the root axioms, has an average violation rate lower than most rules but much higher than $\mathcal{F}^{\text{NN-all}}$. Notably, $\mathcal{F}^{\text{NN-root}}$ has a higher violation rate *on most root axioms* than $\mathcal{F}^{\text{NN-all}}$.

5.2.2 Axioms and Preference Distributions

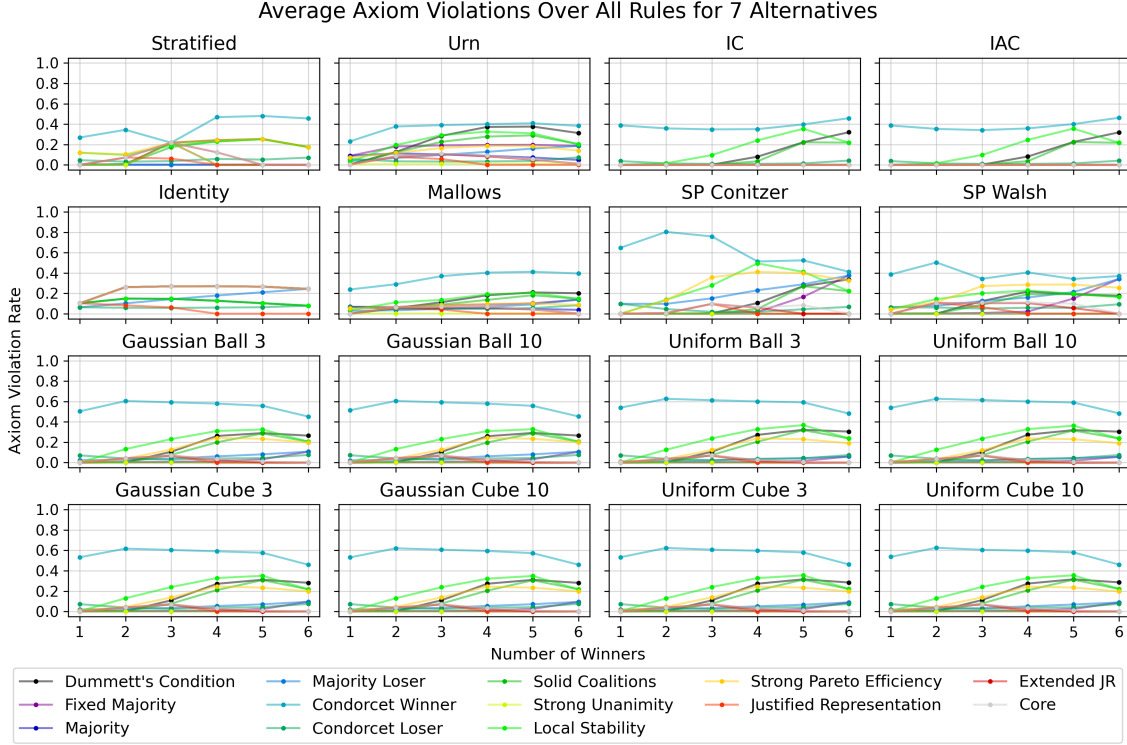


Figure 4: Violation rates of each axiom for every preference distribution, averaged over all rules. This shows trends common across all rules in the relationships between axioms and preference distributions.

We visualize the relationship between axioms and preferences in Figure 4. Some previous trends are visible – the Condorcet Winner axiom is violated frequently in almost all distributions, generally at a similar rate regardless of the number of winners.

Observation: Distributions with medium or high levels of structure exhibit quite different axiomatic patterns than other distributions with the same level of structure. This is most surprising in the large differences across Single-Peaked distributions. In fact, on many individual rules Local Stability is violated on Conitzer’s Single-Peaked distribution at quite a high rate while Walsh’s Single-Peaked preferences do not result in violations of Local Stability.

Observation: Figure 3 shows a spike in maximum violation rate occurring in most distributions when selecting a committee that contains less than half of alternatives. This is likely an artifact of specific axioms: For example, on Identity preferences it is impossible to violate Justified Representation when choosing a committee containing more than half of alternatives.

5.2.3 Voting Rules and Preference Distributions

We now focus primarily on voting rules and preferences. Figure 3 shows the axiom violation rate averaged over all axioms for each different voting rule. As we would expect, highly structured preference distributions (Identity, Single-Peaked, Stratified) have correspondingly higher axiom violation rates by most rules. This is natural as additional structure provides more opportunity to violate axioms, as shown by \mathcal{F}^{Max} which is much higher for these distributions. Similarly, in distributions with low structure (Impartial Culture, Impartial Anonymous Culture) the maximum violation rate is quite low. Again, we find this intuitive: if all alternatives have equal support then any committee becomes an equally reasonable choice and reasonable axioms should not be violated.

In looking at the behaviour of specific rules on each distribution we can make more fine-grained observations:

Observation: On Stratified preferences, $\mathcal{F}^{\text{seq-CC}}$, \mathcal{F}^{MAV} and \mathcal{F}^{CC} have higher violation rates than other rules. When $k > 2$, $\mathcal{F}^{\text{seq-CC}}$ has an AVR very similar to random committees. The “phase shift” change in violations as $k > 3$ alternatives are approved correlates with the election of an alternative in the bottom half of the stratified preferences. Recall that there is some set of $\lfloor \frac{k}{2} \rfloor$ alternatives who are all more preferred by every voter than the other alternatives. When $k > 3$ at least one of these unpopular alternatives is elected which, intuitively, causes violations.

Observation: Tie-breaking methods become important when ties are common. While perhaps an obvious statement, this observation explains the unusually high violation rates of \mathcal{F}^{CC} , $\mathcal{F}^{\text{seq-CC}}$, and $\mathcal{F}^{\text{SNTV}}$ on Identity preferences.

5.2.4 Voting Rule Differences

| | Borda | EPH | SNTV | Bloc | STV | PAV | MES | CC | seq-CC | lex-CC | Monroe | Greedy M. | MAV | RSD |
|-----------|-------|------|------|------|------|------|------|------|--------|--------|--------|-----------|------|------|
| Borda | .000 | - | - | - | - | - | - | - | - | - | - | - | - | - |
| EPH | .248 | .000 | - | - | - | - | - | - | - | - | - | - | - | - |
| SNTV | .463 | .420 | .000 | - | - | - | - | - | - | - | - | - | - | - |
| Bloc | .248 | .021 | .417 | .000 | - | - | - | - | - | - | - | - | - | - |
| STV | .315 | .360 | .299 | .359 | .000 | - | - | - | - | - | - | - | - | - |
| PAV | .255 | .051 | .429 | .065 | .365 | .000 | - | - | - | - | - | - | - | - |
| MES | .271 | .118 | .412 | .130 | .373 | .086 | .000 | - | - | - | - | - | - | - |
| CC | .600 | .483 | .588 | .487 | .572 | .464 | .496 | .000 | - | - | - | - | - | - |
| seq-CC | .568 | .484 | .464 | .492 | .566 | .467 | .431 | .672 | .000 | - | - | - | - | - |
| lex-CC | .310 | .133 | .451 | .144 | .396 | .089 | .117 | .440 | .461 | .000 | - | - | - | - |
| Monroe | .514 | .396 | .531 | .400 | .491 | .376 | .408 | .117 | .618 | .366 | .000 | - | - | - |
| Greedy M. | .334 | .223 | .428 | .232 | .404 | .203 | .170 | .512 | .402 | .218 | .431 | .000 | - | - |
| MAV | .611 | .599 | .694 | .598 | .617 | .597 | .612 | .342 | .813 | .587 | .344 | .635 | .000 | - |
| RSD | .486 | .465 | .586 | .464 | .526 | .467 | .470 | .646 | .629 | .481 | .577 | .484 | .626 | .000 |
| Random | .714 | .714 | .714 | .714 | .714 | .714 | .714 | .714 | .714 | .714 | .714 | .714 | .714 | .714 |
| Min | .158 | .261 | .476 | .254 | .325 | .278 | .300 | .566 | .598 | .331 | .481 | .361 | .561 | .490 |
| Max | .940 | .936 | .865 | .941 | .912 | .930 | .922 | .813 | .748 | .913 | .866 | .904 | .833 | .856 |
| NN-all | .147 | .240 | .484 | .232 | .335 | .258 | .283 | .569 | .593 | .315 | .484 | .349 | .562 | .484 |
| NN-root | .375 | .395 | .539 | .392 | .424 | .403 | .414 | .530 | .661 | .432 | .411 | .443 | .426 | .496 |

Table 2: Difference between rules for 7 alternatives with $1 \leq k < 7$ averaged over all preference distributions. Darker values correspond to larger differences. A difference of 0 between two rules indicates the rules always elect the same committee while a difference of 1 indicates that the rules’ winning committees have maximal overlap.

In Table 2 we show the mean difference between each rule based on the overlap between their elected committees. As expected, randomly chosen committees and \mathcal{F}^{Max} are almost always the most different than committees returned by all other rules. Surprisingly, there are exceptions: \mathcal{F}^{MAV} and $\mathcal{F}^{\text{seq-CC}}$ elect committees with *less* overlap than occurs with a randomly chosen committee. This indicates that these rules optimize for different, and mutually exclusive, goals. Other pairs of rules that are more different from each other than they are random committees can be found in Appendix C and our supplementary files, typically involving \mathcal{F}^{MAV} , $\mathcal{F}^{\text{seq-CC}}$, and \mathcal{F}^{CC} . Oddly, $\mathcal{F}^{\text{seq-CC}}$, and \mathcal{F}^{CC} typically elect highly distinct committees despite similar definitions. Note that all the differences from random committees listed in Table 2 are identical, this is merely an effect of rounding. The values are all close, however, due to the nature of random selection of winners.

On the other hand, we can observe rules with quite similar definitions: $\mathcal{F}^{\text{Bloc}}$, \mathcal{F}^{PAV} and \mathcal{F}^{EPH} all follow some procedure that award an equal number of points to a set of alternatives. Each of these rules have very low distance from the others, indicating that the minor differences in their definitions have relatively little effect on the committees they elect. Unsurprisingly, $\mathcal{F}^{\text{NN-all}}$ is most similar to \mathcal{F}^{Min} . This is understandable given that $\mathcal{F}^{\text{NN-all}}$ is trained to find committees that minimize the number of axiom violations. It is also interesting that $\mathcal{F}^{\text{Borda}}$ is relatively close to \mathcal{F}^{Min} , though, this is expected given that $\mathcal{F}^{\text{Borda}}$ was one of the best performing rules in our experiments. Overall, there does seem to be a trend of better performing rules being more similar to one another, and poorer performing rules being more similar to one another.

5.3 Learning From Learned Rules

As we see in Table 1, both learning configurations (optimizing for all axioms, and for only root axioms) result in rules with strong axiomatic performance. Averaged across all axioms, $\mathcal{F}^{\text{NN-all}}$ has a lower violation rate than any existing rule we evaluated while $\mathcal{F}^{\text{NN-root}}$ has a lower violation rate than all rules but $\mathcal{F}^{\text{Borda}}$. While these results are still above the lowest possible violation rate, as given by \mathcal{F}^{Min} , they show that networks are very capable of identifying committees that provide good axiomatic properties.

By learning two different sets of axioms we are able to both deepen our understanding of which axioms are more difficult to learn and to gauge whether learning “redundant” axioms with $\mathcal{F}^{\text{NN-all}}$ is beneficial. As $\mathcal{F}^{\text{NN-root}}$ has a much lower violation rate, it is clearly beneficial to learn from all axioms. This is intuitive, $\mathcal{F}^{\text{NN-all}}$ has the opportunity to learn to satisfy non-root axioms even in cases when root axioms may not be possible to mutually satisfy. However, even when evaluating *only the root axioms*, $\mathcal{F}^{\text{NN-all}}$ has an average violation rate of 0.0283 compared to an average violation rate of 0.0692 for $\mathcal{F}^{\text{NN-root}}$. Training on the additional axioms appears to provide enough additional signal to the learning process which reduces violations of root axioms.

5.4 Real-World Data

Using networks trained on the Mixed distribution, we applied each $\mathcal{F}^{\text{NN-all}}$ to a selection of real-world data collected from PrefLib [35]. We find that this highlights the sensitivity of learned models to their training data: our Mixed distribution is not a perfect match of the distribution underlying the empirical data we tested on. Nonetheless, we find that $\mathcal{F}^{\text{NN-all}}$ generalizes well, having, for example, a mean AVR of 0.102 compared to an AVR of 0.303 for committees selected at random. For more details and results of these experiments we refer to Appendix B.

5.5 Optimized Positional Scoring Rules

The low axiom violation rate of $\mathcal{F}^{\text{Borda}}$ inspires a natural follow-up question: Is there a positional scoring rule which provides even stronger performance than $\mathcal{F}^{\text{Borda}}$? A positive answer here would be interesting both to practitioners and theorists. The black box nature of the neural networks underlying each \mathcal{F}^{NN} restrict these rules from being applied in settings where interpretability and reliability are vital. On the other hand, positional scoring rules provide an immediate intuition as to the behaviour of the rule that allows them to be applied in practice (e.g., the Plurality rule can be said to reward voter’s favourite alternatives while the Anti-Plurality rule punishes those alternatives least preferred by voters [8]). Theoreticians may find novel positional scoring rules interesting as such rules could point in the direction of new results or points of analysis.

As positional scoring rules have a natural “state” (their score vector) we are easily able to apply generic optimization techniques to find new vectors which may have low axiom violation rates. To do this, we use the `optimal-voting` package [2] which uses simulated annealing to optimize a score vector. We set the loss function equal to the axiom violation rate and run annealing for Mixed preferences using $m = 7$ alternatives only, with one annealing run for each number of winners. We refer to a rule optimized in this manner as \mathcal{F}^{Opt} . Due to computational and time constraints we run for only 1000 steps, using 2000 profiles sampled from the larger set of 25,000 profiles. All rules have an initial state corresponding to the Borda rule. We then evaluate the violation rates of each rule on a different set of 25,000 profiles from the Mixed distribution to ensure that our results do not simply reflect overfitting on the training data.

Table 3 shows the results of annealing using all axioms in the loss function, and using just root axioms. The vectors resulting from annealing are included in Appendix C.1. We highlight that our optimization process is quite limited in both data and number of steps. Despite this, optimization consistently finds some vector that outperforms $\mathcal{F}^{\text{Borda}}$. These vectors also display a common trend; each optimized score vector awards higher points to the first k alternatives, then significantly reduces the number of points awarded. This suggests a trend towards a vector similar to k -approval (where each voter’s most preferred k alternatives receive a single point). However, we note that the k -approval rule itself proves to have far worse axiomatic properties than both \mathcal{F}^{Opt} and $\mathcal{F}^{\text{Borda}}$. These preliminary results leave open the possibility of some novel class of score vector which follows a *different* decreasing sequence than $\mathcal{F}^{\text{Borda}}$.

| k | All Axioms | | Root Axioms | |
|-----|------------------------------|----------------------------|------------------------------|----------------------------|
| | $\mathcal{F}^{\text{Borda}}$ | \mathcal{F}^{Opt} | $\mathcal{F}^{\text{Borda}}$ | \mathcal{F}^{Opt} |
| 1 | 0.010 | 0.008 | 0.019 | 0.016 |
| 2 | 0.011 | 0.011 | 0.021 | 0.018 |
| 3 | 0.022 | 0.018 | 0.036 | 0.032 |
| 4 | 0.030 | 0.025 | 0.044 | 0.043 |
| 5 | 0.034 | 0.028 | 0.042 | 0.044 |
| 6 | 0.025 | 0.022 | 0.030 | 0.027 |

Table 3: Axiom violation rate for $m = 7$ alternatives on both axiom sets by $\mathcal{F}^{\text{Borda}}$ and the positional scoring rule found via simulated annealing, \mathcal{F}^{Opt} .

6 Discussion

We have shown that (1) different multi-winner voting rules elect committees that are distinct from one another and violate axioms at different rates, (2) these differences depend greatly upon the underlying preferences of voters, and (3) it is possible to learn novel rules which result in significantly lower axiom violation rates. These contributions join a growing body of work demonstrating the effectiveness of combining machine learning with social choice [28, 12, 32] and adds to the literature extending axiomatic analysis beyond the worst case [25]. We highlight findings of particular note to

the broader research community along with and opportunities for future work.

Competing Definitions of Proportionality We observe two contrasting types of proportionality among our axioms. Core, EJR, and JR are formally linked (each a weaker version of the former) [31] while Local Stability, Dummett’s, and Solid Coalitions have distinct origins. Most voting rules violate the first three of these axioms similarly, while there is much more variability in AVR on axioms from the second type. Clear distinctions between these two types of proportionality are missing in the literature. For example, \mathcal{F}^{CC} and \mathcal{F}^{Monroe} “explicitly aim at proportional representation” [18] yet these rules violate Dummett’s Condition, Solid Coalitions, and Local Stability frequently while maintaining a low AVR on JR, EJR, and the Core.

Proportionality and Electing Losers The high AVR of the proportional rules \mathcal{F}^{CC} and \mathcal{F}^{Monroe} for the Condorcet Loser and Majority Loser axioms demonstrates a fundamental tension between (some types of) proportionality and avoiding bad alternatives. To satisfy proportionality, it may be necessary to elect an alternative ranked low by a majority of voters. New axioms may aim to bridge this gap by formally considering a balance between proportional representation and electing highly polarizing alternatives.

Individual Excellence and Proportional Rules Rules focused primarily on electing alternatives which are individually popular among voters generally result in fewer violations of axioms focused on proportionality than rules with a state goal of proportionality. If proportionality guarantees can be established for individual excellence rules this would allow benefitting from their superior performance in settings where proportionality is important.

Learning Simple Axioms Targetting both simple and complex axioms resulted in \mathcal{F}^{NN-all} having superior performance to $\mathcal{F}^{NN-root}$. This suggests that future axiom learning tasks could be enhanced by introducing simple dummy axioms that provide some feedback about more complex axioms (e.g., a majority winner must always be in a Condorcet winning set, so the Majority Winner axiom may be seen as informative about Condorcet Winner).

Difference Between Rules Our results can be qualitatively compared with existing maps of multi-winner elections [22]. In particular, in both their work and ours: (1) rules focused on individual excellence elect different committees than proportional rules, and (2) \mathcal{F}^{MAV} is very different from all other rules.

6.1 Future Directions

There are several interesting directions that this work can take. First, our research can be directly extended to support arbitrary sets of intraprofile axioms and extensions to the more general class of *interprofile* axioms [39, 40]. This would allow us to measure violations of additional axioms such as consistency or clone-proofness [8].

Second, we proposed a new metric for evaluating voting rules, the axiom violation rate. While impossibility theorems show that some axioms can never be mutually satisfied; in these cases, an interesting challenge for future research is to develop new rules optimized for low axiom violation rate. Using black-box machine learning techniques may be sufficient, however identifying classes of rule (such as positional scoring rules) which can be optimized in a way that provides a reliable, interpretable outcome will provide additional usability.

Finally, we are interested in extending our metric of axiom violations to allow robust, theoretical conclusions about axiomatic properties – similar to those provided by the PAC-learning framework [42]. We believe that such an approach would complement both the axiomatic and the empirical studies of voting rules.

References

- [1] Cem Anil and Xuchan Bao. Learning to elect. *Advances in Neural Information Processing Systems*, 34:8006–8017, 2021.
- [2] Ben Armstrong. Optimal voting, 2025. URL <https://pypi.org/project/optimal-voting/>.
- [3] Ben Armstrong and Kate Larson. Machine learning to strengthen democracy. In *NeurIPS Joint Workshop on AI for Social Good*, 2019.
- [4] Haris Aziz, Edith Elkind, Piotr Faliszewski, Martin Lackner, and Piotr Skowron. The Condorcet principle for multiwinner elections: From shortlisting to proportionality. *arXiv preprint arXiv:1701.08023*, 2017.
- [5] Duncan Black. On the rationale of group decision-making. *Journal of Political Economy*, 56: 23–34, 1948.
- [6] Niclas Boehmer, Robert Brederbeck, Piotr Faliszewski, Rolf Niedermeier, and Stanisław Szufa. Putting a compass on the map of elections. *arXiv preprint arXiv:2105.07815*, 2021.
- [7] Niclas Boehmer, Piotr Faliszewski, and Sonja Kraiczy. Properties of the Mallows model depending on the number of alternatives: A warning for an experimentalist. In *International Conference on Machine Learning*, pages 2689–2711. PMLR, 2023.
- [8] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- [9] Markus Brill and Jannik Peters. Robust and verifiable proportionality axioms for multiwinner voting. *arXiv preprint arXiv:2302.01989*, 2023.
- [10] Dávid Burka, Clemens Puppe, László Szepesváry, and Attila Tasnádi. Voting: A machine learning approach. *European Journal of Operational Research*, 299(3):1003–1017, 2022.
- [11] Vincent Conitzer. Eliciting single-peaked preferences using comparison queries. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 1–8, 2007.
- [12] Vincent Conitzer, Rachel Freedman, Jobst Heitzig, Wesley H Holliday, Bob M Jacobs, Nathan Lambert, Milan Mossé, Eric Pacuit, Stuart Russell, Hailey Schoelkopf, et al. Social choice for AI alignment: Dealing with diverse human feedback. *arXiv preprint arXiv:2404.10271*, 2024.
- [13] Bernard Debord. Prudent k -choice functions: Properties and algorithms. *Mathematical Social Sciences*, 26(1):63–77, 1993.
- [14] Greg d’Eon and Kate Larson. Testing axioms against human reward divisions in cooperative games. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems*, pages 312–320, 2020.
- [15] M Dummett. *Voting Procedures*. Oxford University Press, 1984.
- [16] Florian Eggenberger and George Pólya. Über die statistik verketteter vorgänge. *ZAMM-Journal of Applied Mathematics and Mechanics/Zeitschrift für Angewandte Mathematik und Mechanik*, 3(4):279–289, 1923.

- [17] Edith Elkind. Restricted preference domains in social choice: Two perspectives. In *Algorithmic Game Theory: 11th International Symposium, SAGT 2018, Beijing, China, September 11-14, 2018, Proceedings 11*, pages 12–18. Springer, 2018.
- [18] Edith Elkind, Piotr Faliszewski, Piotr Skowron, and Arkadii Slinko. Properties of multiwinner voting rules. *Social Choice and Welfare*, 48:599–632, 2017.
- [19] James M Enelow and Melvin J Hinich. *The Spatial Theory of Voting: An Introduction*. CUP Archive, 1984.
- [20] Roy Fairstein, Dan Vilenchik, and Kobi Gal. Learning aggregation rules in participatory budgeting: A data-driven approach. *arXiv preprint arXiv:2412.01864*, 2024.
- [21] Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. Multiwinner voting: A new challenge for social choice theory. *Trends in Computational Social Choice*, 74(2017):27–47, 2017.
- [22] Piotr Faliszewski, Martin Lackner, Krzysztof Sornat, and Stanisław Szufa. An experimental comparison of multiwinner voting rules on approval elections. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pages 2675–2683, 2023.
- [23] Daniel Firebanks-Quevedo. Machine learning? In MY election? It’s more likely than you think: Voting rules via neural networks. *Thesis*, 2020.
- [24] Peter C Fishburn. Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, 33(3):469–489, 1977.
- [25] Bailey Flanigan, Daniel Halpern, and Alexandros Psomas. Smoothed analysis of social choice revisited. In *International Conference on Web and Internet Economics*, pages 290–309. Springer, 2023.
- [26] William V Gehrlein. The Condorcet criterion and committee selection. *Mathematical Social Sciences*, 10(3):199–209, 1985.
- [27] William V Gehrlein and Peter C Fishburn. Condorcet’s paradox and anonymous preference profiles. *Public Choice*, pages 1–18, 1976.
- [28] Noah Golowich, Harikrishna Narasimhan, and David C Parkes. Deep learning for multi-facility location mechanism design. In *IJCAI*, pages 261–267, 2018.
- [29] Georges Théodule Guilbaud. Les théories de l’intérêt général et le problème logique de l’agrégation. *Economie appliquée*, 5(4):501–584, 1952.
- [30] Hanna Kujawska, Marija Slavkovik, and Jan-Joachim Rückmann. Predicting the winners of Borda, Kemeny and Dodgson elections with supervised machine learning. In *Multi-Agent Systems and Agreement Technologies: 17th European Conference, EUMAS 2020, and 7th International Conference, AT 2020, Thessaloniki, Greece, September 14-15, 2020, Revised Selected Papers 17*, pages 440–458. Springer, 2020.
- [31] Martin Lackner and Piotr Skowron. *Multi-Winner Voting with Approval Preferences*. Springer International Publishing, 2023. ISBN 9783031090165. doi: 10.1007/978-3-031-09016-5. URL <http://dx.doi.org/10.1007/978-3-031-09016-5>.

- [32] Marc Lanctot, Kate Larson, Michael Kaisers, Quentin Berthet, Ian Gemp, Manfred Diaz, Roberto-Rafael Maura-Rivero, Yoram Bachrach, Anna Koop, and Doina Precup. Soft Condorcet optimization. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025)*, 2025.
- [33] Colin L Mallows. Non-null ranking models. i. *Biometrika*, 44(1/2):114–130, 1957.
- [34] Leonardo Matone, Ben Abramowitz, Nicholas Mattei, and Avinash Balakrishnan. Deepvoting: Learning voting rules with tailored embeddings. *arXiv preprint arXiv:2408.13630*, 2024.
- [35] Nicholas Mattei and Toby Walsh. Preflib: A library for preferences. In *Algorithmic Decision Theory: Third International Conference, ADT 2013, Bruxelles, Belgium, November 12-14, 2013, Proceedings 3*, pages 259–270. Springer, 2013.
- [36] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems*, 32, 2019.
- [37] Dominik Peters and Piotr Skowron. Proportionality and the limits of welfarism. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 793–794, 2020.
- [38] Jameson Quinn and Bruce Schneier. A proportional voting system for awards nominations resistant to voting blocs. *Preprint per https://www.schneier.com/academic/archives/2016/05/a_proportional_votin.html*. Accessed, 14, 2019.
- [39] Marie Christin Schmidtlein. Voting by axioms. *Master’s thesis. ILLC, University of Amsterdam*, 2022.
- [40] Marie Christin Schmidtlein and Ulle Endriss. Voting by axioms. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, pages 2067–2075, 2023.
- [41] Nicolaus Tideman. The single transferable vote. *Journal of Economic Perspectives*, 9(1):27–38, 1995.
- [42] Leslie G Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.
- [43] Toby Walsh. Generating single peaked votes. *arXiv preprint arXiv:1503.02766*, 2015.

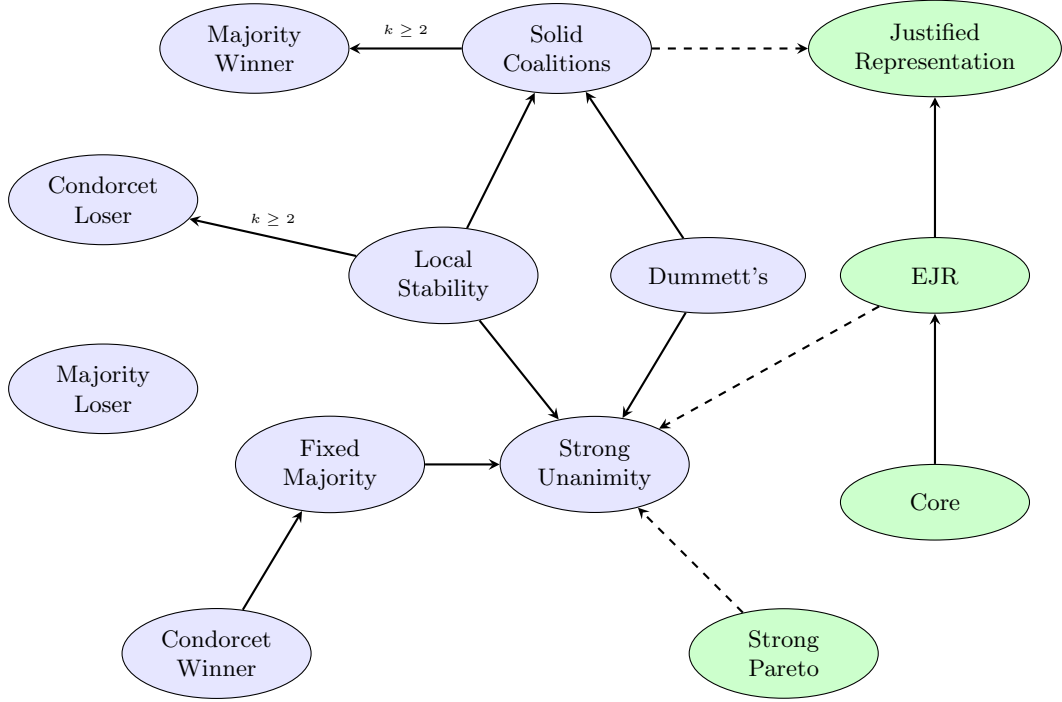


Figure 5: Satisfaction relationships between axioms that we consider. An arrow from node X to node Y indicates that satisfying axiom X also leads to satisfaction of axiom Y, or the equivalent contrapositive: violating Y also implies a violation of X. Blue nodes are those originating in ordinal settings while green nodes originate in approval-based settings. A dashed line indicates that the connection relies upon our assumption that each voter approves of their k top-ranked alternatives.

Root Nodes: Local Stability, Dummett's, Condorcet Winner, Strong Pareto, Core, Majority Loser.

A Details on Axiom Relationships

This section describes the satisfaction relationships between the axioms we consider. We show in Figure 5 that many of the axioms we use are necessarily satisfied if some other axiom is satisfied. Some of these relationships have been previously described while several represent novel connections. We do not claim that these results are complete; other connections may exist which we have not identified. This section serves to provide some intuition as to which axioms are more or less similar to others.

A.1 Known Relationships

We first list pairs of axioms for which we know that one implies the other. As this is not the primary focus of our paper, our proofs remain at a high level and are relatively informal.

Relationship 1. *Condorcet Winner* \implies *Fixed Majority*

A Fixed Majority is a set of voters that is a Condorcet winning set. Thus, satisfying the Condorcet Winner property will also lead to satisfaction of Fixed Majority.

Relationship 2. *If the selected committee has at least $k \geq 2$ alternatives, then Solid Coalitions \implies Majority Winner*

Let X refer to the set of all alternatives ranked first by at least $\frac{n}{k}$ voters. If Solid Coalitions is satisfied then every member of X is in the winning committee. A majority winner is one ranked first by $\frac{n}{2} \geq \frac{n}{k}$ voters. Thus, if Solid Coalitions is satisfied and a majority winner exists, then it is elected. Note that this is not the case when $k = 1$. Then, X contains only an alternative that is unanimously ranked first which may not include a majority winner.

Relationship 3. *If every voter approves of exactly their k top-ranked alternatives, then Solid Coalitions \implies Justified Representation.*

In order to satisfy JR, whenever at least $\frac{n}{k}$ voters approve of an alternative, one of those voters must approve of one of the elected alternatives. SC requires that whenever $\frac{n}{k}$ voters rank some alternative c first, this alternative is in the winning set.

If SC is satisfied then, by assumption, whenever there exists a set of at least $\frac{n}{k}$ voters who all rank an alternative first, they also approve that alternative and elect it.

Relationship 4. *If a locally stable committee exists, then Local Stability \implies Solid Coalitions.*

Consider a locally stable committee C . Say there exists a set of voters V^* with $|V^*| \geq \frac{n}{k}$ who all rank candidate x first. If $x \in C$ then Solid Coalitions is satisfied. If $x \notin C$ then the committee is not locally stable and we have a contradiction. Thus, a locally stable committee will always elect an alternative required by Solid Coalitions.

Relationship 5. *If a locally stable committee exists, then Local Stability \implies Strong Unanimity.*

Consider a locally stable committee C and a profile in which all voters rank the same k alternatives above all others. If C were not exactly the k top-ranked alternatives, then there would be some alternative not in C preferred by all voters, and C would not be locally stable.

Relationship 6. *If every voter approves of exactly their k top-ranked alternatives, then EJR \implies Strong Unanimity.*

If there exists a profile in which all voters rank the same k alternatives above all others (and, by assumption, are exactly the approval set of each voter), these voters form a k -cohesive group and, if EJR is satisfied, will form the winning committee.

Relationship 7. *If every voter approves of exactly their k top-ranked alternatives, then Strong Pareto \implies Strong Unanimity.*

Say there exists a profile in which all voters rank the same k alternatives above all others and, by assumption, form the approval set of each voter. Any committee which is not exactly these k alternatives will be dominated by a committee containing these k alternatives so when voters are unanimous Strong Pareto is sufficient to elect their preferred committee.

Relationship 8. *If a locally stable committee exists with at least $k \geq 2$ members, then Local Stability \implies Condorcet Loser.*

For any preference profile with a Condorcet losing committee C , there exists some alternative $x \notin C$ which is preferred by at least $\frac{n}{2}$ voters to every alternative in C . Under our assumption that $k \geq 2$, a locally stable committee cannot contain any alternative y such that x is preferred to y by more than $\frac{n}{2}$ voters. Thus, a locally stable committee is not a Condorcet Loser.

Relationship 9. *Fixed Majority \implies Strong Unanimity.*

If every voter ranks the same k alternatives on top, those alternatives constitute a fixed majority committee and will be elected if FM is satisfied.

Relationship 10. *Dummett's \implies Solid Coalitions.*

We refer to Elkind et al. [18] for discussion of this relationship.

Relationship 11. *Dummett's \implies Unanimity.*

We refer to Elkind et al. [18] for discussion of this relationship.

Relationship 12. *Core \implies EJR \implies JR.*

We refer to Lackner and Skowron [31] for discussion of these relationships.

A.2 Relationships Known To Not Exist

We also include some counter-examples showing certain relationships that *do not* exist. While it is the case that Figure 5 may be missing edges we show here some of the edges that we thought may exist due to the similar nature of certain axioms, but can be demonstrated not to exist.

Relationship 13. *Condorcet Winner $\not\Rightarrow$ Dummett's and Dummett's $\not\Rightarrow$ Condorcet Winner.*

Say that $n = 10$, $l = 2$, $k = 5$ and consider voters with the following preference orders:

| | |
|---------------------------|---|
| v_0 | $0 \succ 1 \succ 2 \succ 3 \succ 4 \succ 5 \succ 6 \succ 7 \succ 8 \succ 9$ |
| v_1 | $0 \succ 1 \succ 4 \succ 5 \succ 6 \succ 7 \succ 8 \succ 9 \succ 2 \succ 3$ |
| v_2 | $0 \succ 1 \succ 6 \succ 7 \succ 8 \succ 9 \succ 2 \succ 3 \succ 4 \succ 5$ |
| v_3 | $0 \succ 1 \succ 8 \succ 9 \succ 2 \succ 3 \succ 4 \succ 5 \succ 6 \succ 7$ |
| $6 \times v_{\text{maj}}$ | $9 \succ 8 \succ 7 \succ 6 \succ 5 \succ 4 \succ 3 \succ 2 \succ 1 \succ 0$ |

A group of $\frac{l \cdot n}{k} = 4$ voters (v_0, v_1, v_2, v_3) rank the same 2 alternatives, $\{0, 1\}$, on top. These alternatives must be in the winning committee in order to satisfy Dummett's Condition. However, the Condorcet winning set is $\{9, 8, 7, 6, 5\}$. In this case, any possible winning committee will violate at least one of Condorcet Winner and Dummett's condition.

Relationship 14. *EJR $\not\Rightarrow$ Dummett's.*

Consider the following preference profile and say we are choosing $k = 5$ winners. For convenience, a vertical bar separates the approved alternatives and the disapproved alternatives.

| | |
|---------------------------|--|
| v_0 | $0 \succ 6 \succ 2 \succ 3 \succ 4 \mid \succ 5 \succ 1 \succ 7 \succ 8 \succ 9$ |
| v_1 | $0 \succ 7 \succ 2 \succ 3 \succ 4 \mid \succ 5 \succ 6 \succ 1 \succ 8 \succ 9$ |
| v_2 | $0 \succ 8 \succ 2 \succ 3 \succ 4 \mid \succ 5 \succ 6 \succ 7 \succ 1 \succ 9$ |
| v_3 | $0 \succ 9 \succ 2 \succ 3 \succ 4 \mid \succ 5 \succ 6 \succ 7 \succ 8 \succ 1$ |
| $6 \times v_{\text{maj}}$ | $9 \succ 8 \succ 7 \succ 6 \succ 5 \mid \succ 4 \succ 3 \succ 2 \succ 1 \succ 0$ |

In the profile above, note that v_0, v_1, v_2, v_3 form a 2-cohesive group. In order to satisfy, EJR one of these voters must approve of two elected alternatives. Similarly, the other 6 voters form a 3-cohesive group and must approve of three elected alternatives. A committee which would not violate EJR is $3, 4, 5, 6, 7$

Dummett's condition requires that, since 4 voters rank 0 on top, 0 must be in any winning committee. Thus, it is possible to violate Dummett's condition without violating EJR. Satisfaction of EJR does not imply satisfaction of Dummett's condition.

Relationship 15. *Dummett's* $\not\Rightarrow$ *EJR*.

| | |
|---------------------------|--|
| v_0 | $0 \succ 6 \succ 2 \succ 3 \succ 4 \mid \succ 5 \succ 1 \succ 7 \succ 8 \succ 9$ |
| v_1 | $1 \succ 7 \succ 2 \succ 3 \succ 4 \mid \succ 5 \succ 6 \succ 0 \succ 8 \succ 9$ |
| v_2 | $2 \succ 8 \succ 1 \succ 3 \succ 4 \mid \succ 5 \succ 6 \succ 7 \succ 0 \succ 9$ |
| v_3 | $3 \succ 9 \succ 2 \succ 1 \succ 4 \mid \succ 5 \succ 6 \succ 7 \succ 8 \succ 0$ |
| $6 \times v_{\text{maj}}$ | $9 \succ 8 \succ 7 \succ 6 \succ 5 \mid \succ 4 \succ 3 \succ 2 \succ 1 \succ 0$ |

Consider the above profile. Dummett's condition requires that the top 3 preferences of the group of 6 voters must be elected, but does not capture any other group information. A committee elected by a rule satisfying Dummett's condition may be $\{9, 8, 7, 6, 5\}$.

However, v_0, v_1, v_2, v_3 still form a 2-cohesive group – each of the voters approves of $\{3, 4\}$. In order to satisfy EJR, one of these voters must approve of two elected alternatives which is not the case for the winning committee $\{9, 8, 7, 6, 5\}$. Thus, satisfying Dummett's does not imply satisfaction of EJR.

Relationship 16. *Dummett's* $\not\Rightarrow$ *Fixed Majority* and *Fixed Majority* $\not\Rightarrow$ *Dummett's*.

Say that $n = 10$, $l = 2$, $k = 5$ and consider voters with the following preference orders:

| | |
|---------------------------|---|
| v_0 | $0 \succ 1 \succ 2 \succ 3 \succ 4 \succ 5 \succ 6 \succ 7 \succ 8 \succ 9$ |
| v_1 | $0 \succ 1 \succ 4 \succ 5 \succ 6 \succ 7 \succ 8 \succ 9 \succ 2 \succ 3$ |
| v_2 | $0 \succ 1 \succ 6 \succ 7 \succ 8 \succ 9 \succ 2 \succ 3 \succ 4 \succ 5$ |
| v_3 | $0 \succ 1 \succ 8 \succ 9 \succ 2 \succ 3 \succ 4 \succ 5 \succ 6 \succ 7$ |
| $6 \times v_{\text{maj}}$ | $9 \succ 8 \succ 7 \succ 6 \succ 5 \succ 4 \succ 3 \succ 2 \succ 1 \succ 0$ |

As in previous example (D and CW), consider the required winners from each axiom.

Dummett's requires that – since $\frac{ln}{k} = \frac{2 \cdot 10}{5} = 4$ voters rank 0 and 1 first – 0 and 1 must be winners. There is a Fixed Majority which all rank candidates 5 through 9 first. To satisfy FM, these must be the winners. In this case it is possible to satisfy one of Dummett's condition or Fixed Majority, but not both.

B Results on Real World Data

We gathered from PrefLib (an online repository of real-world preference information) all complete preference profiles containing 5, 6, or 7 alternatives [35]. Resulting in, respectively, 15, 19, and 22 preference profiles. For each of these preferences, we identified through brute force search the committees that would violates the fewest and most of our axioms.

As this amount of data is not suitable for a meaningful amount of learning, we measured the Axiom Violation Rate on this data of networks trained on our Mixed preference distributions. Results of this evaluation are found in the next section. These results emphasize clearly that the Mixed preference distribution is mildly informative about the distribution(s) underlying PrefLib data but that there is also a significant gap between the training and testing distributions in this area. That is, $\mathcal{F}^{\text{NN-all}}$ has worse AVR than some rules but also much better AVR than many rules. Complete results of the PrefLib evaluation for each number of alternatives are listed below.

- 5 Alternatives, Real World Data
- 6 Alternatives, Real World Data
- 7 Alternatives, Real World Data

B.1 5 Alternatives, Real World Data

| Method | Mean | Maj W | Maj L | Cond W | Cond L | Pareto | F Maj | Unanimity | Dummett's | JR | EJR | Core | S. Coalitions | Stability |
|-----------|-------------|----------|----------|----------|----------|----------|----------|-----------|-----------|----------|----------|----------|---------------|-----------|
| NN-all | .081 | .071 | .018 | .428 | .018 | .112 | .089 | 0 | .154 | 0 | 0 | 0 | .071 | .089 |
| Min | .004 | 0 | 0 | 0 | 0 | .018 | 0 | 0 | .036 | 0 | 0 | 0 | 0 | 0 |
| Max | .471 | .286 | .357 | .929 | .554 | .679 | .304 | 0 | .661 | .214 | .357 | .357 | .625 | .804 |
| Borda | .015 | 0 | 0 | .054 | 0 | .018 | 0 | 0 | .054 | 0 | 0 | 0 | .036 | .036 |
| EPH | .038 | 0 | 0 | .107 | 0 | 0 | 0 | 0 | .179 | 0 | 0 | 0 | .107 | .107 |
| SNTV | .055 | 0 | .089 | .339 | 0 | .125 | .125 | 0 | .036 | 0 | 0 | 0 | 0 | 0 |
| Bloc | .038 | 0 | 0 | .107 | 0 | 0 | 0 | 0 | .179 | 0 | 0 | 0 | .107 | .107 |
| STV | .019 | 0 | .018 | .161 | 0 | .036 | .036 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| PAV | .044 | 0 | 0 | .143 | 0 | 0 | 0 | 0 | .179 | 0 | 0 | 0 | .125 | .125 |
| MES | .045 | 0 | 0 | .161 | 0 | 0 | 0 | 0 | .179 | 0 | 0 | 0 | .125 | .125 |
| CC | .136 | .071 | .125 | .446 | .018 | .161 | .196 | 0 | .304 | 0 | 0 | 0 | .214 | .232 |
| seq-CC | .146 | .018 | .125 | .554 | .018 | .232 | .196 | 0 | .339 | 0 | 0 | 0 | .196 | .214 |
| lex-CC | .063 | 0 | .036 | .196 | 0 | 0 | .054 | 0 | .214 | 0 | 0 | 0 | .161 | .161 |
| Monroe | .077 | .036 | .054 | .286 | .018 | .071 | .071 | 0 | .196 | 0 | 0 | 0 | .125 | .143 |
| Greedy M. | .059 | 0 | .018 | .232 | 0 | .018 | .036 | 0 | .196 | 0 | 0 | 0 | .125 | .143 |
| MAV | .162 | .125 | .143 | .607 | .054 | .179 | .250 | 0 | .268 | 0 | 0 | 0 | .196 | .286 |
| RSD | .104 | .089 | .036 | .321 | .018 | 0 | .107 | 0 | .214 | .054 | .054 | .054 | .161 | .250 |
| Random | .255 | .196 | .214 | .732 | .107 | .357 | .286 | 0 | .464 | .071 | .089 | .089 | .304 | .411 |

Table 4: Average Axiom Violation Rate for 5 alternatives and $1 \leq k < 5$ winners across Real World Data preferences.

| | Borda | EPH | SNTV | Bloc | STV | PAV | MES | CC | seq-CC | lex-CC | Monroe | Greedy M. | MAV | RSD |
|-----------|-------|------|------|------|------|------|------|------|--------|--------|--------|-----------|------|------|
| Borda | 0 | - | - | - | - | - | - | - | - | - | - | - | - | - |
| EPH | .223 | 0 | - | - | - | - | - | - | - | - | - | - | - | - |
| SNTV | .304 | .375 | 0 | - | - | - | - | - | - | - | - | - | - | - |
| Bloc | .223 | 0 | .375 | 0 | - | - | - | - | - | - | - | - | - | - |
| STV | .187 | .304 | .241 | .304 | 0 | - | - | - | - | - | - | - | - | - |
| PAV | .214 | .036 | .384 | .036 | .321 | 0 | - | - | - | - | - | - | - | - |
| MES | .196 | .116 | .330 | .116 | .304 | .089 | 0 | - | - | - | - | - | - | - |
| CC | .411 | .304 | .482 | .304 | .446 | .304 | .393 | 0 | - | - | - | - | - | - |
| seq-CC | .455 | .491 | .348 | .491 | .491 | .473 | .402 | .616 | 0 | - | - | - | - | - |
| lex-CC | .241 | .071 | .393 | .071 | .330 | .036 | .125 | .277 | .491 | 0 | - | - | - | - |
| Monroe | .330 | .205 | .393 | .205 | .339 | .205 | .295 | .107 | .545 | .205 | 0 | - | - | - |
| Greedy M. | .268 | .214 | .304 | .214 | .330 | .205 | .152 | .429 | .304 | .223 | .321 | 0 | - | - |
| MAV | .643 | .562 | .714 | .562 | .616 | .545 | .616 | .330 | .848 | .518 | .411 | .670 | 0 | - |
| RSD | .366 | .348 | .437 | .348 | .393 | .339 | .304 | .500 | .545 | .348 | .420 | .339 | .616 | 0 |
| Random | .750 | .750 | .705 | .750 | .741 | .768 | .732 | .750 | .705 | .777 | .768 | .750 | .723 | .777 |
| Min | .241 | .223 | .411 | .223 | .223 | .232 | .286 | .348 | .571 | .259 | .268 | .339 | .473 | .402 |
| Max | .946 | .938 | .955 | .938 | .973 | .955 | .955 | .812 | .866 | .938 | .866 | .929 | .714 | .893 |
| NN-all | .409 | .481 | .484 | .481 | .449 | .481 | .445 | .659 | .527 | .499 | .561 | .472 | .748 | .517 |

Table 5: Difference between rules for 5 alternatives with $1 \leq k < 5$ on Real-World preferences.




Figure 6: Axiom violation rate for each axiom on Real World Data preferences with 5 alternatives.

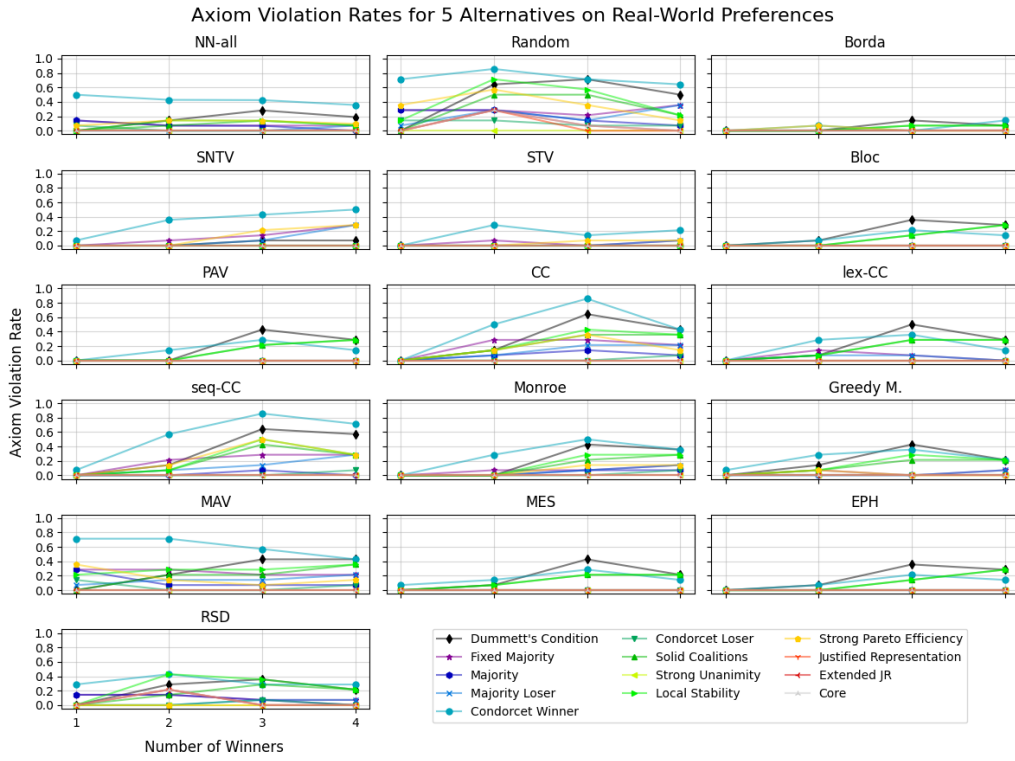


Figure 7: Axiom violation rate for each rule on Real World Data preferences with 5 alternatives.

B.2 6 Alternatives, Real World Data

| Method | Mean | Maj W | Maj L | Cond W | Cond L | Pareto | F Maj | Unanimity | Dummett's | JR | EJR | Core | S. Coalitions | Stability |
|-----------|-------------|----------|----------|-------------|----------|----------|----------|-----------|-----------|----------|----------|----------|---------------|-----------|
| NN-all | .117 | .052 | .033 | .476 | .056 | .101 | .067 | 0 | .242 | .019 | .030 | .030 | .180 | .241 |
| Min | .007 | 0 | 0 | .022 | 0 | .022 | 0 | 0 | .011 | 0 | 0 | 0 | 0 | .033 |
| Max | .385 | .222 | .222 | .744 | .511 | .589 | .144 | .011 | .611 | .167 | .211 | .222 | .556 | .789 |
| Borda | .023 | .011 | 0 | .044 | 0 | .033 | .011 | 0 | .089 | 0 | 0 | 0 | .056 | .056 |
| EPH | .050 | 0 | 0 | .222 | 0 | 0 | 0 | 0 | .178 | 0 | 0 | 0 | .122 | .133 |
| SNTV | .030 | 0 | .011 | .311 | 0 | .044 | .022 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| Bloc | .050 | 0 | 0 | .200 | 0 | 0 | 0 | 0 | .178 | 0 | 0 | 0 | .122 | .144 |
| STV | .019 | 0 | 0 | .189 | 0 | .056 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| PAV | .056 | 0 | 0 | .256 | 0 | 0 | 0 | 0 | .189 | 0 | 0 | 0 | .144 | .144 |
| MES | .056 | 0 | 0 | .256 | 0 | 0 | 0 | 0 | .178 | 0 | 0 | 0 | .144 | .144 |
| CC | .159 | .056 | .122 | .511 | .056 | .167 | .100 | .011 | .400 | 0 | .022 | .033 | .256 | .333 |
| seq-CC | .103 | .033 | .022 | .422 | .033 | .067 | .056 | 0 | .267 | 0 | .011 | .011 | .200 | .211 |
| lex-CC | .072 | 0 | 0 | .356 | 0 | 0 | .022 | 0 | .233 | 0 | 0 | 0 | .156 | .167 |
| Monroe | .117 | .033 | .044 | .467 | .033 | .100 | .067 | 0 | .333 | 0 | 0 | 0 | .189 | .256 |
| Greedy M. | .064 | .011 | 0 | .333 | .011 | 0 | .022 | 0 | .178 | 0 | 0 | 0 | .133 | .144 |
| MAV | .180 | .078 | .156 | .589 | .089 | .244 | .111 | 0 | .378 | .022 | .033 | .033 | .222 | .389 |
| RSD | .110 | .078 | .011 | .522 | .022 | 0 | .078 | 0 | .244 | .011 | .022 | .022 | .156 | .267 |
| Random | .192 | .122 | .156 | .622 | .100 | .256 | .133 | .011 | .367 | .022 | .067 | .067 | .256 | .322 |

Table 6: Average Axiom Violation Rate for 6 alternatives and $1 \leq k < 6$ winners across Real World Data preferences.

| | Borda | EPH | SNTV | Bloc | STV | PAV | MES | CC | seq-CC | lex-CC | Monroe | Greedy M. | MAV | RSD |
|-----------|-------|------|------|------|------|------|------|------|--------|--------|--------|-----------|------|------|
| Borda | 0 | - | - | - | - | - | - | - | - | - | - | - | - | - |
| EPH | .306 | 0 | - | - | - | - | - | - | - | - | - | - | - | - |
| SNTV | .415 | .428 | 0 | - | - | - | - | - | - | - | - | - | - | - |
| Bloc | .300 | .013 | .431 | 0 | - | - | - | - | - | - | - | - | - | - |
| STV | .293 | .363 | .185 | .369 | 0 | - | - | - | - | - | - | - | - | - |
| PAV | .326 | .035 | .428 | .044 | .365 | 0 | - | - | - | - | - | - | - | - |
| MES | .313 | .165 | .335 | .174 | .350 | .135 | 0 | - | - | - | - | - | - | - |
| CC | .565 | .467 | .646 | .476 | .578 | .444 | .533 | 0 | - | - | - | - | - | - |
| seq-CC | .476 | .443 | .322 | .448 | .415 | .422 | .317 | .628 | 0 | - | - | - | - | - |
| lex-CC | .389 | .107 | .474 | .117 | .406 | .072 | .181 | .424 | .426 | 0 | - | - | - | - |
| Monroe | .502 | .400 | .583 | .406 | .520 | .381 | .470 | .080 | .596 | .378 | 0 | - | - | - |
| Greedy M. | .387 | .344 | .319 | .354 | .352 | .322 | .217 | .556 | .320 | .331 | .507 | 0 | - | - |
| MAV | .687 | .689 | .757 | .693 | .717 | .672 | .683 | .337 | .791 | .654 | .350 | .719 | 0 | - |
| RSD | .602 | .580 | .607 | .580 | .589 | .570 | .602 | .661 | .676 | .563 | .611 | .591 | .689 | 0 |
| Random | .696 | .704 | .739 | .694 | .722 | .704 | .715 | .713 | .676 | .704 | .696 | .707 | .722 | .687 |
| Min | .220 | .348 | .380 | .346 | .257 | .361 | .385 | .543 | .517 | .404 | .476 | .435 | .596 | .594 |
| Max | .891 | .900 | .920 | .900 | .894 | .889 | .896 | .706 | .876 | .881 | .746 | .898 | .567 | .819 |
| NN-all | .489 | .583 | .484 | .579 | .523 | .580 | .534 | .725 | .428 | .593 | .702 | .535 | .804 | .618 |

Table 7: Difference between rules for 6 alternatives with $1 \leq k < 6$ on Real-World preferences.




Figure 8: Axiom violation rate for each axiom on Real World Data preferences with 6 alternatives.

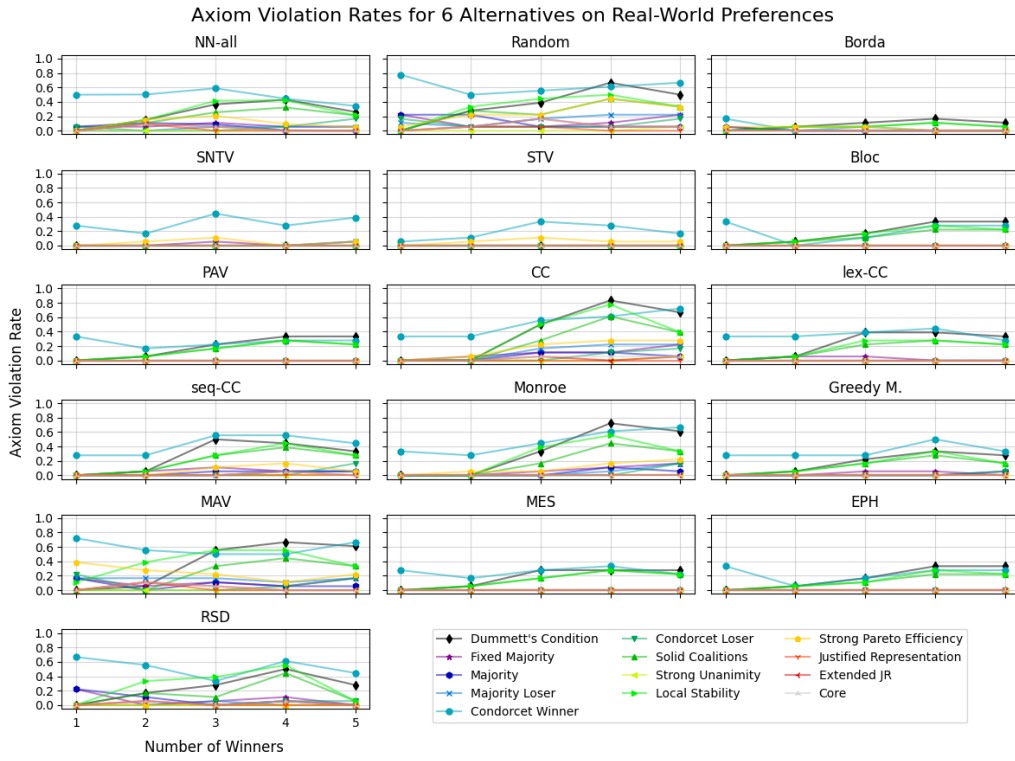


Figure 9: Axiom violation rate for each rule on Real World Data preferences with 6 alternatives.

B.3 7 Alternatives, Real World Data

| Method | Mean | Maj W | Maj L | Cond W | Cond L | Pareto | F Maj | Unanimity | Dummett's | JR | EJR | Core | S. Coalitions | Stability |
|-----------|-------------|----------|----------|-------------|----------|----------|----------|-----------|-------------|----------|----------|----------|---------------|-----------|
| NN-all | .102 | .016 | .039 | .457 | .015 | .108 | .041 | 0 | .200 | .035 | .042 | .046 | .110 | .223 |
| Min | .010 | 0 | 0 | .068 | 0 | .008 | 0 | 0 | .038 | 0 | 0 | 0 | 0 | .015 |
| Max | .534 | .409 | .318 | .909 | .712 | .750 | .235 | .061 | .750 | .288 | .432 | .462 | .697 | .917 |
| Borda | .030 | .008 | 0 | .136 | 0 | 0 | .015 | 0 | .121 | 0 | 0 | 0 | .045 | .068 |
| EPH | .033 | 0 | 0 | .152 | 0 | 0 | 0 | 0 | .136 | 0 | 0 | 0 | .061 | .083 |
| SNTV | .086 | 0 | .061 | .598 | 0 | .167 | .129 | .023 | .076 | 0 | .030 | .030 | 0 | 0 |
| Bloc | .033 | 0 | 0 | .144 | 0 | 0 | 0 | 0 | .129 | 0 | 0 | 0 | .061 | .091 |
| STV | .027 | 0 | .015 | .258 | 0 | .023 | .030 | 0 | .023 | 0 | 0 | 0 | 0 | 0 |
| PAV | .034 | 0 | 0 | .152 | 0 | 0 | 0 | 0 | .144 | 0 | 0 | 0 | .061 | .083 |
| MES | .046 | 0 | 0 | .258 | 0 | 0 | 0 | 0 | .159 | 0 | 0 | 0 | .076 | .106 |
| CC | .177 | .053 | .174 | .674 | .008 | .318 | .159 | .038 | .348 | 0 | .091 | .098 | .159 | .182 |
| seq-CC | .184 | .038 | .068 | .712 | .015 | .242 | .152 | .030 | .417 | 0 | .068 | .076 | .258 | .318 |
| lex-CC | .054 | 0 | 0 | .326 | 0 | 0 | .023 | 0 | .182 | 0 | 0 | 0 | .076 | .098 |
| Monroe | .097 | .008 | .068 | .545 | 0 | .152 | .076 | 0 | .235 | 0 | 0 | 0 | .076 | .106 |
| Greedy M. | .088 | .008 | .015 | .477 | .008 | .015 | .061 | 0 | .227 | 0 | 0 | 0 | .144 | .189 |
| MAV | .149 | .068 | .106 | .712 | .038 | .235 | .167 | 0 | .250 | 0 | .023 | .023 | .114 | .205 |
| RSD | .091 | .030 | .030 | .523 | 0 | 0 | .015 | 0 | .220 | .008 | .015 | .015 | .144 | .189 |
| Random | .303 | .197 | .152 | .841 | .053 | .523 | .227 | .061 | .500 | .106 | .174 | .220 | .371 | .515 |

Table 8: Average Axiom Violation Rate for 7 alternatives and $1 \leq k < 7$ winners across Real World Data preferences.

| | Borda | EPH | SNTV | Bloc | STV | PAV | MES | CC | seq-CC | lex-CC | Monroe | Greedy M. | MAV | RSD |
|-----------|-------|------|------|------|------|------|------|------|--------|--------|--------|-----------|------|------|
| Borda | 0 | - | - | - | - | - | - | - | - | - | - | - | - | - |
| EPH | .150 | 0 | - | - | - | - | - | - | - | - | - | - | - | - |
| SNTV | .422 | .409 | 0 | - | - | - | - | - | - | - | - | - | - | - |
| Bloc | .154 | .008 | .413 | 0 | - | - | - | - | - | - | - | - | - | - |
| STV | .235 | .254 | .333 | .254 | 0 | - | - | - | - | - | - | - | - | - |
| PAV | .152 | .029 | .419 | .033 | .256 | 0 | - | - | - | - | - | - | - | - |
| MES | .203 | .129 | .365 | .133 | .321 | .121 | 0 | - | - | - | - | - | - | - |
| CC | .527 | .455 | .581 | .458 | .496 | .449 | .520 | 0 | - | - | - | - | - | - |
| seq-CC | .520 | .491 | .367 | .495 | .545 | .475 | .409 | .722 | 0 | - | - | - | - | - |
| lex-CC | .220 | .119 | .426 | .122 | .295 | .092 | .145 | .431 | .472 | 0 | - | - | - | - |
| Monroe | .433 | .352 | .514 | .356 | .396 | .342 | .415 | .133 | .628 | .328 | 0 | - | - | - |
| Greedy M. | .345 | .298 | .376 | .306 | .369 | .293 | .212 | .597 | .374 | .306 | .489 | 0 | - | - |
| MAV | .539 | .525 | .652 | .525 | .566 | .520 | .557 | .312 | .770 | .497 | .332 | .646 | 0 | - |
| RSD | .400 | .367 | .515 | .367 | .396 | .376 | .393 | .556 | .571 | .390 | .465 | .461 | .556 | 0 |
| Random | .758 | .732 | .717 | .732 | .721 | .739 | .739 | .773 | .725 | .739 | .768 | .753 | .739 | .725 |
| Min | .181 | .199 | .427 | .196 | .172 | .198 | .265 | .473 | .586 | .253 | .383 | .400 | .508 | .400 |
| Max | .970 | .975 | .914 | .971 | .934 | .972 | .977 | .778 | .848 | .944 | .854 | .937 | .789 | .900 |
| NN-all | .427 | .440 | .507 | .436 | .444 | .440 | .453 | .693 | .492 | .447 | .617 | .465 | .656 | .509 |

Table 9: Difference between rules for 7 alternatives with $1 \leq k < 7$ on Real-World preferences.




Figure 10: Axiom violation rate for each axiom on Real World Data preferences with 7 alternatives.

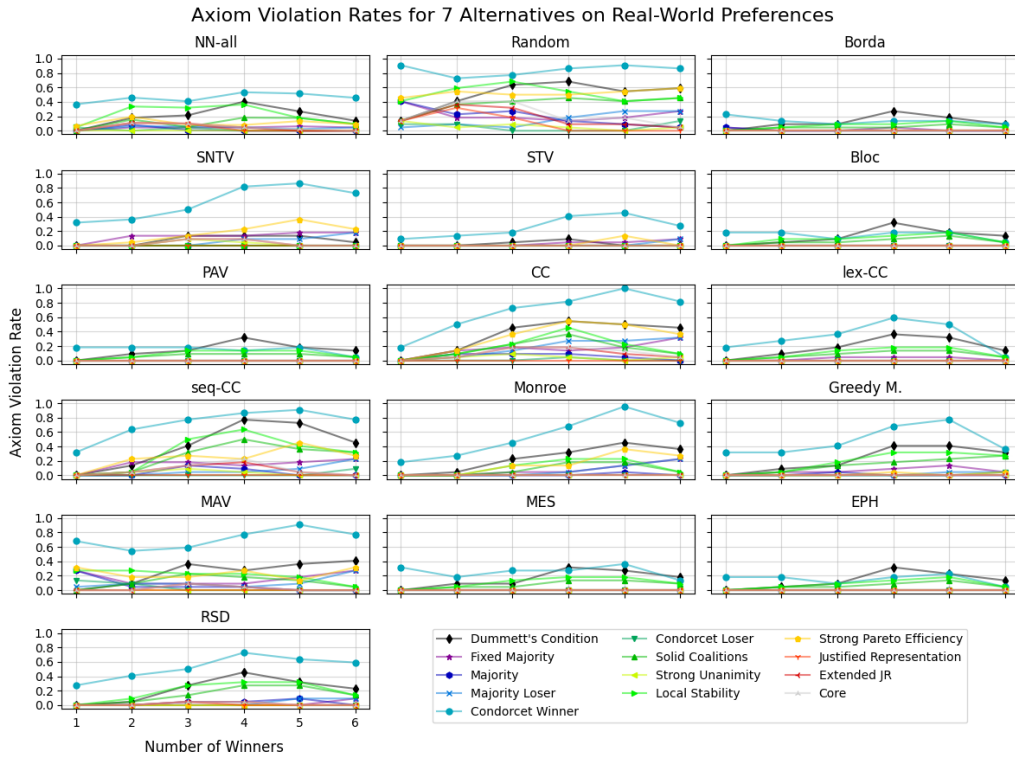


Figure 11: Axiom violation rate for each rule on Real World Data preferences with 7 alternatives.

C Additional Experimental Results

This section includes some additional results for our experiments: the figures matching the body of our paper corresponding to $m = 5$ and $m = 6$ alternatives, as well as exact numerical results for the optimal positional scoring rule experiment.

C.1 Optimized Positional Scoring Rules

Table 10 shows results of annealing for 1000 steps using 2000 profiles sampled from the Mixed distribution, for each axiom set. Each vector is then evaluated on a 25000 profiles also sampled from the Mixed distribution. The score vector resulting from annealing is shown in the final column.

| All Axioms | | | |
|----------------|------------------------------|----------------------------|--------------------------------------|
| k | $\mathcal{F}^{\text{Borda}}$ | \mathcal{F}^{Opt} | Annealed Vector |
| 1 | 0.010 | 0.008 | (1, 0.69, 0.51, 0.30, 0.08, 0.05, 0) |
| 2 | 0.011 | 0.011 | (1, 0.82, 0.50, 0.45, 0.10, 0.01, 0) |
| 3 | 0.022 | 0.018 | (1, 0.77, 0.71, 0.24, 0.16, 0.08, 0) |
| 4 | 0.030 | 0.025 | (1, 0.74, 0.51, 0.46, 0.12, 0.09, 0) |
| 5 | 0.034 | 0.028 | (1, 0.50, 0.44, 0.32, 0.25, 0.05, 0) |
| 6 | 0.025 | 0.022 | (1, 0.54, 0.42, 0.42, 0.37, 0.30, 0) |
| Reduced Axioms | | | |
| k | $\mathcal{F}^{\text{Borda}}$ | \mathcal{F}^{Opt} | Annealed Vector |
| 1 | 0.019 | 0.016 | (1, 0.70, 0.42, 0.36, 0.19, 0.05, 0) |
| 2 | 0.021 | 0.018 | (1, 0.87, 0.54, 0.41, 0.26, 0.01, 0) |
| 3 | 0.036 | 0.032 | (1, 0.85, 0.70, 0.38, 0.13, 0.02, 0) |
| 4 | 0.044 | 0.043 | (1, 0.63, 0.60, 0.51, 0.22, 0.05, 0) |
| 5 | 0.042 | 0.042 | (1, 0.55, 0.47, 0.41, 0.38, 0.08, 0) |
| 6 | 0.030 | 0.027 | (1, 0.71, 0.50, 0.48, 0.43, 0.31, 0) |

Table 10: Score vector resulting from annealing and the axiom violation rate across a test set of 25000 profiles for $m = 7$ alternatives on both axiom sets by $\mathcal{F}^{\text{Borda}}$ and the positional scoring rule found via simulated annealing, \mathcal{F}^{Opt} .

C.2 Results for 5 Alternatives

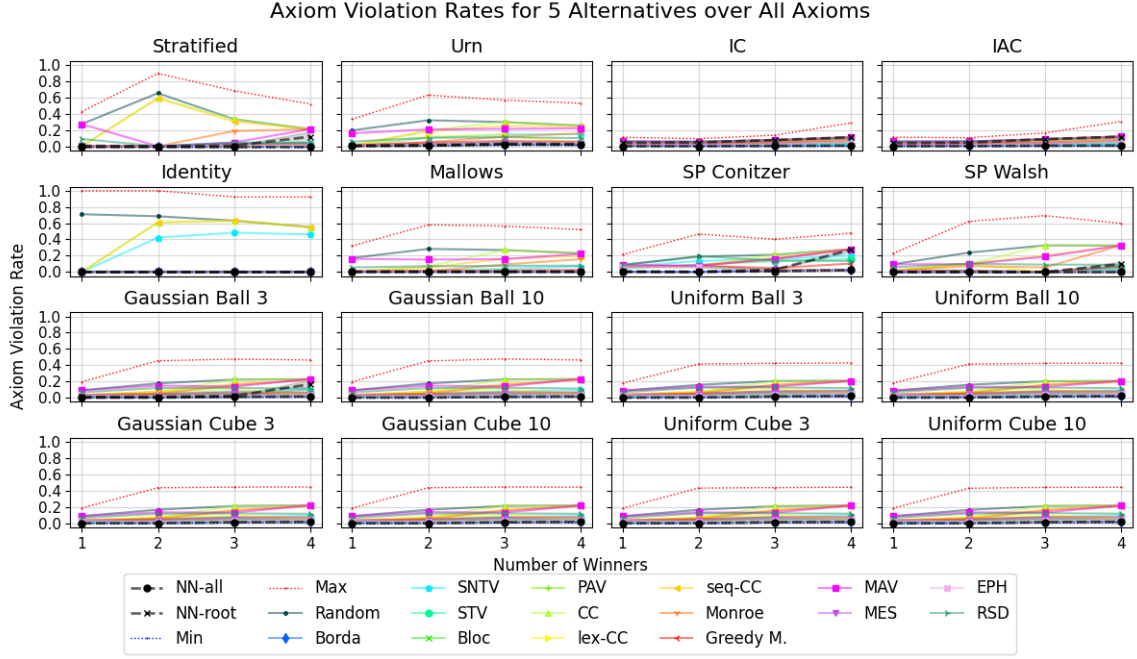


Figure 12: Axiom violation rates for each rule under each individual preference distribution for $m = 5$. In all cases, $\mathcal{F}^{\text{NN-all}}$ has AVR lower than, or similar to, other rules.

| Method | Mean | Maj W | Maj L | Cond W | Cond L | Pareto | F Maj | Unanimity | Dummett's | JR | EJR | Core | S. Coalitions | Stability |
|-----------|------|-------|-------|--------|--------|--------|-------|-----------|-----------|------|------|------|---------------|-----------|
| NN-all | .006 | .000 | .000 | .001 | .000 | .002 | .000 | 0 | .027 | .000 | .000 | .000 | .017 | .025 |
| NN-root | .022 | .002 | .020 | .131 | .013 | .017 | .023 | 0 | .030 | .000 | .000 | .000 | .021 | .028 |
| Min | .003 | 0 | .000 | .013 | .000 | .000 | .000 | 0 | .009 | 0 | 0 | 0 | .004 | .012 |
| Max | .430 | .218 | .425 | .959 | .753 | .458 | .268 | .082 | .478 | .196 | .324 | .324 | .443 | .665 |
| Borda | .012 | .002 | .007 | .096 | 0 | .002 | .016 | 0 | .015 | .000 | .000 | .000 | .009 | .011 |
| EPH | .025 | .000 | .003 | .205 | .005 | 0 | .001 | 0 | .044 | .000 | .000 | .000 | .030 | .038 |
| SNTV | .078 | 0 | .109 | .484 | .016 | .132 | .111 | .038 | .044 | .000 | .041 | .041 | 0 | .003 |
| Bloc | .024 | .000 | .003 | .197 | .005 | 0 | 0 | 0 | .043 | .000 | .000 | .000 | .030 | .040 |
| STV | .034 | 0 | .042 | .298 | .005 | .053 | .039 | 0 | 0 | .000 | .000 | .000 | 0 | .000 |
| PAV | .029 | .001 | .003 | .242 | .005 | 0 | .006 | 0 | .046 | 0 | 0 | 0 | .033 | .036 |
| MES | .031 | .001 | .004 | .256 | .005 | .001 | .009 | 0 | .049 | 0 | 0 | 0 | .035 | .037 |
| CC | .164 | .054 | .163 | .660 | .069 | .217 | .179 | .057 | .230 | 0 | .071 | .071 | .166 | .196 |
| seq-CC | .148 | .046 | .156 | .627 | .044 | .182 | .175 | .057 | .208 | 0 | .059 | .059 | .145 | .167 |
| lex-CC | .047 | .010 | .013 | .362 | .005 | 0 | .033 | 0 | .072 | 0 | 0 | 0 | .055 | .057 |
| Monroe | .098 | .016 | .090 | .525 | .048 | .125 | .086 | 0 | .136 | 0 | .001 | .001 | .111 | .130 |
| Greedy M. | .045 | .002 | .026 | .348 | .006 | .006 | .033 | 0 | .063 | 0 | 0 | 0 | .046 | .051 |
| MAV | .134 | .054 | .135 | .700 | .096 | .128 | .150 | 0 | .147 | .013 | .015 | .015 | .118 | .176 |
| RSD | .089 | .023 | .085 | .536 | .040 | 0 | .070 | 0 | .095 | .023 | .024 | .024 | .074 | .161 |
| Random | .218 | .109 | .215 | .804 | .127 | .285 | .225 | .071 | .250 | .050 | .111 | .112 | .191 | .287 |

Table 11: Average Axiom Violation Rate for 5 alternatives and $1 \leq k < 5$ winners across all preferences.

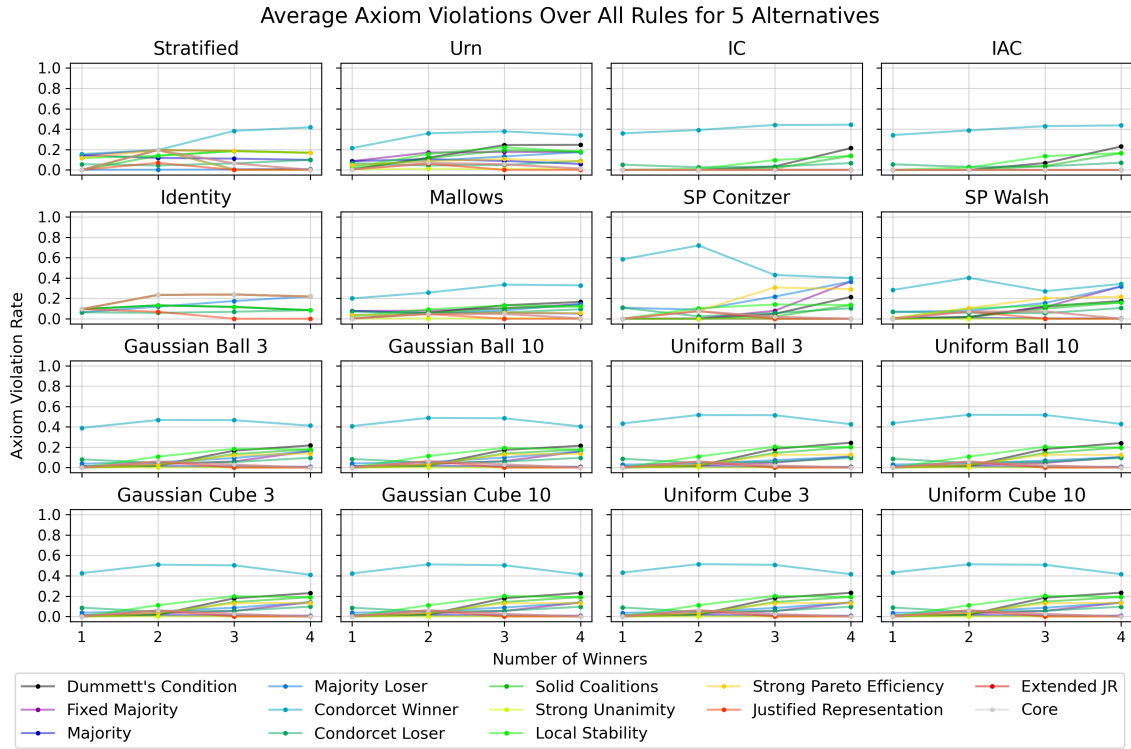


Figure 13: Violation rates of each axiom for every preference distribution, averaged over all rules. This shows trends common across all rules in the relationships between axioms and preference distributions.

| | Borda | EPH | SNTV | Bloc | STV | PAV | MES | CC | seq-CC | lex-CC | Monroe | Greedy M. | MAV | RSD |
|-----------|-------|------|------|------|------|------|------|------|--------|--------|--------|-----------|------|------|
| Borda | 0 | - | - | - | - | - | - | - | - | - | - | - | - | - |
| EPH | .227 | 0 | - | - | - | - | - | - | - | - | - | - | - | - |
| SNTV | .394 | .355 | 0 | - | - | - | - | - | - | - | - | - | - | - |
| Bloc | .229 | .011 | .353 | 0 | - | - | - | - | - | - | - | - | - | - |
| STV | .245 | .303 | .261 | .303 | 0 | - | - | - | - | - | - | - | - | - |
| PAV | .235 | .045 | .365 | .053 | .309 | 0 | - | - | - | - | - | - | - | - |
| MES | .239 | .096 | .344 | .103 | .315 | .066 | 0 | - | - | - | - | - | - | - |
| CC | .549 | .429 | .524 | .432 | .521 | .405 | .439 | 0 | - | - | - | - | - | - |
| seq-CC | .507 | .424 | .409 | .429 | .505 | .405 | .367 | .586 | 0 | - | - | - | - | - |
| lex-CC | .293 | .119 | .397 | .125 | .350 | .079 | .116 | .367 | .396 | 0 | - | - | - | - |
| Monroe | .452 | .329 | .453 | .332 | .427 | .310 | .345 | .124 | .541 | .318 | 0 | - | - | - |
| Greedy M. | .301 | .195 | .363 | .201 | .348 | .171 | .140 | .444 | .343 | .194 | .365 | 0 | - | - |
| MAV | .594 | .576 | .667 | .577 | .595 | .563 | .582 | .376 | .743 | .528 | .352 | .594 | 0 | - |
| RSD | .472 | .450 | .553 | .450 | .499 | .453 | .454 | .614 | .593 | .471 | .542 | .467 | .610 | 0 |
| Random | .700 | .700 | .700 | .700 | .700 | .700 | .700 | .700 | .700 | .700 | .700 | .700 | .700 | .700 |
| Min | .125 | .211 | .407 | .207 | .249 | .230 | .246 | .521 | .525 | .289 | .424 | .310 | .562 | .470 |
| Max | .964 | .948 | .897 | .951 | .936 | .935 | .930 | .808 | .785 | .910 | .867 | .911 | .821 | .852 |
| NN-all | .125 | .201 | .413 | .197 | .256 | .221 | .239 | .522 | .524 | .281 | .425 | .306 | .563 | .468 |
| NN-root | .248 | .286 | .443 | .285 | .314 | .295 | .309 | .480 | .555 | .333 | .348 | .342 | .453 | .456 |

Table 12: Difference between rules for 5 alternatives with $1 \leq k < 5$ averaged over all preference distributions.

C.3 Results for 6 Alternatives

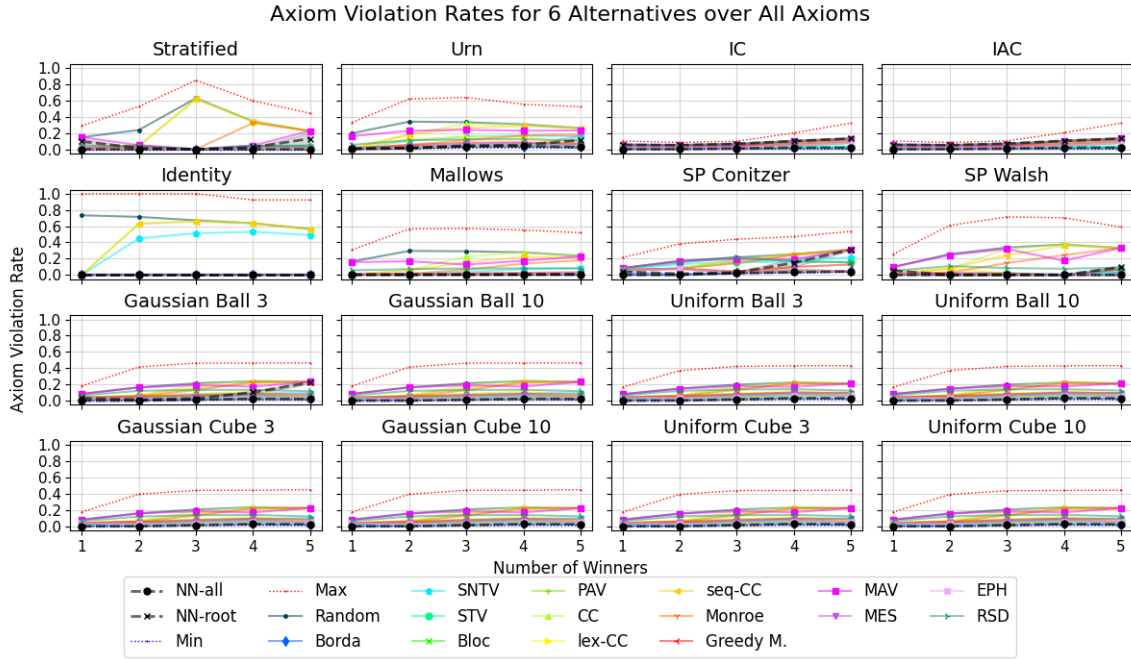


Figure 14: Axiom violation rates for each rule under each individual preference distribution for $m = 6$. In all cases, $\mathcal{F}^{\text{NN-all}}$ has AVR lower than, or similar to, other rules.

| Method | Mean | Maj W | Maj L | Cond W | Cond L | Pareto | F Maj | Unanimity | Dummett's | JR | EJR | Core | S. Coalitions | Stability |
|-----------|------|-------|-------|--------|--------|--------|-------|-----------|-----------|------|------|------|---------------|-----------|
| NN-all | .011 | .000 | .000 | .006 | .000 | .003 | .000 | 0 | .042 | .000 | .000 | .000 | .032 | .055 |
| NN-root | .030 | .001 | .024 | .180 | .012 | .029 | .022 | .000 | .039 | .000 | .000 | .000 | .033 | .056 |
| Min | .005 | 0 | .000 | .024 | 0 | .001 | .000 | 0 | .009 | 0 | 0 | 0 | .006 | .031 |
| Max | .431 | .136 | .375 | .940 | .688 | .553 | .203 | .078 | .520 | .242 | .311 | .327 | .487 | .742 |
| Borda | .016 | .002 | .006 | .113 | 0 | .003 | .013 | 0 | .027 | .000 | .000 | .000 | .019 | .029 |
| EPH | .032 | .000 | .002 | .234 | .003 | .000 | .001 | 0 | .060 | .000 | .000 | .000 | .049 | .067 |
| SNTV | .090 | 0 | .103 | .559 | .011 | .183 | .108 | .045 | .054 | .000 | .048 | .049 | 0 | .007 |
| Bloc | .031 | .000 | .001 | .223 | .003 | 0 | 0 | 0 | .059 | .000 | .000 | .000 | .048 | .073 |
| STV | .041 | 0 | .039 | .373 | .002 | .085 | .032 | 0 | 0 | .000 | .000 | .000 | 0 | .000 |
| PAV | .035 | .001 | .002 | .268 | .003 | 0 | .005 | 0 | .064 | 0 | 0 | 0 | .052 | .064 |
| MES | .039 | .001 | .003 | .301 | .003 | .002 | .008 | 0 | .071 | 0 | 0 | 0 | .057 | .066 |
| CC | .179 | .036 | .154 | .719 | .041 | .291 | .156 | .060 | .264 | 0 | .077 | .080 | .199 | .246 |
| seq-CC | .167 | .032 | .148 | .698 | .032 | .242 | .155 | .060 | .248 | 0 | .072 | .072 | .184 | .224 |
| lex-CC | .054 | .006 | .009 | .402 | .003 | 0 | .027 | 0 | .093 | 0 | 0 | 0 | .076 | .085 |
| Monroe | .115 | .007 | .084 | .599 | .035 | .188 | .072 | 0 | .174 | 0 | .003 | .004 | .149 | .181 |
| Greedy M. | .054 | .002 | .021 | .400 | .004 | .010 | .026 | 0 | .086 | 0 | 0 | 0 | .070 | .086 |
| MAV | .148 | .028 | .128 | .732 | .063 | .223 | .104 | 0 | .178 | .020 | .026 | .026 | .149 | .250 |
| RSD | .096 | .010 | .068 | .571 | .024 | 0 | .046 | 0 | .119 | .023 | .024 | .025 | .100 | .234 |
| Random | .226 | .068 | .189 | .833 | .084 | .354 | .180 | .071 | .285 | .048 | .119 | .123 | .222 | .359 |

Table 13: Average Axiom Violation Rate for 6 alternatives and $1 \leq k < 6$ winners across all preferences.

| | Borda | EPH | SNTV | Bloc | STV | PAV | MES | CC | seq-CC | lex-CC | Monroe | Greedy M. | MAV | RSD |
|-----------|-------|------|------|------|------|------|------|------|--------|--------|--------|-----------|------|------|
| Borda | 0 | - | - | - | - | - | - | - | - | - | - | - | - | - |
| EPH | .238 | 0 | - | - | - | - | - | - | - | - | - | - | - | - |
| SNTV | .427 | .385 | 0 | - | - | - | - | - | - | - | - | - | - | - |
| Bloc | .239 | .016 | .382 | 0 | - | - | - | - | - | - | - | - | - | - |
| STV | .278 | .330 | .280 | .330 | 0 | - | - | - | - | - | - | - | - | - |
| PAV | .243 | .047 | .394 | .058 | .334 | 0 | - | - | - | - | - | - | - | - |
| MES | .255 | .107 | .376 | .117 | .342 | .076 | 0 | - | - | - | - | - | - | - |
| CC | .569 | .449 | .548 | .453 | .542 | .429 | .460 | 0 | - | - | - | - | - | - |
| seq-CC | .538 | .453 | .437 | .460 | .533 | .435 | .399 | .608 | 0 | - | - | - | - | - |
| lex-CC | .301 | .126 | .421 | .135 | .371 | .085 | .115 | .397 | .424 | 0 | - | - | - | - |
| Monroe | .484 | .363 | .490 | .366 | .460 | .343 | .375 | .111 | .568 | .344 | 0 | - | - | - |
| Greedy M. | .318 | .210 | .394 | .217 | .374 | .187 | .155 | .470 | .372 | .202 | .397 | 0 | - | - |
| MAV | .602 | .589 | .673 | .588 | .604 | .585 | .600 | .352 | .781 | .567 | .350 | .619 | 0 | - |
| RSD | .476 | .454 | .565 | .453 | .508 | .456 | .458 | .624 | .608 | .472 | .557 | .472 | .614 | 0 |
| Random | .700 | .700 | .700 | .700 | .700 | .700 | .700 | .700 | .700 | .700 | .700 | .700 | .700 | .700 |
| Min | .142 | .236 | .439 | .230 | .283 | .252 | .272 | .542 | .560 | .310 | .457 | .335 | .564 | .476 |
| Max | .941 | .933 | .870 | .937 | .914 | .924 | .916 | .802 | .762 | .900 | .855 | .896 | .810 | .846 |
| NN-all | .137 | .221 | .445 | .215 | .291 | .238 | .260 | .544 | .557 | .298 | .459 | .327 | .564 | .472 |
| NN-root | .327 | .347 | .499 | .344 | .377 | .353 | .367 | .523 | .622 | .390 | .402 | .402 | .446 | .472 |

Table 14: Difference between rules for 6 alternatives with $1 \leq k < 6$ averaged over all preference distributions.

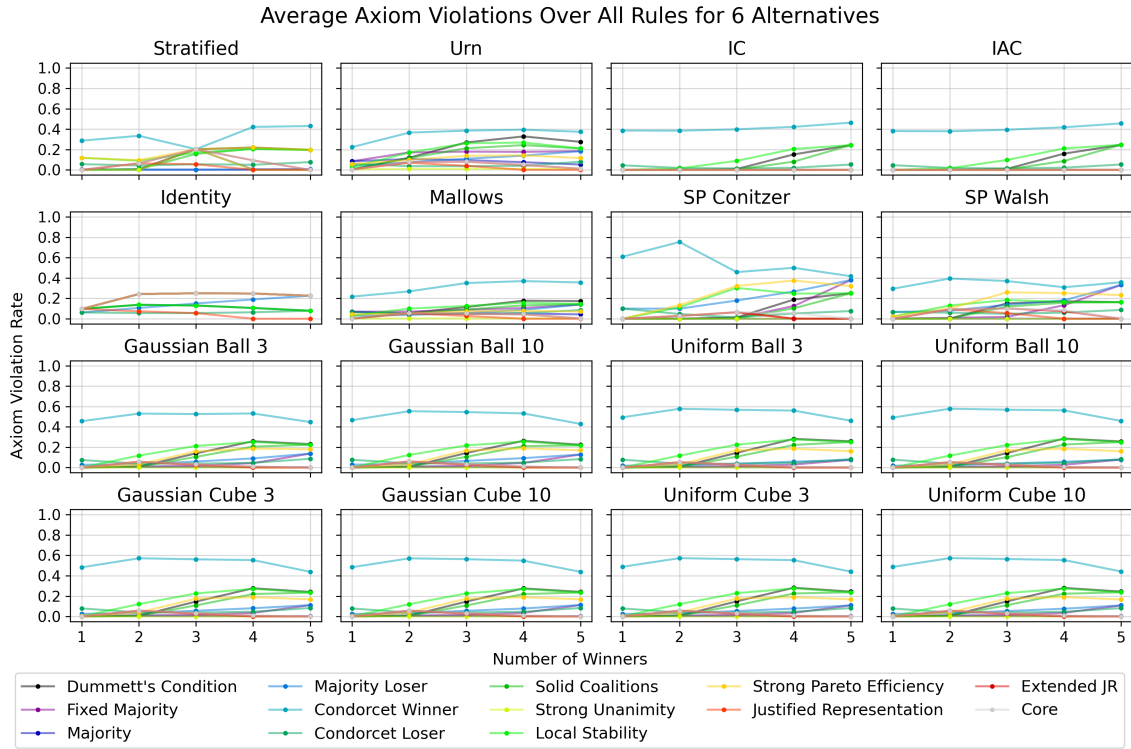


Figure 15: Violation rates of each axiom for every preference distribution, averaged over all rules. This shows trends common across all rules in the relationships between axioms and preference distributions.