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

Joshua Caiata[1], **Ben Armstrong**[2], **Kate Larson**[1]

[1]University of Waterloo
[2]Tulane University
jcaiata@uwaterloo.ca, barmstrong@tulane.edu, kate.larson@uwaterloo.ca

Abstract

Committee-selection problems arise in many contexts and applications; there has been increasing interest in the social choice research community on identifying which properties are satisfied by different multi-winner voting rules. 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. We use this framework to analyze the relationship between multi-winner voting rules and their axiomatic performance under several preference distributions, and propose a method for systematically minimizing axioms 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 subset of alternatives based on their preferences (Lackner and Skowron 2023; Faliszewski et al. 2017). There are numerous properties or axioms we might wish a multi-winner voting rule to satisfy; however, many combinations of axioms are known to be impossible to satisfy simultaneously. Traditional research often focuses on ascertaining which combinations are possible or are satisfied by a voting rule. Such work is motivated by a desire to identify rules which universally satisfy axioms desirable to a particular setting. However, in many instances, a voting rule may not satisfy an axiom, yet rarely violates it in practice.

We propose a data-driven framework to evaluate and explore voting rules axiomatically. To do so, we define a measure of axiom violation more fine-grained than binary satisfaction, moving away from worst-case analysis and towards an average-case evaluation model. We apply this framework to explore the relationship between voter preference distributions and axiomatic properties within the context of common multi-winner voting rules, allowing us to quantify the practical trade-offs that classical theory hides. We consider how to systematically reduce axiom violations across voter preference distributions using machine learning, showing that this approach yields voting rules that fare much better than most traditional voting rules.

Specifically, we make the following contributions:

1. We develop a novel data-driven measure of axiomatic violations, which demonstrate the high sensitivity of many multi-winner voting rules to changes in underlying voter preferences.
2. We identify novel similarities and unexpected differences between existing multi-winner rules in terms of the committees they select and how often they violate axioms.
3. We demonstrate that machine learning can very effectively discover novel multi-winner voting rules with compelling axiomatic performance.

In summary, this paper combines a data-driven approach with machine learning to deepen our understanding of social choice and inform the design of new voting systems. This work is supported by the Supplementary Material in our extended paper (Caiata, Armstrong, and Larson 2025).

1.1 Related Work

Much prior work has developed novel axioms to describe desirable properties, or shown which voting rules satisfy which axioms. Of relevance to our work, Elkind et al. describes axioms and satisfaction results for several multi-winner voting rules (Elkind et al. 2017). 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, those alternatives should be in the winning committee) (Dummett 1984). A similar analysis is taken for rules and axioms based on approval preferences by Lackner and Skowron (2023); Peters and Skowron (2020).

Our approach builds on data-driven axiomatic analysis (d'Eon and Larson 2020; Fairstein, Vilenchik, and Gal 2024), using axioms from the social choice literature. Our results vary greatly based on underlying voter preferences; these are well-studied for the single-winner setting (Elkind 2018) but less is known for the multi-winner setting. We use well-studied distributions shown to approximate human preferences or explore restricted cases (Boehmer et al. 2021). 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 (Faliszewski et al. 2023).

Other recent work has explored the application of machine learning to social choice. Prior work focuses primarily on approximating single-winner rules (Matone et al. 2025; Burka et al. 2022; Kujawska, Slavkovik, and Rückmann 2020). Existing work on learning new rules, primarily in single-winner settings, has studied learning rules under specific axioms (Armstrong and Larson 2019; Hornischer and Terzopoulou 2025), reducing susceptibility to manipulation (Holliday, Kristoffersen, and Pacuit 2025), or maximizing utility without axiomatic focus (Anil and Bao 2021). One paper explored learning multi-winner rules for participatory budgeting with a focus on measures of social welfare (Fairstein, Vilenchik, and Gal 2024). Our work, instead, studies how often voting rules violate axioms in practice and how one might systematically reduce such violations.

2 Social Choice Cornerstones

Traditional research in social choice is based upon three cornerstones: voting rules, voter preferences, and measures of outcome quality. Voting rules aggregate voter preferences following a procedure aimed at maximizing quality. Rules may provide a single winner, a ranking over alternatives, or a set of multiple winners. Voter preferences are the subject of aggregation; often modelled as originating from some fixed distribution or based upon empirical data. Measures of voting rule quality primarily fall into one of two paradigms: *social welfare*, or *axiomatic*. This paper considers rules which elect multiple winners, sampling ordinal or approval preferences from a variety of distributions, and measure quality using a novel measure of axiom violations. We now provide the basic notation used through the paper, and defer full definitions to (Caiata, Armstrong, and Larson 2025).

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 each voter's ranked preferences, we also consider the alternatives of which they most approve. Let $App(v) \subset M$ be the *approval set* of voter $v \in V$, containing the k most preferred alternatives of v, with no information about v's relative preferences over these alternatives. We let $P_{App} = (App(v_1), \dots, App(v_n))$, and when it is clear from the context we refer to either ordinal preference profiles or approval sets as preference profiles 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 associated P and k, $1 \leq k < m$, $\mathcal{F}_k(E) \subseteq \{C \mid C \subset M, |C| = k\}$ is a multi-winner voting rule that returns a family of k-sized subsets of M, called the winning *committees*. If $\mathcal{F}_k$ 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 focus on *resolute* multi-winner voting rules (those with a single winning committee) and assume each rule uses lexicographic tie-breaking to return a single committee.

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 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.

Ordinal Voting Rules

- k-Borda
- Single Non-Transferable Vote (SNTV)
- Single Transferable Vote (STV)

Approval-Based Voting Rules

- Bloc
- Proportional Approval Voting (PAV)
- Chamberlin-Courant (CC), sequential Chamberlin-Courant (seq-CC), lexicographic Chamberlin-Courant (lex-CC)
- Monroe, Greedy Monroe
- Minimax Approval (MAV)
- Method of Equal Shares (MES)
- E Pluribus Hugo (EPH)
- Random Serial Dictator (RSD)

Faliszewski et al. (2017) 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:

Individual Excellence: $\mathcal{F}^{\text{Borda}}$, $\mathcal{F}^{\text{SNTV}}$, $\mathcal{F}^{\text{Bloc}}$, $\mathcal{F}^{\text{EPH}}$
Diversity: $\mathcal{F}^{\text{SNTV}}$, $\mathcal{F}^{\text{CC}}$
Proportionality: $\mathcal{F}^{\text{STV}}$, $\mathcal{F}^{\text{PAV}}$, $\mathcal{F}^{\text{Monroe}}$, $\mathcal{F}^{\text{CC}}$, $\mathcal{F}^{\text{MES}}$, $\mathcal{F}^{\text{EPH}}$

These are subjective categorizations and are not mutually exclusive. For example, $\mathcal{F}^{CC}$ has been described as both *diverse* (Faliszewski et al. 2023) and *proportional* (Elkind et al. 2017). $\mathcal{F}^{MAV}$ 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). We also incorporate three quasi-rules $\mathcal{F}^{\text{Min}}$ and $\mathcal{F}^{\text{Max}}$ which always elect the committee which violate (respectively) the least and the most axioms possible, as well as $\mathcal{F}^{\text{Random}}$ which elects a committee uniformly at random.

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

(Schmidtlein 2022). Axioms can also be loosely categorized based on the stated priorities of their definitions.

Axioms of Individual Excellence

- Majority Winner/Loser (Fishburn 1977)
- Condorcet Winner/Loser (Gehrlein 1985)
- Strong Pareto Efficiency (Lackner and Skowron 2023)
- Fixed Majority (Debord 1993; Elkind et al. 2017)
- Strong Unanimity (Elkind et al. 2017)

Axioms of Diversity and Proportionality

- Local Stability (Aziz et al. 2017)
- Dummett's Condition (Dummett 1984)
- Solid Coalitions (Elkind et al. 2017)
- Core (Lackner and Skowron 2023)
- Justified Representation, Extended Justified Representation (Lackner and Skowron 2023)

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.

Many of these axioms are related to one another; we provide a study of the relationship between these axioms in (Caiata, Armstrong, and Larson 2025). While many of these are known we also contribute several novel findings. We have chosen one set of axioms to study; many other sets of axioms are also of interest. Thus, our results aim to inform both about these specific axioms but also, more broadly, provide a framework for this approach to empirical analysis.

2.3 Preference Distributions

Our experiments consider a wide range of standard preference distributions ($\mathcal{D}$): 8 families of distribution, as well as two additional sets of preferences. These are used both for training novel rules, as well as testing novel and existing rules. We generate ordinal ballots and, when needed by a rule, convert them to approvals where each voter approves of their k top-ranked alternatives. There are many other valid methods of generating approvals which we do not explore. We categorize distributions by the amount of structure within the profiles 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 (Guilbaud 1952).
- Impartial Anonymous Culture (Kuga and Nagatani 1974; Gehrlein and Fishburn 1976).

Moderately Structured Distributions

- Mallows (Mallows 1957; Boehmer et al. 2021)
- Urn (Eggenberger and Pólya 1923).
- Euclidean (Enelow and Hinich 1984).

Highly Structured Distributions

- Identity: All voters in V have identical preferences.
- Single-Peaked (Conitzer 2007; Walsh 2015).
- Stratification (Boehmer et al. 2021).

Additional Preference Distributions

- Mixed: Profiles sampled equally from each artificial distribution listed above.
- PrefLib: Empirical voter preferences hosted on PrefLib (Mattei and Walsh 2013).

3 Data-Driven Analysis

Traditionally, axiomatic analysis of voting rules asks if a rule universally satisfies some axiom. In empirical settings, however, *the gap between worst-case analysis and average-case behaviour can be large*. The key insight, and motivation, for a shift in voting rule analysis is that a fine-grained and empirical lens reveals real, meaningful differences in rule behaviour across distributions that are invisible when only looking at axiomatic satisfaction as a binary question.

We introduce two measures that we use to understand the behaviour of voting rules. Our data-driven analysis is 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 to capture similarity in rule behaviour, separated from differences in the algorithm rules may follow.

3.1 Axiom Violation Rate

Axiom violation rate (AVR) is 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. 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 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 normalization factor $\delta = \frac{m}{m - |m-2k|}$ which ensures the difference between two rules on a given set of profiles $\mathbb{P}$ ranges from 0 to 1.

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

3.3 Learning Voting Rules

We learn two versions of $\mathcal{F}^{\text{NN}}$, a novel voting rule built atop a multi-layer perceptron that is optimized to minimize axiom violation rate. Here we outline the procedure we use to train the model. Specific details on the parameters used to train $\mathcal{F}^{\text{NN}}$ in our experiments are included in Section 4. Our data generation and learning use the following procedure:

1. **Generate Axiom Violation Data**
   i Generate separate training and testing sets by sampling profiles from some distribution $\mathcal{D}$. Randomly rename alternatives to ensure profiles satisfy neutrality.
   ii Find, by exhaustive search, the committee $c = \arg\min_c \sum_{A \in \mathbb{A}} A(P, c)$ which minimizes axiom violations in $\mathbb{A}$. Use lexicographic tie-breaking.
2. **Transform Profiles Into Training Data** To allow our learning process to scale to arbitrary numbers of voters, transform each profile generated in the previous step into three matrices, which are flattened, concatenated, and normalized to form input features.

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

   Weighted Preference Matrix An $m \times m$ matrix $\mathcal{R}^{\text{w}}$ with $\mathcal{R}^{\text{w}}_{ij} = c$ to indicate that c voters prefer a_i over a_j.

   Ranking Matrix An $m \times m$ matrix $\mathcal{R}^{\text{r}}$ where $\mathcal{R}^{\text{r}}_{ij} = c$ indicates that c voters place a_i in rank j.
3. **Learn Novel Rules** Each profile corresponds to a single set of concatenated matrices and a committee that minimizes violations of axioms we are interested in. We use the **L1 Loss** function to train $\mathcal{F}^{\text{NN}}$ on this data.

4 Experimental Results and Analysis

We now describe our experiments, our 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.

4.1 Experimental Parameters

We run experiments to test all rules on all combinations of $m = \{5, 6, 7\}$ alternatives and $k = \{1, 2, ..., m - 1\}$ winners. In all cases we use profiles with $n = 50$ voters[1]. All rules are evaluated on 25,000 profiles. All experiments ran on the Digital Research Alliance of Canada's Graham cluster. We restrict our experiments to a maximum of 7 alternatives as some of the axioms we consider (e.g., the Core) require significant time to evaluate. Our focus is primarily on describing our novel methodology for analyzing voting rules empirically; the issue of scalability is left as future work. We introduce two configurations of our learned rule $\mathcal{F}^{\text{NN}}$:

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.

Learning Configuration 2: Root Axioms We train $\mathcal{F}^{\text{NN-root}}$ on a second set of axioms; those which, if satisfied, imply satisfaction of the other axioms. These are: Majority Loser, Condorcet Winner, Dummett's Condition, Local Stability, Strong Pareto, and the Core. See (Caiata, Armstrong, and Larson 2025) for details of these relationships. We train $\mathcal{F}^{\text{NN-root}}$ on all distributions below, except as noted.

[1]Early experiments sampled n from a normal distribution truncated between 25 and 75. We found n has no impact on performance and use profiles with 50 voters for simplicity.

Learning Parameters For each configuration we train 20 neural networks using PyTorch (Paszke et al. 2019) with 5 hidden layers of 256 nodes. Networks are trained for up to 50 epochs using the Adam optimizer with a learning rate of 1e-4, stopping if there is no improvement of at least 0.0005 over 10 epochs. We generate separate train and test sets of 25,000 examples each for all 255 unique combinations of m, k, and $\mathcal{D}$. Each profile contains 50 voters. 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 (except with the Identity distribution).

Preference Distributions We parameterize each of our preference distributions as follows.

- IC, IAC, Identity take no parameter
- Mallow's with θ sampled uniformly at random as described by Boehmer, Faliszewski, and Kraiczy (2023).
- Urn with α sampled from a Gamma distribution (Boehmer et al. 2021).
- Single-peaked distributions as described by (Conitzer 2007) and (Walsh 2015).
- 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. Due to the similarity of results across Euclidean distributions for $\mathcal{F}^{\text{NN-all}}$ we train $\mathcal{F}^{\text{NN-root}}$ only on 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 $\mathcal{F}^{\text{NN-root}}$.

4.2 Understanding Our Cornerstones

We first report our findings by considering each pair of our cornerstones: axioms, voting rules, and preferences. By examining multiple aspects of our area of focus we can form novel connections.

Axioms and Voting Rules In Table 1 we show the AVR of each rule on each individual axiom. Upon dividing rules and axioms into their informal categories of excellence-based or diverse/proportional some consistent patterns emerge.

Most notably, excellence-based rules generally have low axiom violation rates, even on axioms relating to diversity/proportionality. *Axioms that are formally satisfied by some proportional rules (i.e., JR, EJR) are almost never violated by most excellence-based rules.* Other proportionality-based axioms, such as Solid Coalitions, are violated *more frequently* by all diverse/proportional rules except $\mathcal{F}^{\text{STV}}$. The following may explain this finding: Alternatives which are liked individually are 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.

Axioms and Preference Distributions In considering the relationship between axioms and preferences we also identify trends. Figure 1 show that distributions with moderate or high structure have highly variable AVR (see also, (Caiata,

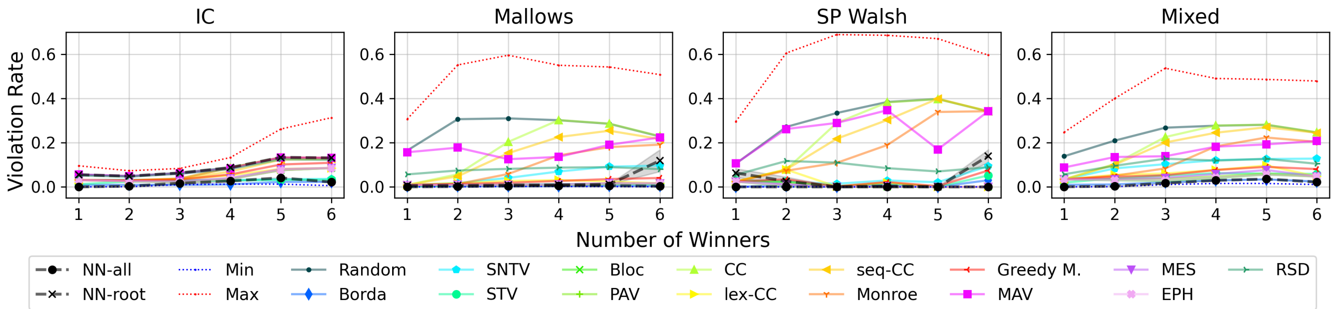

Figure 1: AVR on all axioms for $m = 7$ on distributions with low, moderate, and high structure, and mix of all distributions.

Armstrong, and Larson 2025)). Surprisingly, we observe differences even between Single-Peaked distributions. We also observe a spike in maximum AVR when selecting a committee that contains less than half of alternatives. This may be an artifact of specific axioms: e.g., on Identity preferences it is impossible to violate Justified Representation when choosing a committee with a majority of alternatives.

Voting Rules and Preference Distributions We now focus primarily on voting rules and preferences. Figure 1 shows the axiom violation rate averaged over all axioms for each different voting rule on selected distributions. As we would expect, highly structured preference distributions (Identity, Single-Peaked, Stratified) have correspondingly higher axiom violation rates by most rules. This is unsurprising as additional structure provides more opportunity to violate axioms, as shown by $\mathcal{F}^{\text{Max}}$ which is much higher for these distributions. Similarly, in distributions with low structure (IC, IAC) 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.

Plots shared in our supplementary material enable fine-grained observations. e.g., $\mathcal{F}^{\text{CC}}$, $\mathcal{F}^{\text{seq-CC}}$, and $\mathcal{F}^{\text{SNTV}}$ have unexpectedly high AVR on Identity preferences (Caiata, Armstrong, and Larson 2025). In fact, this is an artifact of lexicographic tie-breaking. Randomized tie-breaking could affect these results and is interesting future work.

Voting Rule Differences Table 2 shows the mean difference between each rule based on the overlap between their elected committees. As expected, randomly chosen committees and $\mathcal{F}^{\text{Max}}$ are almost always the most different than committees returned by all other rules. Surprisingly, there are exceptions: $\mathcal{F}^{\text{MAV}}$ and $\mathcal{F}^{\text{seq-CC}}$ elect committees with *less* overlap than with a randomly chosen committee. This indicates that these rules optimize for different, and mutually exclusive, goals. Other pairs of rules more different from each other than from random committees can be found when considering only a single preference distribution (included in our supplemental files). Note that all the differences from random committees listed in Table 2 appear identical. While these values are close to each other, the identicalness is an artifact of rounding.

On the other hand, we also find rules that are similar despite, on their surface, following different algorithms: $\mathcal{F}^{\text{Bloc}}$, $\mathcal{F}^{\text{PAV}}$ and $\mathcal{F}^{\text{EPH}}$ all follow some procedure that awards an equal number of points to a set of alternatives. Each of these rules have very low distance from the others, indicating that moderate differences in their algorithmic behaviour have relatively little effect on the committees they elect. $\mathcal{F}^{\text{NN-all}}$ is most similar to $\mathcal{F}^{\text{Min}}$. This is expected given that $\mathcal{F}^{\text{NN-all}}$ is trained to find exactly the committees that minimize the number of axiom violations. It is also interesting that $\mathcal{F}^{\text{Borda}}$ is relatively close to $\mathcal{F}^{\text{Min}}$. This, too, is expected given that $\mathcal{F}^{\text{Borda}}$ has among the lowest AVR. Overall, there is a trend of rules with low AVR electing committees similar to one another, while rules with high AVR have more variation.

4.3 Learning From Learned Rules

Table 1 shows that both axiom configurations result in rules with generally low AVR. Averaged across all axioms, $\mathcal{F}^{\text{NN-all}}$ has a lower AVR than any existing rule we evaluated while $\mathcal{F}^{\text{NN-root}}$ has a lower AVR than all rules but $\mathcal{F}^{\text{Borda}}$. While these results are still higher than the lowest possible violation rate ($\mathcal{F}^{\text{Min}}$), they show that networks are capable of identifying committees with desired axiomatic properties.

By learning different sets of axioms we both deepen our understanding of which axioms are more difficult to learn, and gauge the benefit of learning "redundant" axioms with $\mathcal{F}^{\text{NN-all}}$. As $\mathcal{F}^{\text{NN-all}}$ has a 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 a mean AVR of 0.0283 compared to an average violation rate of 0.0692 for $\mathcal{F}^{\text{NN-root}}$. Training on the additional axioms provides additional signal to the learning process which reduces violations of root axioms.

4.4 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 (Mattei and Walsh 2013). 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,

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 (Elkind et al. 2017; Lackner and Skowron 2023). Due to differences in tie-breaking with previous work, some edge cases do not match prior theoretical results.

for example, a mean AVR of 0.102 compared to an AVR of 0.303 for committees selected at random.

4.5 Optimized Positional Scoring Rules

The low AVR of $\mathcal{F}^{\text{Borda}}$ inspires a natural question: Is there a positional scoring rule with even stronger performance than $\mathcal{F}^{\text{Borda}}$? A positive answer would provide a rule which is both highly optimized and interpretable in a way that neural networks are not. To address this question we use the `optimal-voting` package (Armstrong 2025) to generate improved positional score vectors. After only a small number of annealing steps we reliably generate score vectors that improve upon $\mathcal{F}^{\text{Borda}}$ by a small amount. Though mild, these results suggest further optimization in this regard may yield interesting results. Full results of this experiment are included in (Caiata, Armstrong, and Larson 2025).

5 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 underlying voter preferences, and (3) it is possible to learn novel rules with significantly lower axiom violation rates. These contributions complement other recent work demonstrating the efficacy of combining machine learning with social choice (Golowich, Narasimhan, and Parkes 2018; Conitzer et al. 2024; Lanctot et al. 2025) and adds to the literature extending axiomatic analysis beyond the worst case (Flanigan, Halpern, and Psomas 2023). We highlight findings of particular note to the broader research community as well as 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) (Lackner and Skowron 2023) while Local Stability, Dummett's, and Solid Coalitions have distinct origins. Most rules violate the first three axioms similarly, while there is much more variability in AVR on axioms from the second group. Clear distinctions between these two types of proportionality are missing in the literature. For example, $\mathcal{F}^{\text{CC}}$ and $\mathcal{F}^{\text{Monroe}}$ "explicitly aim at proportional representation" (Elkind et al. 2017) 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 proportional rules $\mathcal{F}^{\text{CC}}$ and $\mathcal{F}^{\text{Monroe}}$ for the Condorcet/Majority Loser axioms demonstrates a fundamental tension between

	Borda	EPH	SNTV	Bloc	STV	PAV	MES	CC	seq-CC	lex-CC	Monroe	Greedy M.	MAV	RSD
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.

proportionality and avoiding bad alternatives. To be proportional, it may be necessary to elect an alternative ranked low by a majority of voters. New axioms may bridge this gap by formally considering a balance between proportionality and electing highly polarizing alternatives.

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

Learning Simple Axioms Targeting both simple and complex axioms results in $\mathcal{F}^{\text{NN-all}}$ having lower AVR than $\mathcal{F}^{\text{NN-root}}$. This suggests that future axiom learning tasks could be enhanced by introducing simple dummy axioms that inform about more complex axioms (e.g., a Fixed Majority committee is a Condorcet Winner, so the FM axiom may be seen as informative about Condorcet Winner).

Difference Between Rules Our results can be compared with maps of multi-winner elections (Faliszewski et al. 2023). In particular, in both their work and ours: (1) rules focused on individual excellence elect different committees than proportional rules, and (2) $\mathcal{F}^{\text{MAV}}$ is quite unique.

5.1 Future Directions

There are several interesting directions for this work. First, our research can be directly applied to evaluate arbitrary sets of intraprofile axioms, and extended to the more general class of *interprofile* axioms (Schmidtlein 2022). This would allow measuring violations of additional axioms such as consistency or clone-proofness (Brandt et al. 2016).

Second, we proposed a new metric for evaluating voting rules, the axiom violation rate. 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 a low axiom violation rate. Black-box machine learning techniques may suffice for this; however, identifying classes of rule (such as positional scoring rules) which are optimizable and provide reliable, interpretable outcomes 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 given by PAC-learning (Valiant 1984). Such an approach would complement both the axiomatic and the empirical study of voting rules.

Acknowledgements

Ben Armstrong was partially supported by NSF Award IIS-III-2107505 and the Tulane University Center of Excellence for Community-Engaged Artificial Intelligence (CEAI).

References

Anil, C.; and Bao, X. 2021. Learning to Elect. *Advances in Neural Information Processing Systems*, 34: 8006–8017.

Armstrong, B. 2025. Optimal Voting.

Armstrong, B.; and Larson, K. 2019. Machine Learning to Strengthen Democracy. In *NeurIPS Joint Workshop on AI for Social Good*.

Aziz, H.; Elkind, E.; Faliszewski, P.; Lackner, M.; and Skowron, P. 2017. The Condorcet Principle for Multiwinner Elections: From Shortlisting to Proportionality. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*, 84–90.

Boehmer, N.; Bredereck, R.; Faliszewski, P.; Niedermeier, R.; and Szufa, S. 2021. Putting a Compass on the Map of Elections. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, 59–64.

Boehmer, N.; Faliszewski, P.; and Kraiczy, S. 2023. Properties of the Mallows Model Depending on the Number of Alternatives: A Warning for an Experimentalist. In *International Conference on Machine Learning*, 2689–2711. PMLR.

Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D. 2016. *Handbook of Computational Social Choice*. Cambridge University Press.

Burka, D.; Puppe, C.; Szepesváry, L.; and Tasnádi, A. 2022. Voting: A Machine Learning Approach. *European Journal of Operational Research*, 299(3): 1003–1017.

Caiata, J.; Armstrong, B.; and Larson, K. 2025. What Voting Rules Actually Do: A Data-Driven Analysis of Multi-Winner Voting. *arXiv preprint arXiv:2508.06454*.

Conitzer, V. 2007. Eliciting Single-Peaked Preferences Using Comparison Queries. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems*, 1–8.

Conitzer, V.; Freedman, R.; Heitzig, J.; Holliday, W. H.; Jacobs, B. M.; Lambert, N.; Mossé, M.; Pacuit, E.; Russell, S.; Schoelkopf, H.; et al. 2024. Position: Social Choice Should Guide AI Alignment in Dealing with Diverse Human Feedback. In *Proceedings of the 41st International Conference on Machine Learning*, 9346–9360.

Debord, B. 1993. Prudent k-Choice Functions: Properties and Algorithms. *Mathematical Social Sciences*, 26(1): 63–77.

d'Eon, G.; and Larson, K. 2020. Testing Axioms Against Human Reward Divisions in Cooperative Games. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems*, 312–320.

Dummett, M. 1984. *Voting Procedures*. Oxford University Press.

Eggenberger, F.; and Pólya, G. 1923. Über die Statistik Verketteter Vorgänge. *ZAMM-Journal of Applied Mathematics and Mechanics/Zeitschrift für Angewandte Mathematik und Mechanik*, 3(4): 279–289.

Elkind, E. 2018. 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*, 12–18. Springer.

Elkind, E.; Faliszewski, P.; Skowron, P.; and Slinko, A. 2017. Properties of Multiwinner Voting Rules. *Social Choice and Welfare*, 48: 599–632.

Enelow, J. M.; and Hinich, M. J. 1984. *The Spatial Theory of Voting: An Introduction*. CUP Archive.

Fairstein, R.; Vilenchik, D.; and Gal, K. 2024. Learning Aggregation Rules in Participatory Budgeting: A Data-Driven Approach. *arXiv preprint arXiv:2412.01864*.

Faliszewski, P.; Lackner, M.; Sornat, K.; and Szufa, S. 2023. An Experimental Comparison of Multiwinner Voting Rules on Approval Elections. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, 2675–2683.

Faliszewski, P.; Skowron, P.; Slinko, A.; and Talmon, N. 2017. Multiwinner Voting: A New Challenge for Social Choice Theory. *Trends in Computational Social Choice*, 74(2017): 27–47.

Fishburn, P. C. 1977. Condorcet Social Choice Functions. *SIAM Journal on Applied Mathematics*, 33(3): 469–489.

Flanigan, B.; Halpern, D.; and Psomas, A. 2023. Smoothed Analysis of Social Choice Revisited. In *International Conference on Web and Internet Economics*, 290–309. Springer.

Gehrlein, W. V. 1985. The Condorcet Criterion and Committee Selection. *Mathematical Social Sciences*, 10(3): 199–209.

Gehrlein, W. V.; and Fishburn, P. C. 1976. Condorcet's Paradox and Anonymous Preference Profiles. *Public Choice*, 1–18.

Golowich, N.; Narasimhan, H.; and Parkes, D. C. 2018. Deep Learning for Multi-Facility Location Mechanism Design. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*, 261–267.

Guilbaud, G. T. 1952. Les Théories de l'intérêt général et le Problème Logique de l'agrégation. *Economie appliquée*, 5(4): 501–584.

Holliday, W. H.; Kristoffersen, A.; and Pacuit, E. 2025. Learning to Manipulate Under Limited Information. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 13915–13925.

Hornischer, L.; and Terzopoulou, Z. 2025. Learning How to Vote with Principles: Axiomatic Insights into the Collective Decisions of Neural Networks. *Journal of Artificial Intelligence Research*, 83.

Kuga, K.; and Nagatani, H. 1974. Voter Antagonism and the Paradox of Voting. *Econometrica: Journal of the Econometric Society*, 1045–1067.

Kujawska, H.; Slavkovik, M.; and Rückmann, J.-J. 2020. 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*, 440–458. Springer.

Lackner, M.; and Skowron, P. 2023. *Multi-Winner Voting with Approval Preferences*. Springer International Publishing. ISBN 9783031090165.

Lanctot, M.; Larson, K.; Kaisers, M.; Berthet, Q.; Gemp, I.; Diaz, M.; Maura-Rivero, R.-R.; Bachrach, Y.; Koop, A.; and

Precup, D. 2025. Soft Condorcet Optimization. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025)*.

Mallows, C. L. 1957. Non-Null Ranking Models. I. *Biometrika*, 44(1/2): 114–130.

Matone, L.; Abramowitz, B.; Armstrong, B.; Mattei, N.; and Balakrishnan, A. 2025. DeepVoting: Learning Voting Rules with Tailored Embeddings. *arXiv preprint arXiv:2408.13630*.

Mattei, N.; and Walsh, T. 2013. Preflib: A Library for Preferences. In *Algorithmic Decision Theory: Third International Conference, ADT 2013, Bruxelles, Belgium, November 12-14, 2013, Proceedings 3*, 259–270. Springer.

Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An Imperative Style, High-Performance Deep Learning Library. *Advances in Neural Information Processing Systems*, 32.

Peters, D.; and Skowron, P. 2020. Proportionality and the Limits of Welfarism. In *Proceedings of the 21st ACM Conference on Economics and Computation*, 793–794.

Schmidtlein, M. C. 2022. Voting by Axioms. *Master's thesis. ILLC, University of Amsterdam*.

Valiant, L. G. 1984. A Theory of the Learnable. *Communications of the ACM*, 27(11): 1134–1142.

Walsh, T. 2015. Generating Single Peaked Votes. *arXiv preprint arXiv:1503.02766*.