The Thirty-Ninth AAAI Conference on Artificial Intelligence (AAAI-25)

Proportional Representation in Practice: Quantifying Proportionality in Ordinal Elections

Tuva Bardal[1], **Markus Brill**[1], **David McCune**[2], **Jannik Peters**[3]

[1]University of Warwick, Coventry, UK
[2] William Jewell College, Liberty, MO, United States
[3] National University of Singapore, Singapore
tuva.bardal@warwick.ac.uk, markus.brill@warwick.ac.uk, mccuned@william.jewell.edu, peters@nus.edu.sg

Abstract

Proportional representation plays a crucial role in electoral systems. In ordinal elections, where voters rank candidates based on their preferences, the *Single Transferable Vote (STV)* is the most widely used proportional voting method. STV is considered proportional because it satisfies an axiom requiring that large enough "solid coalitions" of voters are adequately represented. Using real-world data from local Scottish elections, we observe that solid coalitions of the required size rarely occur in practice. This observation challenges the importance of proportionality axioms and raises the question of how the proportionality of voting methods can be assessed beyond their axiomatic performance. We address these concerns by developing quantitative measures of proportionality. We apply these measures to evaluate the proportionality of voting rules on real-world election data. Besides STV, we consider *SNTV*, the *Expanding Approvals Rule*, and *Sequential Ranked-Choice Voting*. We also study the effects of ballot truncation by artificially completing truncated ballots and comparing the proportionality of outcomes under complete and truncated ballots.

Datasets — https://github.com/mggg/scot-elex

1 Introduction

Proportional representation is a core principle of modern electoral systems, its goal being to ensure that elected officials proportionally represent the composition of the electorate. Most countries achieve proportional representation through party-list systems, where voters cast their ballots for political parties rather than individual candidates, and legislative seats are allocated in proportion to the number of votes each party receives. However, several countries—mostly from the former British commonwealth—use a system in which voters can vote for individual candidates using ranked preferences. In such settings, the electorate is usually divided into a number of districts, each of which elects a few (usually 2–6) candidates to represent the district in the given governing body. To achieve proportional representation in district elections which do not use a party-list system, most localities use the *Single Transferable Vote (STV)*, a voting rule dating back to the 19th century (Tideman 1995). STV is generally considered to guarantee proportional representation since it satisfies a property known as *proportionality for solid coalitions (PSC)* formulated by Dummett (1984). This property, in essence, guarantees that any sufficiently large group of voters that rank the same candidates (not necessarily in the same order) above all other candidates is given an amount of representation commensurate with its size. Such groups are referred to as *solid coalitions*.

Despite the prominence of PSC in the literature, some have challenged the idea that PSC is a good way to capture the notion of proportional representation rigorously. For instance, the requirement that all voters in a solid coalition must all share the exact same candidate set in their ballot prefix makes PSC highly non-robust (Tideman 2006). This issue has been considered from a purely theoretical point of view (Aziz and Lee 2020; Brill and Peters 2023), as well as through experiments on synthetic data (Brill and Peters 2023). There exist several other properties that seek to guarantee proportional representation in ways similar to that of PSC (Aziz et al. 2017; Brill and Peters 2023), and that, from the aforementioned perspective, are more robust (Brill and Peters 2023). However, such properties are purely qualitative, and like PSC they usually rely rigidly on a "threshold of representation," making them blind to groups of voters that come close to being "sufficiently large."

While proportionality axioms have received much attention in the theoretical literature, empirical investigation of the force of axioms formulated to guarantee proportional representation has been limited due to a scarcity of real-world ballot data. However, McCune and Graham-Squire (2024) recently compiled a real-world dataset consisting of 1100 Scottish local council elections from the period 2007–2022. In these elections, STV is used to elect members for local councils in Scotland. We observe that in this data, solid coalitions of sufficient size rarely occur. Consequently, for most elections in the dataset, PSC places no or few restrictions on the winning committee, and most outcomes are therefore proportional according to the property. Since PSC often undergirds the claim that STV is proportional, the toothlessness of PSC in practice raises the question of how proportional the method really is.[1] More generally, we con-

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

[1]As demonstrated by Brill and Peters (2023), STV fails proportionality axioms that are stronger than PSC.

sider the issue of how to assess the proportionality of a committee when axioms have little to no effect on the outcome. As a first step towards answering this question, we define quantitative versions of several proportionality axioms suggested in the literature, based on the idea that we should be able to relax size-requirements imposed on groups of voters whenever there are few or no groups of sufficient size. This approach not only allows us to measure the extent to which an axiom is satisfied, but also makes it possible to strengthen axioms in situations where they have little to no effect. We furthermore use the data from the Scottish local council elections to assess experimentally the proportionality of different voting rules according to our measures.

In practice, voters tend not to form cohesive groups large enough for standard proportionality axioms to place any significant requirement on the winning committee. The issues of ballot truncation and small cohesive groups have been discussed previously (Brill and Peters 2023; Hoffman et al. 2024; Marsh and Plescia 2016), but such work did not have access to actual ballot data. Our contribution is that we investigate how to adapt proportionality axioms to such a real-world setting, and we provide empirical results from a large dataset of real-world STV elections.

Related Work The formal study of proportional representation in multiwinner voting was initiated by Dummett (1984). In recent years, proportional representation has been mainly studied in *approval-based* multiwinner voting (Lackner and Skowron 2022). Proportionality for ordinal elections has been studied by Aziz and Lee (2020, 2021, 2022), Brill and Peters (2023), and Delemazure and Peters (2024). Varying the size of voter groups that deserve representation has been considered by Janson (2018), who focuses on worst-case bounds, and Jiang, Munagala, and Wang (2020), who focus on relaxations of axioms. By contrast, in this paper, we focus on strengthening axioms and on bounds that are specific to an instance (rather than worst-case over all instances). Finally, ballot truncation has been considered by Hoffman et al. (2024) for multiwinner elections, but is more commonly studied in the context of single-winner voting (Kilgour, Grégoire, and Foley 2020; Tomlinson, Ugander, and Kleinberg 2023; Burnett and Kogan 2015; Graham-Squire and McCune 2023).

2 Preliminaries

For $t \in \mathbb{N}$, let $[t] = \{1, \ldots, t\}$. Throughout the paper, we assume that we are given a set $N = [n]$ of *voters* and a set $C = \{c_1, \ldots, c_m\}$ of *candidates*. The preferences of voters are *top-truncated*. That is, each voter $i \in N$ chooses a set $A_i \subseteq C$ of candidates and a *strict* ranking $\succ_i : A_i \times A_i$ over these candidates, where $|A_i|$ may be less than m. For notational purposes, we assume that $c \succ_i c'$ for any $c \in A_i$ and $c' \notin A_i$, while voters are indifferent among candidates they do not rank. Finally, let k denote the number of candidates that need to be selected. We refer to subset $W \subseteq C$ of candidates of size $|W| = k$ as a *committee*. A *(multiwinner voting) instance* I is a collection of voters, candidates, voter preferences, and the committee size.

2.1 Proportionality Axioms

For $\ell \in [k]$, we say that a group $N' \subseteq N$ of voters is *ℓ-large* if $|N'| \geq \ell \frac{n}{k}$. Generally speaking, in proportional multiwinner voting if N' is ℓ-large and N' is "cohesive" in some sense, then ℓ of the candidates supported by N' should receive seats on the winning committee.

The most prominent proportionality axiom for ranked preferences is *proportionality for solid coalitions (PSC)* introduced by Dummett (1984). Given a subset $N' \subseteq N$ of voters and $C' \subseteq C$ of candidates, N' forms a *solid coalition* over C' if for any pair of candidates $c_j \in C'$ and $c_r \in C \setminus C'$, it holds that $c_j \succ_i c_r$ for all $i \in N'$. In other words, C' forms a prefix of the ranking $\succ_i$ of every voter in N'. Using the notion of solid coalitions, we can now define PSC.[2]

Definition 1 (PSC)**.** *A committee W satisfies* proportionality for solid coalitions (PSC) *if for any $\ell \in [k]$ and any ℓ-large group $N' \subseteq N$ of voters forming a solid coalition over $C' \subseteq C$, it holds that $|C' \cap W| \geq \min(|C'|, \ell)$.*

As a potential alternative to PSC and possible generalization of the Condorcet principle to proportional representation, Aziz et al. (2017) introduced the concept of *local stability*. Intuitively, local stability postulates that no group of voters of size at least $\frac{n}{k}$ should find an unselected candidate they all prefer to everyone in the committee.[3] Notably, committees satisfying local stability need not exist.

Definition 2 (LS)**.** *A committee W satisfies* local stability (LS) *if there is no 1-large group of voters $N' \subseteq N$ and candidate $c \notin W$ with $c \succ_i c'$ for all $i \in N'$ and $c' \in W$.*

We also consider two proportionality axioms that have been proposed in the setting of approval-based multiwinner voting: *EJR+* (Brill and Peters 2023) and *priceability* (Peters and Skowron 2020). These can be translated to the ordinal setting using a construction due to Brill and Peters (2023).[4] The definition of priceability involves a system of inequalities and a *price* $p \in \mathbb{R}^+$; we call a committee *priceable* if there is a solution to this system with $p < \frac{n}{k}$.

2.2 Voting Rules

A voting rule maps each instance to one or more *winning committees*. We briefly introduce the voting rules we study. For more details, we refer to the full version of this paper.

The *Single Transferable Vote (STV)* is a family of rules, with different versions of STV used in different jurisdictions (Tideman 1995). Following McCune and Graham-Squire (2024), we describe the version that is used in Scottish local elections. STV proceeds in rounds and starts by assigning each voter $i \in N$ a weight $w_i = 1$. In each round, STV checks whether the candidate with the most weighted first-place votes has at least q (weighted) first-place votes, where

[2]A related axiom is *generalized PSC* (Aziz and Lee 2020), which generalizes PSC to the case of weak orders. For top-truncated preferences, PSC and generalized PSC are equivalent.

[3]The same concept was independently studied by Jiang, Munagala, and Wang (2020) in the more general context of core stability.

[4]Since we are only dealing with ordinal preferences in this paper, we omit the "rank" prefix used by Brill and Peters (2023). We do not consider the axiom *(rank-)PJR+* for computational reasons.

$q = \lfloor \frac{n}{k+1} \rfloor + 1$ is the *quota*. If that is the case, this candidate is elected and the voters ranking this candidate first are reweighted proportionally such that the total weight of the election decreases by exactly q. If there is no such candidate, the candidate with the least weighted first-place votes is removed. Besides this version of STV, which we refer to as as *Scottish STV*, we also consider *Meek-STV* (Hill, Wichmann, and Woodall 1987).

The *Expanding Approvals Rule (EAR)* is another family of proportional multiwinner voting rules (Aziz and Lee 2020). Just like STV, it uses a quota q and starts by assigning each voter a weight $w_i = 1$. It then iterates over all possible ranks from first to last. For each such rank and so-far unselected candidate c it checks whether the total weight of the voters giving a rank of at most r to c is at least q. If there is such a candidate, it takes any such candidate into the committee and decreases the collective weight of the corresponding voters by q. In our implementation of the rule we select the candidate with the largest total weight and decrease the weights as in Scottish STV. If there is no such candidate, r is increased.[5]

The *Single Non-Transferable Vote (SNTV)* selects the k candidates with the highest first-place vote count. SNTV is sometimes referred to as k*-plurality*.

Sequential Ranked-Choice Voting (seq-RCV), which is currently used in the US state of Utah (McCune et al. 2024), executes the single-winner RCV procedure k times. Single-winner RCV (a.k.a. *instant runoff voting*) iteratively deletes the candidate with the fewest first-place votes until only a single candidate is left.

We include seq-RCV in the rules we examine because it is not proportional but *majoritarian* (i.e., a group consisting of barely more than half of the electorate can force the entire committee to consist of their candidates) and therefore can provide context for our results around methods like STV. We note that across all 1070 multiwinner Scottish elections in our dataset there are only nine in which seq-RCV chooses a winner set incompatible with PSC. Thus, if PSC is the standard by which a rule is judged to be proportional, seq-RCV is virtually proportional in practice. However, seq-RCV often produces outcomes which are wildly non-proportional under any intuitive notion of "proportional" (McCune et al. 2024), and this provides additional motivation for why we should explore alternatives to standard PSC in practice.

While SNTV does not satisfy any of the axioms discussed in Section 2.1, it is considered a "semi-proportional" method (Amy 2000). Such methods aim to give some representation to minorities, albeit not proportional to their support.

The remaining three voting rules are proportional: Both versions of STV satisfy PSC, and EAR satisfies an even stronger axiom called *rank-PJR+*. The latter axiom was introduced by Brill and Peters (2023), who also showed that EAR winning committees are always priceable.

[5]The treatment of unranked candidates in EAR allows for different interpretations, as noted in Remark 3 by Aziz and Lee (2020). When a voter ranks a strict subset A_i of candidates, the set $C \setminus A_i$ of unranked candidates (i.e., the last equivalence class) can be included either *(i)* as soon as the ranked candidates are exhausted, or *(ii)* only in the final step of the method. We tested both variants in our experimental analysis. The performance differences between the two versions were minor, with variant *(ii)* performing slightly better. Therefore, we focus on variant *(ii)* here.

	$< 25\%$	$< 50\%$	$< 100\%$
PSC	42 (3.9%)	275 (25.7%)	776 (72.5%)
EJR+	64 (6.0%)	357 (33.4%)	1005 (93.9%)
LS	173 (16.2%)	650 (60.7%)	1061 (99.1%)
Priceability	61 (5.7%)	354 (33.1%)	1003 (93.7%)

Table 1: For each axiom, the row values correspond to the number of elections in the dataset where the satisfaction rate is strictly less than the percentage value at the top of the column. For instance, in only 42 out of 1070 elections it is the case that less than 25% of all committees satisfy PSC.

3 Scottish Local Council Elections

For local governance, Scotland is partitioned into 32 "council areas," each of which is governed by a council. In turn, each council area is divided into wards, each of which elects a number of councilors to represent the ward on the council. The number of candidates running and the number of seats available in a typical election are not large; most elections satisfy $m \in \{6, 7, 8, 9\}$ and $k \in \{3, 4\}$. Since 2007, all wards have used Scottish STV to choose their representatives. Elections are held every five years.

McCune and Graham-Squire (2024) collected ballot data from 1100 Scottish local council elections between the years 2007 and 2022. Out of the 1100 elections, 1070 satisfy $k > 1$; we only consider these 1070 instances.

Notably, voters are not required to provide full rankings and ballots are often heavily truncated: across the 5,485,379 total ballots cast from all elections, approximately 14% rank only a single candidate, and a majority, 58%, rank fewer than k. In contrast, only 13% of ballots are complete (where by "complete" we mean a ballot that contains $m-1$ or m candidates). We refer to McCune and Graham-Squire (2024) for more details about the dataset.

We evaluated the force of proportionality axioms on the ballot data from the elections by calculating the number of outcomes excluded by the axioms on each instance. Out of 1070 elections, there are 294 (27.5%) for which every committee of size k satisfies PSC, 592 (55.3%) where there is only one solid coalition which earns a single seat under PSC, and 184 (17.2%) where multiple solid coalitions are deserving of seats by PSC. Thus, in real-world elections PSC does not place significant restrictions on the winning committee. The other axioms we consider are generally more discerning than PSC, however none of the axioms identify a unique outcome on any of the elections we consider. We give an example to illustrate both how PSC may fall short in excluding outcomes and how the number of compatible committees may differ between the axioms. An overview of the number of outcomes that satisfy each of the axioms considered over all elections in the dataset can be found in Table 1.

Example 1. *Consider the 2012 council election of Midlothian, ward 2, with* $n = 5132$ *voters,* $m = 7$ *candidates, and*

$k = 3$ seats. The candidates, their party affiliations, and their first-place vote counts are listed in Table 2.

Candidate	Party	First-Place Votes
D. Milligan (**DM**)	Labour	1,574
L. Milliken (**LM**)	Labour	525
J. Aitchison (**JA**)	Independent	382
B. Constable (**BC**)	SNP	1,257
T. Munro (**TM**)	SNP	358
I. Baxter (**IB**)	Greens	671
E. Cummings (**EC**)	Conservative	365

Table 2: Candidates and vote counts for Example 1.

*There are $\binom{m}{k} = \binom{7}{3} = 35$ possible outcomes in this election. The largest (non-trivial) solid coalition is over the two Labour candidates **DM** and **LM** and has size 1624. This coalition consists of the 1218 voters who cast a ballot of the form **DM** $\succ$ **LM** $\succ \dots$ and of the 406 voters who cast a ballot of the form **LM** $\succ$ **DM** $\succ \dots$. Interestingly, the size of this coalition is much smaller than the total number of voters who ranked a Labour candidate first. The reasons are that some Labour voters rank only one candidate on their ballots and many voters cast split-ticket ballots. The next largest solid coalition has size 1277 and is over the two SNP candidates **BC** and **TM**, again barely more than the first-place votes of **BC**. The largest solid coalition over more than two candidates consists of 554 voters who support the two SNP candidates as well as **IB**. Solid coalitions over four or more candidates are extremely small. The threshold for a coalition to be 1-large is $\lceil \frac{n}{k} \rceil = 1711$, and therefore any of the 35 possible winning committees satisfies PSC. In comparison, 24 out of 35 committees satisfy rank-EJR+ and priceability, while 12 out of 35 outcomes are locally stable.*

4 Quantifying Proportionality

In this section, we turn proportionality axioms into quantitative proportionality measures. The main idea behind the construction of measures consists in (1) defining a parameterized version of the proportionality axiom by introducing a multiplicative factor on the size constraint, and (2) identifying the smallest parameter for which the parameterized axiom is satisfied by the given committee. We start with PSC and consider other axioms in Section 4.4.

4.1 Quantifying PSC

Recall that a group $N' \subseteq N$ of voters is called ℓ-large if $|N'| \geq \ell \frac{n}{k}$. A parameterized version of this notion can be obtained by introducing a multiplicative factor.

Definition 3. *Consider an instance with n voters and committee size k. For $\alpha \in \mathbb{R}^+$ and $\ell \in [k]$, a group $N' \subseteq N$ of voters is ℓ_α-*large *if $|N'| \geq \alpha \cdot \ell \frac{n}{k}$.*

That is, the value of the parameter α changes the size requirement a group needs to fulfil in order to be deemed worthy of ℓ representatives. If $\alpha < 1$, then the size constraint is relaxed, as groups of size smaller than $\ell \frac{n}{k}$ deserve ℓ representatives. On the other hand, if $\alpha > 1$, then a group of voters must have larger size to deserve representation. PSC requires that ℓ-large groups need to be represented appropriately; consequently, replacing "ℓ-large" with "ℓ_α-large" in the definition of PSC makes the axiom more demanding for $\alpha < 1$ and less demanding for $\alpha > 1$.

Definition 4 (α-PSC)**.** *Let $\alpha \in \mathbb{R}^+$. A committee W satisfies* α-PSC *if for any $\ell \in [k]$ and any ℓ_α-large subset $N' \subseteq N$ of voters forming a solid coalition over $C' \subseteq C$, it holds that $|C' \cap W| \geq \min(|C'|, \ell)$.*

Clearly, 1-PSC is equivalent to (original) PSC and lowering the value of α makes the axiom more demanding: If $\alpha_1 \leq \alpha_2$, then α_1-PSC implies α_2-PSC. For a given committee W, we are therefore interested in the smallest value of α such that W satisfies α-PSC. Formally, let[6]

$$\alpha_{\text{PSC}}(W) = \inf\{\alpha \colon W \text{ satisfies } \alpha\text{-PSC}\}.$$

We refer to $\alpha_{\text{PSC}}(W)$ as the *PSC value of W*. Observe that W satisfies PSC if and only if $\alpha_{\text{PSC}}(W) < 1$.

Furthermore, we call the minimum achievable α-value for an instance I the *PSC value of I* and denote it with

$$\alpha^*_{\text{PSC}}(I) = \min_{W \subseteq C \colon |W| = k} \alpha_{\text{PSC}}(W).$$

Choosing a winning committee which achieves $\alpha^*_{\text{PSC}}(I)$ might be normatively desirable in a proportional context because such a committee fulfills the spirit of PSC in the absence of large solid coalitions.

Example 2. *Consider again the instance from Example 1. Here, 1624 voters form a solid coalition over $\{$**DM**, **LM**$\}$, followed by solid coalitions of size 1574, 1277, 1257, 671, over $\{$**DM**$\}$, $\{$**BC**, **TM**$\}$, $\{$**BC**$\}$, and $\{$**IB**$\}$, respectively. The α-values at which the solid coalitions become ℓ_α-large for $\ell \in \{1, 2\}$ are given in the table below.*

ℓ	{DM, LM}	{DM}	{BC, TM}	{BC}	{IB}
1	0.949	0.920	0.746	0.734	0.392
2	0.474	–	0.373	–	–

Table 3: Solid coalitions and α-values in the election from Example 1.

*Consider the committees $W = \{$**DM**, **LM**, **BC**$\}$ and $W' = \{$**DM**, **BC**, **IB**$\}$. Committee W is chosen by Meek-STV and EAR, while W' is chosen by Scottish STV. Both committees satisfy PSC, however we can distinguish the committees based on their PSC values: $\alpha_{PSC}(W) = 0.392$ and $\alpha_{PSC}(W') = 0.474$. In fact, W achieves the PSC value for the instance (i.e., $\alpha^*_{PSC}(I) = 0.392$), as any smaller value would additionally force candidate **IB** to be included. From the point of view of solid coalitions, W is perhaps the better choice of winning committee because the solid coalition $\{$**DM**, **LM**$\}$ is more than double the size of any solid coalition containing **IB**.*

[6] We use infimum rather than minimum since the set of values for which α-PSC holds is an open interval of the form $(\alpha_{\text{PSC}}, +\infty)$.

4.2 Computing the PSC Value of a Committee

As already noted, decreasing the value of α leads to more representation demands by solid coalitions. We can identify exactly the values of α that makes a group ℓ_α-deserving:

$$N' \text{ is } \ell_\alpha\text{-large} \quad \Leftrightarrow \quad \alpha \leq \frac{|N'|}{n} \cdot \frac{k}{\ell}.$$

Let $\alpha^\ell_{(N',C')} = \frac{|N'|}{n} \cdot \frac{k}{\ell}$ denote the value of α for which the solid coalition (N', C') becomes ℓ_α-large. Then, the values

$$\alpha^1_{(N',C')},\, \alpha^2_{(N',C')},\, \dots,\, \alpha^{|C'|}_{(N',C')}$$

are exactly the thresholds of α-values for which the group starts to become deserving of $1, 2, \dots, |C'|$ many representatives under α-PSC. (Values $\alpha^\ell_{(N',C')}$ with $\ell > |C'|$ are irrelevant because the group's deservingness is upper bounded by $|C'|$ according to the definition of α-PSC.)

In our algorithm for computing the PSC value of a committee, we compute these values for all solid coalitions. For each subset $C' \subseteq C$, there is a unique *maximal* group $N_{C'}$ of voters that solidly supports C'. The group $N_{C'}$ consists of all voters ranking all candidates in C' over all other candidates. Clearly, it is sufficient to consider only maximal solid coalitions. Let $\mathcal{S}$ denote the set of all maximal solid coalitions. It is not hard to see that $|\mathcal{S}|$ is polynomial in the size of the profile and that we can efficiently enumerate all maximal solid coalitions by iterating over the prefixes of the voters.

Given the set $\mathcal{S}$ of all maximal solid coalitions, we can now collect all threshold values for α. Define T as the set that contains the relevant values for each solid coalition, i.e.,

$$T = \bigcup_{(N',C') \in \mathcal{S}} \{\alpha^1_{(N',C')}, \alpha^2_{(N',C')}, \dots, \alpha^{|C'|}_{(N',C')}\}.$$

Here, each threshold value $\alpha^\ell_{(N',C')}$ is associated with a "PSC constraint" of the form $|W \cap C'| \geq \ell$.

Theorem 1. *Given an instance and a committee W, the PSC value of W can be computed in polynomial time.*

Proof. First calculate the set $\mathcal{S}$ of all maximal solid coalitions and the set T of relevant thresholds. Consider the threshold values in T in non-increasing order. When considering $\alpha^\ell_{(N',C')}$, check whether $|W \cap C'| \geq \ell$. If yes, go to the next threshold value. If not, we know that $\alpha_{\text{PSC}}(W) = \alpha^\ell_{(N',C')}$, because $\alpha^\ell_{(N',C')}$ is the largest value of α for which the corresponding PSC constraint is not satisfied by W. □

4.3 Computing the PSC Value of an Instance

Computing the minimal possible α-value that is achievable in an instance by *any* committee is more challenging. We first show that the problem is NP-hard.[7]

Theorem 2. *Given an instance and a value $\alpha < 1$, deciding whether α-PSC is satisfiable is NP-complete.*

[7] See the full version of this paper for the proof. The problem is trivial for $\alpha \geq 1$, as a committee satisfying PSC always exists.

As a consequence, computing the PSC value of an instance is NP-hard. In order to compute PSC values in our experiments (see Section 5), we employ integer linear programming (ILP). The approach is similar to the one used in Section 4.2: We compute the set T of threshold values and then consider these values in non-increasing order. When considering $\alpha^\ell_{(N',C')}$, we add the constraint $|W \cap C'| \geq \ell$ to our ILP and check whether the resulting ILP is feasible. If yes, we consider the next threshold in T. If not, we have found the PSC value of the instance, as $\alpha^\ell_{(N',C')}$ is the largest value of α for which α-PSC is not satisfiable.

Formally, the ILP has a binary variable $x_c \in \{0, 1\}$ for each candidate $c \in C$ and a constraint $\sum_{c \in C} x_c \leq k$ ensuring that at most k candidates are selected. Constraints of the form $|W \cap C'| \geq \ell$ can be encoded as $\sum_{c \in C'} x_c \geq \ell$.

We remark that this algorithm has similarities to the *D'Hondt apportionment method* (Balinski and Young 1982). In the full version of this paper, we develop a description of this algorithm which gives rise to the idea of "apportionment for non-disjoint parties," which might be of independent interest.

4.4 Quantifying Other Axioms

Generalizing the quantification approach to local stability and EJR+ is straightforward. Similarly to PSC, we can replace each ℓ-large group by an ℓ_α-large group leading to the following two definitions.

Definition 5 (α-LS)**.** *A committee W satisfies* α-local stability (α-LS) *if there is no 1_α-large group $N' \subseteq N$ of voters and $c \notin W$ such that $c \succ_i c'$ for all $i \in N'$ and $c' \in W$.*

While committees satisfying local stability (i.e., 1-LS) do not necessarily exist, Charikar et al. (2024) have recently shown that every instance admits a 9.8217-LS committee.

For the definition of α-EJR+, we let

$$\text{rank}(i, c) = |\{c' \in A_i : c' \succ_i c\}| + 1$$

denote the rank that voter i assigns to candidate $c \in A_i$. For unranked candidates $c \notin A_i$, we let $\text{rank}(i, c) = m$.

Definition 6 (α-EJR+)**.** *A committee W satisfies* α-EJR+ *if there is no $\ell \in [k]$, ℓ_α-large group $N' \subseteq N$ of voters, unselected candidate $c \notin W$, and rank $r \in [m]$ such that*

(i) $\text{rank}(i, c) \leq r$ *for all* $i \in N'$
(ii) $|\{c' \in C\colon \text{rank}(i, c') \leq r\} \cap W| < \ell$ *for all* $i \in N'$.

The definition of priceability is already parameterized, with a price of $p < \frac{n}{k}$ implying PSC (Brill and Peters 2023). It is easy to generalize this implication to show that if the lowest possible price is p, the corresponding committee satisfies $p\frac{n}{k}$-PSC. Thus, we say that a committee satisfies α-*priceability* if the smallest price p for which the committee is priceable satisfies $p \leq \alpha\frac{n}{k}$.

For all three notions, the minimal α-value achieved by a given committee can be computed in polynomial time. For local stability and EJR+, it is sufficient to iterate over the unchosen candidates and compare the size of the coalitions that would want to deviate to these candidates. The optimal price for priceability can be computed via a linear program.

	M-STV	EAR	SNTV	seq-RCV
S-STV	108 (10.1%)	262 (24.5%)	277 (25.9%)	485 (45.3%)
M-STV	–	230 (21.5%)	333 (31.1%)	415 (38.8%)
EAR		–	452 (42.2%)	459 (42.9%)
SNTV			–	599 (56.0%)

Table 4: Number of instances on which the rules disagree.

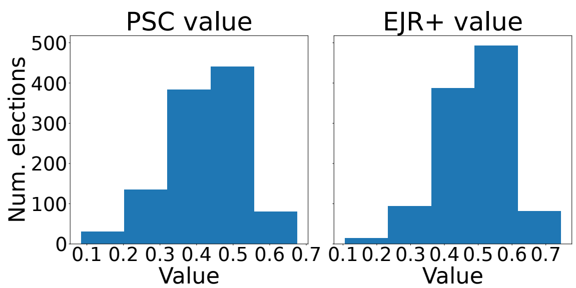

Figure 1: Histograms of PSC values and EJR+ values achievable in our elections, rounded to one decimal place.

We note that it is already NP-complete to decide whether any locally stable committee exists (Aziz et al. 2017). Further, the construction in the proof of Theorem 2 also applies to both priceability and EJR+, showing that computing the minimal α-value for these two measures is also NP-hard.

5 Experimental Results

To assess the measures defined in Section 4, we conducted several experiments on the 1070 election instances from the dataset discussed in Section 3. We highlight some of our results in this section, mostly focusing on PSC; all remaining results can be found in the full version of this paper.

We considered the following voting rules: Scottish STV (*S-STV*), Meek STV (*M-STV*), EAR, SNTV, and seq-RCV. Table 4 shows how often these rules disagree with each other (i.e., choose different committees) on our data. We observe that S-STV and M-STV agree very frequently, but not in all elections. Further, both STV variants agree with SNTV in nearly 70% of the elections, i.e., in most elections both STV variants simply select the k candidates with the most first-place votes. This is slightly less for EAR, which agrees with SNTV in only 58% of the elections. Further, seq-RCV seems to be the rule that is most different from the other rules, agreeing with SNTV in only 45% of the cases.

Minimal Values For each instance, we computed the optimal α-value for each measure (see Figure 1 for histograms of values for PSC and EJR+). For all measures, the majority of minimal α-values lie roughly in the range 0.4 to 0.6, the PSC values overall being somewhat lower than for the other measures (EJR+ and priceability in particular). Interestingly, we observe that while a 1-LS or 1-EJR+ committee is not guaranteed to exist in general, they always exist for the instances in our dataset. This observation is similar in spirit to the observation that Condorcet winners almost always exist in real-world elections (McCune and McCune 2024).

	PSC		EJR+		Priceability		LS	
	opt.	dist.	opt.	dist.	opt.	dist.	opt.	dist.
S-STV	856	0.20	826	0.23	**870**	**0.19**	829	0.23
M-STV	819	0.24	**842**	**0.22**	840	0.22	754	0.30
EAR	677	0.38	737	0.33	709	0.35	656	0.40
SNTV	**901**	**0.16**	752	0.30	832	0.23	**935**	**0.13**
seq-RCV	552	0.50	646	0.40	586	0.46	459	0.60

Table 5: For each rule and each axiom, *(i)* "opt." refers to the number of instances for which the rule achieves the optimal α-value and *(ii)* "dist." refers to the average distance between the outcome of the rule and the outcome with optimal α-value (measured in terms of number of candidates that need to be exchanged). The best values in each column appear in bold.

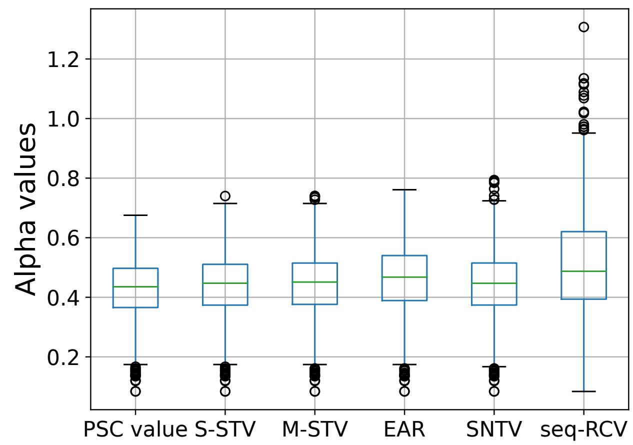

Figure 2: The PSC values achieved by voting rules, together with optimal PSC values (shown in the leftmost column).

Distance from Optimality For each voting rule, we counted *(i)* how often the rule achieves the optimal α-value and *(ii)* the average distance between the committees chosen by the voting rule and the committees optimizing the α-value.[8] The results, presented in Table 5, reveal which voting rules are "most aligned" with each of the four measures. In particular, SNTV is most aligned with the PSC and LS measures, and the STV rules are most aligned with the EJR+ and priceability measures. The strong performance of SNTV can be considered surprising insofar as the rule does not satisfy any proportionality guarantees. A possible explanation for the good values achieved by SNTV (which outperforms EAR according to all four measures) can be found in the structure of our data: often, most of the constraints that a quantified proportionality axiom like α-PSC imposes involve top-ranked candidates only, and SNTV — by definition — selects the candidates with the most first-place votes.

Values Achieved by Rules We furthermore computed the spread of α-values achieved by different voting rules over

[8] For *(ii)*, we define the distance between two committees as half of their symmetric difference.

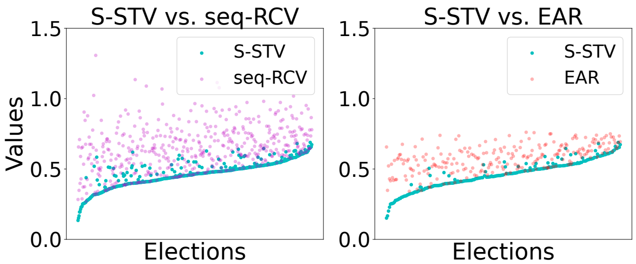

Figure 3: *Left:* PSC values achieved by Scottish STV and seq-RCV for the 485 elections where the rules disagree. *Right:* PSC values achieved by Scottish STV and EAR over the 262 elections where the rules disagree. Elections are ordered by increasing optimal PSC value.

the set of all instances and compared these values to the spread of optimal α-values. For PSC values, the results are presented as a box plot in Figure 2. Overall, all proportional rules — and the semi-proportional SNTV — perform similarly in terms of approximating optimal values, with the range of values for each of the rules coming close to those of the optimal values.[9] Somewhat surprisingly, EAR — the rule satisfying the strongest proportionality axioms (see Section 2.2) — does slightly worse than the other proportional rules. (This is also apparent in Table 5.)

Furthermore, SNTV does slightly better than the other rules w.r.t. PSC values (and the same is true for LS). A reason for that, as already discussed in the context of Table 5, is that the measures often require the k most popular candidates to be chosen: on average, 57% of the constraints corresponding to the optimal PSC value are over singleton sets of candidates, and thus correspond directly to first-place votes.

Pairwise Comparisons Finally, we considered pairwise comparisons of voting rules w.r.t. the α-values they achieve. In these comparisons, we only consider instances on which the two rules under consideration output different committees. We focus on two comparisons w.r.t. PSC values: S-STV vs. seq-RCV and S-STV vs. EAR (Figure 3). In both of these cases, S-STV does better in terms of PSC values. However, as one might expect, the overall difference in values between S-STV and the non-proportional seq-RCV is much more pronounced than the difference between S-STV and EAR. In particular, seq-RCV fails 1-PSC in 9 instances.

6 The Effect of Ballot Truncation

To understand how ballot truncation affects our results, we created ballot data with complete rankings based on the Scottish data. To create this synthetic data, we employed an iterative process that is described in the full version of this paper. Basically, when extending partial ballots of length r to length $r+1$, we consider the frequency of ballots of length at least $r+1$ which agree on the first r entries.

[9]The outlier value at $\alpha \approx 0.08$ stems from the 2012 election of North Lanarkshire, Ward 9, where 3 out of 4 candidates needed to be elected. In this election, all rules and measures choose the same committee and the only unselected candidate is greatly unpopular.

	$< 25\%$	$< 50\%$	$< 100\%$
PSC	49 (4.6%)	305 (28.5%)	837 (78.2%)
EJR+	81 (7.5%)	430 (40.2%)	1030 (96.3%)
LS	240 (22.4%)	749 (70.0%)	1066 (99.6%)
Priceability	115 (10.7%)	503 (47.0%)	1036 (96.8%)

Table 6: Analog of Table 1 for completed ballots.

We then reran all experiments from Section 5 for the 1070 synthetic election instances with complete rankings. Overall, the results for completed instances are very similar to the results for the original (truncated) instances. For instance, in 21.8% of these elections any committee of size k is compatible with PSC (compared to 27.5% in the truncated case; see Table 6). This implies that the effect of ballot truncation is rather limited for the elections we study. This is a bit of a surprise, as these results suggest that the primary reason PSC has little discriminatory power in real-world elections is *not* that voters truncate their ballots; rather, even if preferences are completed, voters do not form sufficiently large cohesive groups.

In general, as expected, the achieved α-values are slightly larger in the completed instances, with, for instance, SNTV also sometimes violating EJR+. Further, we observe that in most instances $\frac{k}{k+1}$ is a lower bound on the lowest possible priceability value, and that most priceability values achieved by the rules are clustered around that threshold.

With completed preferences, seq-RCV violates PSC in 55 elections and achieves α-values of up to 1.6 for local stability. This suggests that non-proportional methods become even more noticeably non-proportional when preferences are not truncated.

7 Conclusion

Given the absence of large cohesive groups of voters in real-world political elections, we proposed adaptations of several established proportionality axioms in which we loosen size constraints of cohesive groups, thereby creating new ways to quantify proportionality in practice. Our results show that while delivering separations in theory, in practice these proportionality measures seem to behave similarly for proportional rules. A majoritarian method like seq-RCV, on the other hand, performs poorly w.r.t. our measures. We also found that SNTV, a very simple rule without proportionality guarantees, performs well in practice in most cases. This study is a first attempt to grapple with the meaning of empirical proportionality using a large real-world dataset.

There are multiple ways to build upon our work. First, one could try to reconcile theory and practice by coming up with an axiomatic explanation for the performance of STV that goes beyond PSC. Second, our results motivate the search for proportionality axioms (or measures) that are better suited for assessing the real-world performance of voting rules. Finally, it would be interesting to obtain ballot data from some of the various other jurisdictions that use STV and to check whether those elections exhibit the same effects that we observed in the Scottish election dataset.

Acknowledgments

This work was partially supported by the Singapore Ministry of Education under grant number MOE-T2EP20221-0001.

References

Amy, D. J. 2000. *Behind the Ballot Box: A Citizen's Guide to Voting Systems*. Bloomsbury Academic.

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 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 84 – 90.

Aziz, H.; and Lee, B. E. 2020. The expanding approvals rule: improving proportional representation and monotonicity. *Social Choice and Welfare*, 54: 1–45.

Aziz, H.; and Lee, B. E. 2021. Proportionally Representative Participatory Budgeting with Ordinal Preferences. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, 5110–5118. AAAI Press.

Aziz, H.; and Lee, B. E. 2022. A characterization of proportionally representative committees. *Games and Economic Behavior*, 133: 248–255.

Balinski, M. L.; and Young, H. P. 1982. *Fair Representation: Meeting the Ideal of One Man, One Vote*. Yale University Press.

Brill, M.; and Peters, J. 2023. Robust and Verifiable Proportionality Axioms for Multiwinner Voting. In *Proceedings of the 24th ACM Conference on Economics and Computation (ACM-EC)*, 301. ACM Press. Full version arXiv:2302.01989 [cs.GT].

Burnett, C.; and Kogan, V. 2015. Ballot (and voter) "exhaustion"' under Instant Runoff Voting: An examination of four ranked-choice elections. *Electoral Studies*, 37: 41–49.

Charikar, M.; Lassota, A.; Ramakrishnan, P.; Vetta, A.; and Wang, K. 2024. Six Candidates Suffice to Win a Voter Majority. Technical report, arXiv:2411.03390 [cs.GT].

Delemazure, T.; and Peters, D. 2024. Generalizing Instant Runoff Voting to Allow Indifferences. In *Proceedings of the 25th ACM Conference on Economics and Computation (ACM-EC)*. ACM Press. Forthcoming.

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

Graham-Squire, A.; and McCune, D. 2023. An Examination of Ranked-Choice Voting in the United States, 2004–2022. *Representation*, 1–19.

Hill, I. D.; Wichmann, B. A.; and Woodall, D. R. 1987. Algorithm 123: Single Transferable Vote by Meek's Method. *The Computer Journal*, 30(3).

Hoffman, C.; Kauba, J.; Reidy, J.; and Weighill, T. 2024. Proportionality in multi-winner RCV elections: A simulation study with ballot truncation. *Communications in Statistics - Simulation and Computation*. Forthcoming.

Janson, S. 2018. Thresholds quantifying proportionality criteria for election methods. Technical report, arXiv:1810.06377 [cs.GT].

Jiang, Z.; Munagala, K.; and Wang, K. 2020. Approximately Stable Committee Selection. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, 463–472. ACM.

Kilgour, D.; Grégoire, J.; and Foley, A. 2020. The prevalence and consequences of ballot truncation in ranked-choice elections. *Public Choice*, 184: 197–218.

Lackner, M.; and Skowron, P. 2022. *Multi-Winner Voting with Approval Preferences*. Springer.

Marsh, M.; and Plescia, C. 2016. Split-ticket voting in an STV system: choice in a non-strategic context. *Irish Political Studies*, 31(2): 163–181.

McCune, D.; and Graham-Squire, A. 2024. Monotonicity Anomalies in Scottish Local Government Elections. *Social Choice and Welfare*, 63(1): 69–101.

McCune, D.; Martin, E.; Latina, G.; and Simms, K. 2024. A comparison of sequential ranked-choice voting and single transferable vote. *Journal of Computational Social Science*, 7(1): 643–670.

McCune, D.; and McCune, L. 2024. Does the choice of preferential voting method matter? An empirical study using ranked choice elections in the United States. *Representation*, 60(1): 1–16.

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

Tideman, N. 1995. The Single Transferable Vote. *Journal of Economic Perspectives*, 9(1): 27–38.

Tideman, N. 2006. *Collective Decisions And Voting: The Potential for Public Choice*. Ashgate.

Tomlinson, K.; Ugander, J.; and Kleinberg, J. 2023. Ballot Length in Instant Runoff Voting. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, 5841–5849. AAAI Press.