

Proportionality in Practice: Quantifying Proportionality in Ordinal Elections¹

Tuva Bardal, Markus Brill, David McCune, and Jannik Peters

Abstract

In real-world elections with ranked preferences, the *Single Transferable Vote (STV)* is the most widely used proportional voting method. STV is considered proportional because it satisfies an axiom known as *proportionality for solid coalitions (PSC)*, 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 the *Single Non-Transferable Vote*, 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.

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 aim to 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, many localities use the *Single Transferable Vote (STV)* [35]. STV is generally considered to be proportional because it satisfies the axiom *proportionality for solid coalitions (PSC)* formulated by Dummett [19]. In essence, PSC 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 voters in a solid coalition must all share the exact same candidate set in their ballot prefix makes PSC highly non-robust [36]. This issue has been considered from a purely theoretical point of view [2, 13], as well as through experiments on synthetic data [13]. There exist several other properties that seek to guarantee proportional representation in ways similar to that of PSC [6, 13], and that, from the aforementioned perspective, are more robust [13]. 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

¹This paper has been presented at the *39th AAAI Conference on Artificial Intelligence (AAAI 2025)*.

limited due to a scarcity of real-world ballot data. However, McCune and Graham-Squire [31] recently compiled a real-world dataset consisting of 1070 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 possible committees 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.² More generally, we consider 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 the Scottish elections, voters are allowed to only rank some candidates, i.e., they may *top-truncate* their ballots. This may impact the number and size of cohesive groups. The issues of ballot truncation and small cohesive groups have been discussed previously [13, 25, 30], but such work did not have access to actual ballot data. In order to understand the effect of ballot truncation on our proportionality measures, we generated ballot data with complete rankings based on the Scottish election data and reran all our experiments on this generated data (see Section 6).

Related Work The formal study of proportional representation in multiwinner voting was initiated by Dummett [19]. In recent years, proportional representation has been studied as a fairness criterion in several different decision-making scenarios. Most prominently, starting with Aziz et al. [5] and their notion of *extended justified representation (EJR)*, a wide body of work covers proportionality in approval-based multiwinner voting [29]. Proportionality for ordinal elections has been studied by Aziz and Lee [2, 3, 4], Brill and Peters [13], and Delemazure and Peters [18]. Varying the size of voter groups that deserve representation has been considered by Janson [26], who focuses on worst-case bounds, and Jiang et al. [27], 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).

There are a number of works focused on the experimental evaluation of (proportional) multiwinner voting rules. For instance, Boehmer et al. [10] experimentally evaluate the use of proportional multiwinner voting rules on the Polkadot blockchain, based on real-world data from Polkadot. Further, in the participatory budgeting literature, several papers have analyzed and compared rules on real-world datasets [9, 22, 12]. Besides this, most experimental papers have so far focused on synthetic data; see the survey by Boehmer et al. [11]. For instance, Elkind et al. [20] evaluate voting rules on randomly selected 2-dimensional Euclidean instances.

Finally, the issue of ballot truncation in multiwinner ranked-choice elections has been considered by Hoffman et al. [25], but has otherwise received little attention. In the single-winner instant runoff setting, Kilgour et al. [28] and Tomlinson et al. [37] study the effect on the election winner as we vary the amount of truncation, while Burnett and Kogan [15] and Graham-Squire and McCune [23] analyze how ballot exhaustion can cause the election winner not to secure a majority of total votes cast.

²As demonstrated by Brill and Peters [13], STV fails proportionality axioms that are stronger than PSC.

2 Preliminaries

For $t \in \mathbb{N}$, let $[t] = \{1, \dots, t\}$. Throughout the paper, we assume that we are given a set $N = [n]$ of voters and a set $C = \{c_1, \dots, c_m\}$ 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 [19]. 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.³

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. [6] 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.⁴ 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 $c' \in W$.*

There are several axioms that have been proposed in the setting of approval-based multiwinner voting which can be translated to the ordinal setting using a construction due to Brill and Peters [13]. We first define EJR+ [13].⁵ Note that EJR+ is not always satisfiable.

Definition 3 (EJR+). *A committee W satisfies EJR+ if there is no $\ell \in [k]$, ℓ -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) $|\{c': \text{rank}(i, c') \leq r\} \cap W| < \ell$ for all $i \in N'$.

³A related axiom is *generalized PSC* [2], which generalizes PSC to the case of weak orders. For top-truncated preferences, PSC and generalized PSC are equivalent.

⁴The same concept was independently studied by Jiang et al. [27] in the more general context of core stability.

⁵Since we are only dealing with ordinal preferences in this paper, we omit the “rank” prefix used by Brill and Peters [13]. We do not consider the axiom (*rank*-)PJR+ for computational reasons.

As an alternative to EJR+ we take *priceability* as defined by Brill and Peters [13] and Peters and Skowron [34]. Priceability is always satisfiable, and in essence tries to compute a fractional matching between voters and candidates.⁶

Definition 4 (Priceability). *A committee W is priceable if for each voter $i \in N$ there is a price function $p_i: C \rightarrow [0, 1]$ and a price $p \in \mathbb{R}^+$ such that*

- $\sum_{c \in C} p_i(c) \leq 1$ for all $i \in N$
- $\sum_{i \in N} p_i(c) \leq p$ for each $c \in W$
- $\sum_{i \in N} p_i(c) = 0$ for each $c \notin W$
- $\sum_{i \in N: \text{rank}(i, c) \leq r} (1 - \sum_{c' \in W: \text{rank}(i, c') > r} p_i(c')) \leq p$ for all $r \in [m]$ and $c \notin W$.

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 detailed descriptions, we refer to Appendix A.

The *Single Transferable Vote* (STV) is a family of rules, with different versions of STV used in different jurisdictions [35]. Following McCune and Graham-Squire [31], 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 $q = \lfloor \frac{n}{k+1} \rfloor + 1$ 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 *Scottish STV*, we also consider *Meek-STV* [24].

The *Expanding Approvals Rule* (EAR) is another family of proportional multiwinner voting rules [2]. 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.⁷

The *Single Non-Transferable Vote* (SNTV) selects the k candidates with the highest first-place vote count [21]. SNTV is sometimes referred to as *k-plurality*.

Sequential Ranked-Choice Voting (seq-RCV), which is currently used in the US state of Utah [33], 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*, and therefore can provide context for our results around methods like STV. We note that across all 1070

⁶We minimally deviate from the original definition of Peters and Skowron [34] by requiring that at most p can be paid for a candidate (instead of exactly p). This allows us to associate a price p with any committee. In particular, a price $p \geq \frac{n}{k}$ is now also possible, which will allow us to use priceability to quantify the proportionality of “disproportional” committees as well. Using the same proof as Brill and Peters [13] it follows that if a committee is priceable with a price $p < \frac{n}{k}$ the committee also satisfies the axiom PJR+ [13].

⁷The treatment of unranked candidates in EAR allows for different interpretations, as noted in Remark 3 by Aziz and Lee [2]. 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.

	Necessary fixed candidates		
	0	1	> 1
PSC	294 (27.5%)	592 (55.3%)	184 (17.2%)
EJR+	65 (6.1%)	697 (65.1%)	308 (28.8%)
LS	9 (0.9%)	441 (41.2%)	620 (57.9%)
Priceability	67 (6.3%)	702 (65.6%)	301 (28.1%)

Table 1: For each axiom, the table shows the number of elections (i) where the axiom imposes no constraints at all, (ii) where it is possible to satisfy the axiom by fixing one winning candidate and filling the remaining seats arbitrarily, and (iii) where it is only possible to satisfy the axiom by fixing multiple winning candidates on the committee. Equivalently, the “0” column gives the number of elections where all $\binom{m}{k}$ outcomes satisfy the given axiom, the “1” column gives the number of elections where the number of outcomes satisfying the axiom lies in the interval $[\binom{m-1}{k-1}, \binom{m}{k})$, and the “> 1” column gives the number of elections where the number of outcomes satisfying the axiom is strictly lower than $\binom{m-1}{k-1}$.

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” [33], 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 [1]. 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 the axiom (*rank*)-PJR+, introduced by Brill and Peters [13], which is stronger than PSC. Furthermore, committees chosen under EAR are always priceable [13].

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 [31] 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 [31] 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 more discerning than PSC, but all exhibit some degree of toothlessness. 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 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** $\succ$... and of the 406 voters who cast a ballot of the form **LM** $\succ$ **DM** $\succ$ 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 5. 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 6 (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, let⁸

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

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

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

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 $\{\text{DM}, \text{LM}\}$, followed by solid coalitions of size 1574, 1277, 1257, 671, over $\{\text{DM}\}$, $\{\text{BC}, \text{TM}\}$, $\{\text{BC}\}$, and $\{\text{IB}\}$, 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	$\{\text{DM}, \text{LM}\}$	$\{\text{DM}\}$	$\{\text{BC}, \text{TM}\}$	$\{\text{BC}\}$	$\{\text{IB}\}$
1	0.949	0.920	0.746	0.734	0.392
2	0.474	–	0.373	–	–

Table 3: The largest solid coalitions with corresponding α -values in the election from Example 1. A full overview of all solid coalitions in the election can be found in Appendix B.

Consider the committees $W = \{\text{DM}, \text{LM}, \text{BC}\}$ and $W' = \{\text{DM}, \text{BC}, \text{IB}\}$. 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: $\alpha_{\text{PSC}}(W) = 0.392$ and $\alpha_{\text{PSC}}(W') = 0.474$. In fact, W achieves the PSC value for the instance (i.e., $\alpha_{\text{PSC}}^*(I) = 0.392$), 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 $\{\text{DM}, \text{LM}\}$ is more than double the size of any solid coalition containing IB .

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:

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

⁸We use infimum rather than minimum since the set of values for which α -PSC holds is an open interval of the form $(\alpha_{\text{PSC}}, +\infty)$.

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

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

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 $\alpha_{(N', C')}^\ell$ 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.,

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

Here, each threshold value $\alpha_{(N', C')}^\ell$ 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 $\alpha_{(N', C')}^\ell$, check whether $|W \cap C'| \geq \ell$. If yes, go to the next threshold value. If not, we know that $\alpha_{\text{PSC}}(W) = \alpha_{(N', C')}^\ell$, because $\alpha_{(N', C')}^\ell$ is the largest value of α for which the corresponding PSC constraint is not satisfied by W . $\square$

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.⁹

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

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 $\alpha_{(N', C')}^\ell$, 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 $\alpha_{(N', C')}^\ell$ 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* [8]. In Appendix D, we develop a description of the algorithm which gives rise to the idea of “apportionment for non-disjoint parties,” which might be of independent interest.

⁹The proof can be found in Appendix C. The problem is trivial for $\alpha \geq 1$, as a committee satisfying PSC always exists.

	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.

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 7 (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. [16] 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 8 (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 $\text{rank}(i, c) \leq r$ for all $i \in N'$ and $|\{c' \in C : \text{rank}(i, c') \leq r\} \cap W| < \ell$ for all $i \in N'$.*

The definition of priceability is already parameterized, with a price of $p < \frac{n}{k}$ implying PSC [13]. 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.

We note that it is already NP-complete to decide whether any locally stable committee exists [6]. 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 Appendix E.

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.




Figure 1: Histograms of optimal α -values achievable in our elections, rounded to one decimal place.

	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.

Minimal Values For each instance, we computed the optimal α -value for each measure (see Figure 1 for histograms of values). 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 [32].

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.¹⁰ 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 the set of all instances and compared these values to the spread of optimal α -values. For PSC and EJR+ values, the results are presented as box plots 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.¹¹ Somewhat

¹⁰For (ii), we define the distance between two committees as half of their symmetric difference.

¹¹The 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.




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

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.




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.

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

	Necessary fixed candidates		
	0	1	> 1
PSC	233 (21.8%)	609 (56.9%)	228 (21.3%)
EJR+	40 (3.7%)	642 (60.0%)	388 (36.3%)
LS	4 (0.4%)	343 (32.0%)	723 (67.6%)
Priceability	34 (3.2%)	570 (53.3%)	466 (43.5%)

Table 6: Analog of Table 1 for completed ballots.

described in Appendix F.1. 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.

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 the results suggest that the primary reason PSC has little discriminatory power in these elections is *not* that voters truncate their ballots; rather, even if preferences are completed, voters do not form sufficiently large cohesive groups. Note, though, that our completion method attempts to extend the original data in a faithful manner, and that the results could be sensitive to the choice of completion method.

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.

References

- [1] D. J. Amy. *Behind the Ballot Box: A Citizen's Guide to Voting Systems*. Bloomsbury Academic, 2000.
- [2] H. Aziz and B. E. Lee. The expanding approvals rule: improving proportional representation and monotonicity. *Social Choice and Welfare*, 54:1–45, 2020.
- [3] H. Aziz and B. E. Lee. Proportionally representative participatory budgeting with ordinal preferences. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5110–5118. AAAI Press, 2021.
- [4] H. Aziz and B. E. Lee. A characterization of proportionally representative committees. *Games and Economic Behavior*, 133:248–255, 2022.
- [5] H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 48(2):461–485, 2017.
- [6] H. Aziz, E. Elkind, P. Faliszewski, M. Lackner, and P. Skowron. The Condorcet principle for multiwinner elections: from shortlisting to proportionality. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 84 – 90, 2017.
- [7] M. L. Balinski and H. P. Young. The webster method of apportionment. *Proceedings of the National Academy of Sciences (PNAS)*, 77(1):1–4, 1980.
- [8] M. L. Balinski and H. P. Young. *Fair Representation: Meeting the Ideal of One Man, One Vote*. Yale University Press, 1982.
- [9] N. Boehmer, P. Faliszewski, Ł. Janeczko, and A. Kaczmarczyk. Robustness of participatory budgeting outcomes: Complexity and experiments. In *Proceedings of the 16th International Symposium on Algorithmic Game Theory (SAGT)*, pages 161–178, 2023.
- [10] N. Boehmer, M. Brill, A. Cevallos, J. Gehrlein, L. Sánchez-Fernández, and U. Schmidt-Kraepelin. Approval-based committee voting in practice: A case study of (over-)representation in the Polkadot blockchain. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, pages 9519–9527, 2024.
- [11] N. Boehmer, P. Faliszewski, Ł. Janeczko, A. Kaczmarczyk, G. Lisowski, G. Pierczyński, S. Rey, D. Stolicki, S. Szufa, and T. Was. Guide to numerical experiments on elections in computational social choice. In *Proceedings of the 33rd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 7962–7970, 2024.
- [12] N. Boehmer, P. Faliszewski, Ł. Janeczko, D. Peters, G. Pierczyński, Š. Schierreich, P. Skowron, and S. Szufa. Evaluation of project performance in participatory budgeting. In *Proceedings of the 33rd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2678–2686, 2024.
- [13] M. Brill and J. Peters. Robust and verifiable proportionality axioms for multiwinner voting. In *Proceedings of the 24th ACM Conference on Economics and Computation (ACM-EC)*, page 301. ACM Press, 2023. Full version arXiv:2302.01989 [cs.GT].
- [14] M. Brill, J.-F. Laslier, and P. Skowron. Multiwinner approval rules as apportionment methods. *Journal of Theoretical Politics*, 30(3):358–382, 2018.
- [15] C. Burnett and V. Kogan. Ballot (and voter) “exhaustion” under instant runoff voting: An examination of four ranked-choice elections. *Electoral Studies*, 37:41–49, 2015.
- [16] M. Charikar, A. Lassota, P. Ramakrishnan, A. Vetta, and K. Wang. Six candidates suffice to win a voter majority. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing (STOC)*, pages 1590 – 1601, 2025.
- [17] V. Dančišin. Misinterpretation of the hagenbach-bischoff quota. *Annales Scientia Politica*, 2(1): 75–78, 2013.
- [18] T. Delemazure and D. Peters. Generalizing instant runoff voting to allow indifferences. In *Proceedings of the 25th ACM Conference on Economics and Computation (ACM-EC)*, page 50. ACM Press, 2024.
- [19] M. Dummett. *Voting Procedures*. Oxford University Press, 1984.
- [20] E. Elkind, P. Faliszewski, J.-F. Laslier, P. Skowron, A. Slinko, and N. Talmon. What do multiwinner voting rules do? An experiment over the two-dimensional Euclidean domain. In *Proceedings of the*

- 31st AAAI Conference on Artificial Intelligence (AAAI), pages 494–501. AAAI Press, 2017.
- [21] P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Multiwinner voting: A new challenge for social choice theory. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 2. AI Access, 2017.
 - [22] P. Faliszewski, J. Fils, D. Peters, G. Pierczyński, P. Skowron, D. Stolicki, S. Szufa, and N. Talmon. Participatory budgeting: Data, tools, and analysis participatory budgeting: Data, tools, and analysis. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2667–2674, 2023.
 - [23] A. Graham-Squire and D. McCune. An examination of ranked-choice voting in the united states, 2004–2022. *Representation*, pages 1–19, 2023.
 - [24] I. D. Hill, B. A. Wichmann, and D. R. Woodall. Algorithm 123: Single transferable vote by meek’s method. *The Computer Journal*, 30(3), 1987.
 - [25] C. Hoffman, J. Kauba, J. Reidy, and T. Weighill. Statistical models of ballot truncation in ranked choice elections. *Communications in Statistics - Simulation and Computation*, 2024. Forthcoming.
 - [26] S. Janson. Thresholds quantifying proportionality criteria for election methods. Technical report, arXiv:1810.06377 [cs.GT], 2018.
 - [27] Z. Jiang, K. Munagala, and K. Wang. Approximately stable committee selection. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 463–472. ACM, 2020.
 - [28] D.M. Kilgour, J. Grégoire, and A. Foley. The prevalence and consequences of ballot truncation in ranked-choice elections. *Public Choice*, 184:197–218, 2020.
 - [29] M. Lackner and P. Skowron. *Multi-Winner Voting with Approval Preferences*. Springer, 2022.
 - [30] M. Marsh and C. Plescia. Split-ticket voting in an STV system: choice in a non-strategic context. *Irish Political Studies*, 31(2):163–181, 2016.
 - [31] D. McCune and A. Graham-Squire. Monotonicity anomalies in scottish local government elections. *Social Choice and Welfare*, 63(1):69–101, 2024.
 - [32] D. McCune and L. McCune. Does the choice of preferential voting method matter? an empirical study using ranked choice elections in the united states. *Representation*, 60(1):1–16, 2024.
 - [33] D. McCune, E. Martin, G. Latina, and K. Simms. A comparison of sequential ranked-choice voting and single transferable vote. *Journal of Computational Social Science*, 7(1):643–670, 2024.
 - [34] D. Peters and P. Skowron. Proportionality and the limits of welfarism. In *Proceedings of the 21st ACM Conference on Economics and Computation (ACM-EC)*, pages 793–794. ACM Press, 2020.
 - [35] N. Tideman. The single transferable vote. *Journal of Economic Perspectives*, 9(1):27–38, 1995.
 - [36] N. Tideman. *Collective Decisions And Voting: The Potential for Public Choice*. Ashgate, 2006.
 - [37] K. Tomlinson, J. Ugander, and J. Kleinberg. Ballot length in instant runoff voting. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5841–5849. AAAI Press, 2023.

Tuva Bardal
 University of Warwick
 Coventry, United Kingdom
 Email: tuva.bardal@warwick.ac.uk

Markus Brill
 University of Warwick
 Coventry, United Kingdom
 Email: markus.brill@warwick.ac.uk

David McCune
William Jewell College
Liberty, Missouri, United States
Email: mccuned@william.jewell.edu

Jannik Peters
National University Singapore
Singapore
Email: peters@nus.edu.sg

Appendix

The appendix is structured as follows. In Appendix A, we provide more details on the voting rules STV and EAR. Appendix B lists all solid coalitions of the election instance from Example 1. In Appendix C, we give a proof of Theorem 2, and in Appendix D we discuss connections between the algorithm described in Section 4.3 and apportionment methods. Appendix E provides more details on the experiments discussed in Section 5. Finally, Appendix F is dedicated to experiments on completed ballot data (see Section 6).

A Definitions of Voting Rules

Single Transferable Vote. The rule used most frequently in real-world elections is the method of *single transferable vote* (STV). STV is actually a family of rules, with different versions of STV used in different jurisdictions. Following McCune and Graham-Squire [31] we define it here in the way it is used in Scottish local elections. This version uses the Hagenbach-Bischoff quota $q = \lfloor \frac{n}{k+1} \rfloor + 1$ (which is different from the Droop quota $\lceil \frac{n}{k+1} \rceil$; see 17). STV then proceeds in rounds. It starts off 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. If that is the case, this candidate is elected and the voters ranking this candidate first are reweighted proportionally (if there were voters with total weight w ranking the candidate first, each voter with weight w_i is reweighted by a factor of $w_i \frac{w-q}{w}$). Thus, the total weight of the election is decreased by exactly q . If there is no such candidate, the candidate with the least weighted first-place votes is removed. We refer to this version of STV as *Scottish STV*. We also use the version of STV known as *Meek-STV*, a commonly suggested STV alternative. There are two primary differences between Meek and Scottish STV. Under Meek, candidates who have previously won a seat can receive vote transfers in later rounds, and the quota can decrease as the transfer process unfolds. Due to the inherent complexity of Meek-STV we defer its description to Hill et al. [24].

Expanding Approvals Rule. The *expanding approvals rule* (EAR) is another family of proportional multiwinner voting rules, suggested by Aziz and Lee [2] as an alternative to STV. Just like STV it uses a quota q and starts off by assigning each voter a weight $w_i = 1$. It then iterates over all possible ranks $r \in [m]$. 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 collectively decreases the weight of all such 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 gets increased by 1.

B Additional Data for Example 1

Consider again the instance discussed in Example 1, given in the table below, where capital letters A, B, C, ... are used to denote candidates; the for each candidate, the corresponding letter is given in the rightmost column. The following list contains all 125 maximal solid coalitions for this instance. Each solid coalition (N', C') appears in the format $(C': |N'|)$, and the list is ordered by the size $|N'|$ of coalitions.

(EF: 1624), (E: 1574), (CG: 1277), (C: 1257), (B: 671), (BCG: 554), (F: 525), (ABCDEFG: 460), (BEF: 405), (A: 382), (D: 365), (G: 358), (AEF: 345), (CEFG: 292), (BC: 239), (AE: 228), (BD: 228), (CEF: 216), (CEG: 212), (ABCEFG: 200), (BCEFG: 197), (CE: 167), (BE: 159), (AB: 158), (BCDEFG: 156), (EFG: 140), (ACG: 132), (DEF: 117), (CFG: 107), (ABCG: 107), (ABEF: 107), (BF: 102), (ABD: 101), (CDG: 82), (AC: 79), (ACEFG: 78),

(BCDG: 77), (AD: 73), (ABDEF: 72), (ABE: 68), (BDEF: 65), (BCE: 64), (CDEFG: 64), (BCEG: 63), (DE: 62), (EG: 62), (ABC: 61), (ABCDG: 61), (BG: 59), (ABCDEF: 53), (BCEF: 49), (ABF: 48), (BCD: 47), (CF: 45), (AF: 44), (ABCEF: 42), (ABCDEG: 42), (CD: 40), (ACE: 40), (ACDEFG: 40), (BDE: 39), (ABCEG: 35), (ABCDG: 34), (BCFG: 33), (FG: 29), (BCF: 29), (BDF: 29), (DF: 28), (ACEF: 28), (ACEG: 28), (BEFG: 27), (ABCFG: 25), (BCDEF: 25), (ABEFG: 23), (ABDE: 22), (AG: 21), (AEFG: 21), (BCDEG: 21), (ABCE: 20), (ADEF: 20), (ABG: 18), (BDG: 18), (BCDFG: 17), (BEG: 16), (BFG: 15), (ACDG: 15), (AEG: 14), (ABDEFG: 14), (ADE: 13), (ACFG: 13), (DG: 12), (ABCD: 12), (CDEF: 12), (ACD: 10), (ADF: 9), (ABDF: 9), (ABFG: 9), (CDEG: 9), (ABCDE: 8), (ACF: 7), (CDF: 7), (BCDE: 7), (CDFG: 7), (DEFG: 7), (ADG: 6), (CDE: 6), (ABEG: 5), (ACDEF: 5), (BDEFG: 5), (AFG: 4), (DEG: 4), (ABCF: 4), (BCDF: 4), (ABCDF: 4), (ACDEG: 4), (ACDFG: 4), (ADEF: 4), (DFG: 3), (ACDE: 3), (ABDG: 2), (ACDF: 2), (ADEG: 2), (BDFG: 2), (BDEG: 1).

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

C Proof of Theorem 2

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

Proof. Membership in NP follows from Theorem 1, as any W with $\alpha_{\text{PSC}}(W) \leq \alpha$ witnesses the satisfiability of α -PSC.

To show hardness, we reduce from 3-Hitting Set. Recall that in the 3-Hitting Set problem, we are given a set S , a collection of size-three subsets $S_1, \dots, S_j$ of S and an integer $h < j$, and the goal is to find a set $H \subseteq S$ of size $|H| = h$ such that $H \cap S_i \neq \emptyset$ for all $i \in [j]$. For each $i \in [j]$, let $S_i = \{s_i^1, s_i^2, s_i^3\}$. For a given instance $(S, \{S_1, \dots, S_j\}, h)$, we construct a corresponding election instance as follows.

Let $C = S \cup D$, where $D = \{d_1, \dots, d_j\}$ is a set of j dummy candidates. For each $i \in [j]$, there are two voters v_i^1, v_i^2 with $A_i = S_i \cup \{d_i\}$ and preferences

$$v_i^1 : d_i \succ s_i^1 \succ s_i^2 \succ s_i^3 \quad \text{and} \quad v_i^2 : d_i \succ s_i^2 \succ s_i^3 \succ s_i^1.$$

Thus, $n = 2j$. Finally, we let $k = h + j$ and $\alpha = \frac{k}{n} < 1$.

Since $\alpha \frac{n}{k} = 1$, every individual voter on its own constitutes a 1_α -large solid coalition and, for each $i \in [j]$, the voters v_i^1 and v_i^2 together form a 2_α -large solid coalition over the prefix $\{d_i\} \cup S_i$. Hence, a committee satisfying α -PSC must contain all dummy candidates and at least one candidate from each of the sets $S_1, \dots, S_j$. It follows that there is a committee of size k that satisfies α -PSC if and only if there exists a hitting set of size h .

This proof can easily be extended to any $\alpha < 1$ by either cloning the voters or adding new sets consisting only of a single otherwise unused element. $\square$

D An Apportionment Perspective

The algorithm used to compute the PSC value of an instance (see Section 4.3) has similarities to the *D'Hondt apportionment method* [8]. Apportionment methods distribute parliamentary seats among

parties based on the parties' vote counts in a party-list election. In particular, the D'Hondt method can be computed by (1) constructing a table that contains the vote counts of all parties, the vote counts divided by 2, the vote counts divided by 3, and so on; and (2) iteratively assigning a seat to the party corresponding to the next-highest number in the table, until all k seats have been assigned.¹²

The ILP-based algorithm in Section 4.3 can be phrased in a similar way, based on the observation that the threshold values $\alpha_{(N', C')}^\ell = \frac{|N'|}{n} \cdot \frac{k}{\ell}$ are proportional to $\frac{|N'|}{\ell}$. Therefore, we could interpret the maximal solid coalitions as 'parties' and construct a table of the parties' sizes, the parties' sizes divided by 2, and so on. In particular, the numbers corresponding to a solid coalition (N', C') are $|N'|, \frac{|N'|}{2}, \frac{|N'|}{3}, \dots, \frac{|N'|}{|C'|}$. Then, we can iterate over the numbers in this table in non-increasing order, just like the D'Hondt method would do.

There are two main differences between the algorithm from Section 4.3 and D'Hondt's method. First, our algorithm allocates "representation guarantees" instead of seats: When considering $\frac{|N'|}{\ell}$ corresponding to solid coalition (N', C') , the constraint $|W \cap C'| \geq \ell$ is added. Second, we may give out more than k representation guarantees: Since the candidate sets of different solid coalitions are not necessarily disjoint, candidates can satisfy more than one solid coalition at once. We only stop allocating further representation guarantees if doing so would lead to an infeasible collection of guarantees (as determined by our ILP).

This alternative description of the algorithm for computing the PSC value of an instance gives rise to the idea of *apportionment for non-disjoint parties*. Interestingly, any divisor method can be employed instead of the D'Hondt method. For example, the *Sainte-Laguë* method (aka Webster), which satisfies attractive properties [7], uses divisors $1, 3, 5, 7, \dots$ instead of $1, 2, 3, 4, \dots$.

E Detailed Analysis of Experiments

E.1 Agreement of Rules

To complement the results on how often rules disagree, we calculated the average distance between each pair of rules. The results are given in Table 7. The pairwise distance between rules is low overall. As we have seen in Table 4, except for SNTV and seq-RCV, each pair agrees in a majority of cases. Moreover, the committees returned by the different rules often differ by only one candidate when they disagree.

	S-STV	M-STV	EAR	SNTV	seq-RCV
S-STV	0	0.10	0.25	0.26	0.46
M-STV	-	0	0.22	0.31	0.39
EAR	-	-	0	0.44	0.44
SNTV	-	-	-	0	0.58
seq-RCV	-	-	-	-	0

Table 7: Average distance between pairs of rules.

E.2 Alignment of Measures

In order to check how closely aligned our four measures are, we calculated, for each election, the optimal committee w.r.t. each measure (i.e., the committee minimizing α). Table 8 shows how much the committees optimizing different measures differ from each other. The difference between the

¹²For example, see Brill et al. [14, page 362].

committees optimizing the LS value and committees optimizing other measures is overall higher than the difference between the other measures. One possible explanation for this is that the LS measure is mostly restricted to only consider first-place candidates, while the other measures impose constraints more broadly.

	EJR+	Priceability	LS
PSC	193 (0.18)	91 (0.08)	254 (0.24)
EJR+	–	137 (0.12)	346 (0.33)
Priceability		–	275 (0.26)

Table 8: Number of elections where measure-minimizing committees disagree and average distance between those committees (in parenthesis).

E.3 Values Achieved by Rules

The LS and priceability values achieved by rules are given in Figure 4.




Figure 4: The LS and priceability values achieved by voting rules

F Experimental Results for Completed Ballots

This section contains the results of experiments with completed ballot data, generated based on the Scottish data set. We first describe the method that we used to complete the data in Section F.1. As mentioned in Section 6, completing the data had limited effect on the experiments we performed. In Section F.2 we give an example to illustrate why the method described in Section F.1 may not necessarily yield elections with (many) more sufficiently large solid coalitions. We give the results of the experiments on completed data in Sections F.3 to F.6.

F.1 Generating Completed Ballots

We used the following method to create this synthetic data. The process involves a series of steps, each time extending partial ballots of length r to length $r + 1$ based on the probability distribution of ballots which agree on the first r entries.

To be more concrete, suppose there are 10 ballots of the form ABC . To increase these ballots by length 1, we consider all ballots of the form $ABC*$ which have length at least 4. Suppose there are 38 such

ballots as shown in Table 9. We extend the ballots of the form ABC proportionally, (column 3) and round to a whole number using Hamilton’s apportionment method (column 4).

Ballot	Number	Prop.	Num. ballots
ABCD	9	$(9/38) \cdot 10 = 2.368$	2
ABCE	12	$(12/38) \cdot 10 = 3.158$	3
ABCF	17	$(17/38) \cdot 10 = 4.474$	5
Total	38	10	10

Table 9: Extending ballots of length 3 to length 4.

In theory, this process can be iterated to create complete ballots, assuming a sufficient set of complete ballots. However, in practice it does not make sense to keep extending preferences on a given ballot if the number of ballots of length at least $r + 1$ is not sufficiently large in comparison to the number of ballots of length r . Thus, each ballot was extended until either the ballot contained a complete ordering of the candidates, or until the number of ballots of length $r + 1$ was less than 10% of the number of ballots of length r .

If after this process, a ballot was not complete, we then completed the ballot by choosing between the candidates not yet listed on the ballot uniformly at random.

F.2 Example Instance

As an example of why the number of large cohesive groups in elections does not increase substantially when moving from truncated data to data with full preferences, consider the 2022 election of Glasgow, Ward 20 (see https://en.wikipedia.org/wiki/2022_Glasgow_City_Council_election#Baillieston). In this election with $k = 3$, the two Labour candidates in the real instance with truncated ballots have 38.4% of the first-place votes and form a solid coalition consisting of 0.313% of the electorate (and thus too small to deserve representation under PSC). Furthermore, out of the voters that rank either of the Labour candidates first, 2.75% only vote for this one candidate. Our completion method assigns the other Labour candidate the second place in most of these ballots, but not all of them, yielding a new solid coalition over these candidates that consists of 32.4% of the electorate — barely below the $\frac{n}{3}$ threshold imposed by PSC.

F.3 Minimal Values

We computed the optimal α -value for each measure on the completed data. The histograms of encountered values is shown in Figure 5. The minimal α -values become slightly higher for all measures. The difference in values is most pronounced for EJR+ and priceability. In particular, the priceability value for most instances is lower bounded by $\frac{k}{k+1}$: essentially, for lower values the entire electorate would be able to afford more than k candidates at rank m , hence violating priceability.

F.4 How often do rules produce optimal committees?

For each voting rule, we counted how often the rule achieves each optimal α -value. The results can be found in Table 10, where, for each measure, the value associated with the best-performing rule is highlighted. The overall performance of the rules was slightly worse on the completed data. As in the truncated case, SNTV is most aligned with PSC and LS. On the other hand, on the full preferences, EAR is most aligned with EJR+ and priceability, and now outperforms Scottish STV and Meek STV. Note also that EAR performs better w.r.t. priceability value on the completed data than on the truncated data.




Figure 5: Histograms of optimal values, rounded to one decimal place (completed preferences).

	PSC	EJR+	Priceab.	LS
S-STV	763	591	667	743
M-STV	748	637	675	700
EAR	584	645	739	579
SNTV	860	454	543	906
seq-RCV	409	618	559	336

Table 10: For each rule and each axiom, the number of instances for which the rules achieves the optimal α -value for completed preferences. The best values in each column appear in bold

F.5 Values Achieved by Rules

We compared the spread of α -values achieved by rules to the optimal α -values for each measure. The results are presented in Figure 6 and Figure 7. We observe that the α -values of all rules become slightly higher in the complete case – this is in line with the minimal possible values increasing somewhat. Furthermore, the number of seq-RCV committees that have a value of $\alpha > 1$ increases for each measure. In other words, for each axiom, there are more instances where seq-RCV returns a committee that does not satisfy the axiom. We also observe a difference in the spread of priceability values for all rules, compared to that in the truncated case. This has to do with the $\frac{k}{k+1}$ lower bound on the priceability value: the priceability values achieved by rules cluster around this area, to the point where any priceability value that does not is considered an outlier.




Figure 6: The PSC and EJR+ values achieved by voting rules (completed preferences).




Figure 7: The LS and priceability values achieved by voting rules (completed preferences).

	M-STV	EAR	SNTV	seq-RCV
S-STV	92	306	376	601
M-STV	–	303	399	555
EAR		–	536	570
SNTV			–	731

Table 11: Number of instances pairs of rules disagree on (completed preferences).

F.6 Pairwise Comparisons

We again considered the disagreement and distance between rules. The results are given in Tables 11 and 12. On the completed data, S-STV and M-STV agree slightly more often than in the truncated case. Aside from this, disagreement is higher than on the truncated data for all pairs of rules. This difference is perhaps most noticeable for pairwise comparisons that include seq-RCV, which disagrees with every other rule on $t > 100$ more instances than in the experiments on the truncated data.

	M-STV	EAR	SNTV	seq-RCV
S-STV	0.08	0.30	0.35	0.60
M-STV	–	0.30	0.38	0.50
EAR		–	0.53	0.56
SNTV			–	0.76

Table 12: Average distance between pairs of rules as fraction of seats (completed preferences).

Furthermore, we considered the difference in PSC values for pairs of rules, restricted to instances on which the rules disagree. We again focused on S-STV vs. seq-RCV and S-STV vs. EAR (Figure 8). As was the case on the truncated data, the difference between seq-RCV and S-STV is more pronounced than what is the case for S-STV and EAR. As mentioned previously, seq-RCV fails 1-PSC more often for completed ballots.




Figure 8: *Left:* Differences in PSC values between Scottish STV and Sequential RCV when elections where the rules agree are excluded (completed preferences). *Right:* Differences in PSC values between Scottish STV and EAR when elections where the rules agree are excluded (completed preferences). Elections are ordered by increasing optimal PSC value.