Mathematical Social Sciences 140 (2026) 102506

Contents lists available at ScienceDirect

Mathematical Social Sciences

journal homepage: www.elsevier.com/locate/mss

New fairness criteria for truncated ballots in multi-winner ranked-choice elections

Adam Graham-Squire [a,*], Matthew I. Jones [b,c], David McCune [d]

[a] Department of Mathematical Sciences, High Point University, 1 University Parkway, High Point, NC, 27268, United States of America
[b] Department of Mathematics, Colby College, 5839 Mayflower Hill Dr, Waterville, ME 04901, United States of America
[c] Yale Institute for Network Science, Yale University, 17 Hillhouse Ave, New Haven, CT, 06511, United States of America
[d] Department of Mathematics and Data Science, William Jewell College, 500 College Hill, Liberty, MO, 64068-1896, United States of America

ARTICLE INFO

MSC:
primary 91B12
secondary 91B14

Keywords:
Single transferable vote
Fairness criteria
Empirical results
Truncated ballots

ABSTRACT

In real-world elections where voters cast preference ballots, voters often provide only a partial ranking of the candidates. Despite this empirical reality, prior social choice literature frequently analyzes fairness criteria under the assumption that all voters provide a complete ranking of the candidates. We introduce new fairness criteria for multiwinner ranked-choice elections concerning truncated ballots. In particular, we define notions of the *independence of losing voters blocs* and *independence of winning voters blocs*, which state that the winning committee of an election should not change when we remove truncated ballots which rank only losing candidates, and the winning committee should change in reasonable ways when removing ballots which rank only winning candidates. Of the voting methods we analyze, the Chamberlin-Courant rule performs the best with respect to these criteria, the expanding approvals rule performs the worst, and the method of single transferable vote falls in between.

1. Introduction

This article introduces new fairness criteria for multiwinner ranked-choice elections which stipulate how voting methods should behave when certain types of truncated ballots are removed from the ballot data. To give an informal motivation for these criteria, consider the following scenario: an election contains seven candidates $C_1, \dots, C_7$ and voters cast preference ballots (possibly providing only partial preferences) to determine a winning committee of size three. Suppose the winning committee under a given voting method is $\{C_1, C_2, C_3\}$ but when we remove a handful of ballots which rank only candidate C_7, the winning committee changes to $\{C_4, C_5, C_6\}$. It seems normatively undesirable that voters who cast ballots ranking only a losing candidate, and therefore achieve no representation, can determine the composition of the winner set by participating in the election. Put another way, it seems strange that candidate C_1's ability to win a seat depends on the participation of voters who care only about C_7 (and are presumably indifferent about the other candidates). Similarly, suppose we remove some ballots which rank only C_1 and the winning committee changes to $\{C_1, C_4, C_5\}$. It seems undesirable that voters who support only C_1, and retain their candidate in the winner set if they abstain, could affect which other voters achieve representation. We articulate three related fairness criteria along these lines. The first is the *independence of losing voter blocs (ILVB) criterion*, which stipulates that if we remove truncated ballots which rank only losing candidates then the winning committee should not change. The other two criteria are analogues of the ILVB criterion but for winning candidates. The *independence of winning voter blocs (IWVB) criterion* requires that if we remove truncated ballots which rank only winning candidates then the only allowable change to the winning committee is that some subset of these winning candidates is replaced in the committee. Similarly, the *IWVB* criterion* stipulates that if we remove truncated ballots which rank only winning candidates and all of these candidates retain their seats after this removal, then the winning committee should not change. We provide worst-case analyses for several voting methods with respect to these criteria and give empirical results using a large dataset of real-world multiwinner elections.

Much of the previous social choice literature studies the case of single-winner elections, often under the assumption that voters provide a complete ranking of the candidates. The modern theory of fairness criteria and single-winner elections dates to the pioneering work of Kenneth Arrow (Arrow, 1951), which assumes complete preference information. For a summary of voting criteria in the single-winner case with complete preferences, see Nurmi (1999). Some researchers have proposed criteria explicitly built around the concept

* Corresponding author.
E-mail addresses: agrahams@highpoint.edu (A. Graham-Squire), mijones@colby.edu (M.I. Jones), mccuned@william.jewell.edu (D. McCune).

https://doi.org/10.1016/j.mathsocsci.2026.102506
Received 11 August 2024; Received in revised form 1 November 2025; Accepted 8 January 2026
Available online 10 January 2026

of truncated ballots, such as non-manipulability by sincere preference truncation (Brams, 1982; Fishburn and Brams, 1984). Another example is Woodall's "mono-add-plump" criterion, which states that if candidate A is the winner of an election and we add ballots which rank A first and rank no other candidates, then A should still win Woodall (1994). But fairness criteria explicitly concerning truncated ballots tend to receive less attention in the social choice literature.

The multiwinner ranked-choice setting is less-studied than the single-winner, and correspondingly there has been less study of fairness criteria for multiwinner elections. While some single-winner criteria can be adapted to the multiwinner setting with little to no modification, many classical criteria (such as the Condorcet criterion) do not translate as easily. For an introduction to fairness criteria for multiwinner voting methods, see Elkind et al. (2017), Felsenthal and Moaz (1992). A complicating factor in the multiwinner setting is the differing goals of various multiwinner voting methods. Some aim for proportional representation and others do not, for example. We focus on multiwinner voting rules designed for proportional representation, whereas much of the prior literature focuses on axioms regarding the meaning of proportional representation (Aziz and Lee, 2020; Brill and Peters, 2023; Dummett, 1984; Monroe, 1995; Sáchez-Fernández et al., 2017). There are different definitions of "proportional" in a multiwinner ranked-choice setting, and in Section 5, we focus on Dummett's notion of proportionality for solid coalitions (Dummett, 1984). Our independence of voter blocs criteria are closely related to the independence of irrelevant alternatives criterion for single-winner elections and its variants, including notions of the so-called "spoiler effect" (Arrow, 1951; Börgers, 2010; Heckelman and Chen, 2013; McCune and Wilson, 2024; Miller, 2019). Additionally, certain kinds of no-show paradoxes (Fishburn and Brams, 1983) are a special case of violations of our ILVB criterion.

In this article, the primary voting method of interest is the version of single transferable vote (STV) used in Scottish local government elections because these elections are the data source for our empirical results. To give context for our results concerning Scottish STV, we also study several other proportional methods such as Meek STV (Meek, 1994), the Chamberlin-Courant voting rule (Chamberlin and Courant, 1983), and the expanding approvals rule proposed by Aziz and Lee (Aziz and Lee, 2020). In our worst-case analysis we show that, for each voting method other than Chamberlin-Courant, it is possible to construct elections in which the removal of some *bullet votes* (ballots that rank a single candidate) creates a new winner set disjoint from the original winner set. For the Chamberlin-Courant rule, we prove the winning committee does not change when removing truncated ballots that only rank losers, or removing ballots that rank only winners who retain their seats.

We find that in approximately 30% of elections from our real-world Scottish dataset, at least one voting method returns a violation of our voting bloc criteria, with the percentage for a fixed voting method varying substantially. For example, about 2% of elections return a violation of the loser bloc criterion for Meek STV, which rises to 5% for the expanding approvals rule. The Chamberlin-Courant rule and its variants have the fewest violations while the expanding approvals rule has the most. Instead of just selecting a different candidate from the same party, these violations can have significant political ramifications. The vast majority of violations result in seats changing political parties, altering the political landscape and perhaps influencing policy outcomes.

The paper is structured as follows. In Section 2 we provide all necessary preliminary information. In particular, we define our voting methods and voter-bloc fairness criteria. We also provide a description of the real-world dataset of elections we use. In Section 3 we provide worst-case analyses for violations of our criteria for all voting methods defined in Section 2. Section 4 details our search for violations of our criteria in a large dataset of real-world multiwinner elections. In Section 5 we briefly discuss the relationship between proportionality and our voter bloc criteria, and Section 6 concludes.

2. Preliminaries

We study only ranked-choice elections where each voter provides a (possibly partial) linear ranking of candidates using a preference ballot. Ballots are combined into a *preference profile* P, which provides the number of ballots cast of each type. Let m denote the number of candidates in an election and k denote the size of the winner set, which equals the number of available legislative seats. We refer to an ordered pair (P, k) as an *election*.

A *multiwinner voting method* (or multiwinner voting rule) is a function which takes as input an election and outputs a winning committee of size k. Even though a voting method may output multiple committees due to ties, empirical examples of ties in elections are virtually nonexistent, so we avoid the issue of ties throughout our work and assume a unique winner set. We denote the winning committee under a given voting rule $W(P, k)$. The voting rule under consideration will be clear any time we use such notation, so we do not incorporate the voting method into our winner set notation.

2.1. Multiwinner voting methods

We focus on the multiwinner electoral setting where the goal of a voting method is to provide some form of proportional representation. Accordingly, we restrict attention to multiwinner rules designed with proportionality in mind. The six rules we study are:

- Scottish STV.
- Meek STV.
- EAR OM: The Expanding Approvals Rule (EAR) with truncated ballots processed using the "optimistic method".
- EAR PM: EAR with truncated ballots processed using the "pessimistic method".
- CC OM: The Chamberlin–Courant (CC) rule with truncated ballots processed using the "optimistic method".
- CC PM: The CC rule with truncated ballots processed using the "pessimistic method".

Formal definitions of the methods can be found in the Appendix. In brief, the main differences between the two types of STV are (1) in Meek STV, the quota can be reduced in later rounds as truncated ballots are exhausted but in Scottish STV the quota is fixed, and (2) in Meek STV, previously elected candidates can receive vote transfers but in Scottish STV such candidates receive no transfers after their election. The EAR methods are multiwinner proportional analogs of the single-winner Bucklin voting rule, where lower preferences are systematically revealed and supporters of winning candidates must "pay a price" for their candidate's election. CC rules assign each voter a representative using Borda scores, and thus can be considered a multiwinner version of the Borda count.

We chose these methods for three main reasons. First, the STV and EAR variants satisfy classical proportionality axioms such as Dummett's Proportionality for Solid Coalitions (Dummett, 1984). Second, Scottish STV and Meek STV are used in real-world elections (in Scotland and New Zealand, respectively), making them especially relevant for our empirical analysis in Section 4. Finally, while the Chamberlin–Courant rules do not satisfy most classical proportionality axioms, they are designed to capture proportional representation in a broader, more informal sense (Chamberlin and Courant, 1983). Of course, we could include more proportional or semi-proportional methods in our study, but we feel six methods is enough to provide comparative analysis while keeping the narrative streamlined.

2.2. Our dataset: Scottish local government elections

For the purposes of local government, Scotland is partitioned into 32 council areas, each of which is governed by a council. The councils provide a range of public services that are typically associated with local governments, such as education, waste management, and road maintenance. A council area is divided into wards, each of which elects a set number of councilors to represent the ward on the council. The number of councilors representing each ward is determined primarily by the ward's population, but typically a ward has three or four seats. Every five years each ward holds an election where all seats available in the ward are filled using Scottish STV. Every Scottish ward has used STV for local government elections since 2007.

The data was collected for (McCune and Graham-Squire, 2024), and is now publicly available at https://github.com/mggg/scot-elex. The dataset contains 1100 elections, of which 1070 satisfy $k > 1$. Most of these elections satisfy $k \in \{3,4\}$ and $m \in \{6,7,8,9\}$; see McCune and Graham-Squire (2024) for more details about the dataset.

Voters are allowed to provide a complete ranking of the candidates, but voluntary truncation is very common. Approximately 14.0% of ballots in the dataset rank only a single candidate and 58.0% rank fewer than k candidates. By contrast, only 13.2% provide a complete ranking, where by "complete ranking" we mean a ballot which ranks m or $m-1$ candidates. Because the ballots are so truncated in general, this dataset is a valuable resource for examining fairness criteria involving truncated ballots.

2.3. Losing and winning voter bloc criteria

We now define our new criteria, beginning with independence of losing voter blocs.

Definition 1. Let $\mathcal{L}$ be a subset of losing candidates and $\mathcal{B}(\mathcal{L})$ be a set of ballots such that only candidates from $\mathcal{L}$ are ranked. A voting method satisfies **independence of losing voter blocs** (ILVB) if whenever we remove the ballots $\mathcal{B}(\mathcal{L})$ from the election, creating the modified profile P', then $W(P,k) = W(P',k)$.

The motivation behind ILVB is that if voters cast truncated ballots which rank only losing candidates, these voters achieve no representation and thus removing their ballots should have no effect on the composition of the winning committee. We illustrate a violation of ILVB with a 2012 STV election from the Scottish elections database.

Example 1. The 2012 election in the fifth ward of the East Ayrshire council area unfolded as shown in the top of Table 1. The table is a "votes by round" table, which shows the number of votes controlled by each candidate in a given round of the STV process. We display such tables to summarize the election because the preference profile is usually much too large. Holden is by far the weakest candidate by any reasonable measure of "weak", and is quickly eliminated after Knapp achieves quota. Then Scott is narrowly eliminated by 3.5 votes and the resulting winners are Knapp, Ross, and Todd.

If we remove 20 bullet votes for Holden (recall that a bullet vote is a ballot which ranks only one candidate) then the election unfolds as shown in the bottom of Table 1. Surprisingly, removing a handful of bullet votes for a losing candidate (and, in this case, a very weak losing candidate) causes a change in the winner set. The removal of these ballots lowers the quota by five votes, allowing for four additional votes to transfer to Scott in the second round, and subsequently Ross has the fewest votes in the third round and is eliminated. Those 20 Holden voters, with no expressed preference for Ross or Scott, were pivotal in Ross winning a seat.

Analysis of the full preference profile for this election arguably demonstrates additional paradoxical behavior beyond a violation of the ILVB criterion. The election contains candidates from three parties:

Table 1
An example of a violation of ILVB. The top table shows the vote totals in each round for the original 2012 election in the fifth ward of the East Ayrshire council area. The bottom table shows the corresponding vote totals when we remove 20 bullet votes for Holden. A bold number represents when a candidate surpasses quota.

Original Election, quota = 833				
Candidate	Votes by Round			
Holden (Con)	135	138.7		
Knapp (Lab)	**1250**			
Ross (SNP)	735	759.4	789.4	**924.4**
Scott (Lab)	417	743.9	785.9	
Todd (SNP)	791	814.4	822.7	**949.7**

Modified Election, quota = 828				
Candidate	Votes by Round			
Holden (Con)	115	118.7		
Knapp (Lab)	**1250**			
Ross (SNP)	735	759.6	789.6	
Scott (Lab)	417	747.8	789.9	**835.4**
Todd (SNP)	791	814.6	823.0	**1474.4**

Conservative (Con), Labor (Lab), and the Scottish National Party (SNP). As far as we are aware, it is common knowledge in Scottish politics that the Conservative party is generally to the right of the other two. Furthermore, of the 135 voters who rank Holden first, 55 rank a Lab candidate second and 27 rank an SNP candidate second, so if we were to arrange the parties on a 1-dimensional axis then the Conservative candidate would be on the right, the SNP would be on the left, and Labor would be in the center. By removing bullet votes for Holden, we make the overall electorate more left-leaning, yet the winning committee becomes less left-leaning.

In this article we focus on proportional or semi-proportional multiwinner voting rules, but the ILVB criterion has relevance in the single-winner context as well. Of course, proportionality is not a meaningful concept when electing a single winner, but the ILVB criterion may be normatively important when viewed as a version of Arrow's independence of irrelevant alternatives (IIA) criterion (Arrow, 1951) for truncated ballots. To see why, consider the single-winner August 2022 Special House Election in Alaska, an election which marked the first use of STV in the state (for deeper analysis of this election, see Clelland, 2026; Graham-Squire and McCune, 2024). The election contained the three candidates Nick Begich, Sarah Palin, and Mary Peltola. Peltola won the election using STV, thereby earning Alaska's only US House seat. If a few thousand ballots which ranked only Palin were removed from the ballot data then the winner would have been Begich. The conventional wisdom concerning this election is that Peltola is best characterized as a moderate left candidate, Begich as moderate right, and Palin as far right. Thus, the Alaska ILVB violation is "paradoxical" in the same way Table 1 is paradoxical: in the Alaska House election, if we remove some far-right voters then the election winner moves farther to the right. This may be normatively undesirable.

Our next criterion is similar to ILVB but focuses on the removal of ballots ranking only winning candidates, and the motivation behind the criterion is the same. If, for example, A is a winning candidate and we remove some bullet votes for A, it seems normatively undesirable that other winning candidates can lose their seats. If anything, those other winning candidates should "better" represent their constituencies with the removal of ballots for A. Unlike the ILVB criterion, the criteria defined below are not relevant to the single-winner context.

Definition 2. Let $\mathcal{W}$ be a subset of winning candidates, $1 \leq |\mathcal{W}| < k$, and $\mathcal{B}(\mathcal{W})$ be a set of ballots such that only candidates from $\mathcal{W}$ are ranked. A voting method satisfies **independence of winning voter blocs** (IWVB) if whenever A is a winning candidate, $A \notin \mathcal{W}$, and

Table 2
An example of a violation of IWVB*. The top table shows the original 2022 election in the eighth ward of the North Ayrshire council area. The bottom table shows the election when we remove 199 bullet votes for McDonald.

Original Election, quota = 1007

Candidate	Votes by Round						
Burns (SNP)	**1470**						
Collins (Grn)	171	219.2	227.6	239.1	278.9		
Craig (SFP)	64	70.3	76.5				
Jackson (LD)	106	111.7	130.4	139.9			
Johnson (SNP)	323	696.5	706.3	718.6	730.3	874.0	
McDonald (Lab)	**1096**						
Stephen (Con)	795	797.8	818.5	839.8	885.3	913.6	**1097.1**

Modified Election, quota = 957

Candidate	Votes by Round						
Burns (SNP)	**1470**						
Collins (Grn)	171	217.8	228.1	253.5	256.2		
Craig (SFP)	64	69.9					
Jackson (LD)	106	110.5	117.5				
Johnson (SNP)	323	727.1	738.0	745.7	747.9	885.6	**975.8**
McDonald (Lab)	897	922.1	929.2	**970.6**			
Stephen (Con)	795	797.8	818.1	841.8	846.7	871.1	

we remove the ballots $\mathcal{B}(\mathcal{W})$ from the election, creating the modified profile P', then $A \in W(P', k)$.

The primary reason we define IWVB by focusing on the single candidate A is that if $\mathcal{B}(\mathcal{W})$ is large enough then we should expect some candidates from $\mathcal{W}$ to lose their seats. Thus, IWVB simply insists that winning candidates not ranked on these ballots retain their seats and continue to represent their constituencies. In a setting where proportional representation is the goal, this criterion may be too stringent; as we show below, all of the six methods we study violate this criterion. To preserve some notion of proportionality, changes to the ballot set involving winning candidates may necessitate changes to the winner set which affect candidates not in $\mathcal{W}$. To address this possible objection to IWVB, we define a weaker version, IWVB*, which states what should occur if none of the candidates in $\mathcal{W}$ lose their seats when the ballots from $\mathcal{B}(\mathcal{W})$ are removed. Unlike IWVB, CC OM and CC PM satisfy this criterion, although these methods are not usually viewed as strictly proportional in the sense of methods like STV.

Definition 3. Let $\mathcal{W}$ be a subset of winning candidates, $1 \leq |\mathcal{W}| < k$, and $\mathcal{B}(\mathcal{W})$ be a set of ballots such that only candidates from $\mathcal{W}$ are ranked. A voting method satisfies **IWVB*** if whenever we remove the ballots in $\mathcal{B}(\mathcal{W})$ (creating a new profile P') and $\mathcal{W} \subseteq W(P', k)$, then $W(P, k) = W(P', k)$.

That is, if we remove truncated ballots which rank only a subset of winners, $\mathcal{W}$, and the removal of these ballots does not cause any of the candidates in $\mathcal{W}$ to lose their seats, then the overall winning committee should not change.

We illustrate a violation of IWVB* (and hence also a violation of IWVB) with another election from the Scottish elections database.

Example 2. The 2022 election in the eighth ward of the North Ayrshire council area unfolded as shown in the top of Table 2. There are seven candidates from six different parties: Conservative (Con), Green (Grn), Labor (Lab), Liberal Democrats (LD), Scottish Family (SFP), and the SNP. If we remove 199 bullet votes for Nairn McDonald, the Labor candidate, then the election unfolds as shown in the bottom of Table 2. We might expect McDonald to lose a seat, but instead removing this unilateral support for McDonald causes Susan Johnson to replace Angela Stephen in the winner set.

We make two observations about this example. First, with Scottish STV, a violation of a voter bloc criterion can manifest in different ways. In Example 1, the actual election and the modified election unfold in the same sequential order until the final round. That is, removing bullet votes for Holden did not change the order in which the election unfolded in intermediate rounds: in either election Knapp is elected in the first round, Holden is eliminated in the second round, etc. By contrast, in Example 2 the removal of votes for McDonald causes the election to unfold in a different order, so that McDonald is not elected until the fourth round in the modified election. This change in the order of how the election unfolds ultimately leads to a violation of IWVB*. Second, the party dynamics in this election demonstrate why IWVB might be normatively desirable. It seems strange that the SNP can double its representation in the winning committee if voters who care only about one Labor candidate decide not to vote. The voters who cast bullet votes for McDonald are presumably indifferent between all other candidates; why should their votes determine if a Conservative or SNP candidate earns the third seat? If anything, the presence of these left-leaning voters should support the election of a left-leaning candidate from the SNP, but instead these Labor voters are necessary for a Conservative to win the final seat.

In contrast to the ILVB criterion, the winner bloc criteria make sense only in the multiwinner setting, and we feel they have the most normative weight when the goal of the voting rule is to achieve proportional representation.

When introducing new fairness criteria, our first question should be: do any reasonable voting methods satisfy these criteria? In our case the answer is Yes, since any positional scoring rule such as k-Borda or k-plurality satisfies both ILVB and IWVB (the reason is that a candidate's score does not decrease when removing ballots on which that candidate is not ranked), but these methods do not aim to achieve proportional representation. As we discuss below, depending on what is meant by "proportional", there are proportional methods which satisfy ILVB and IWVB*, but we are unable to find a ranked-choice proportional method which satisfies IWVB. Because positional scoring rules satisfy ILVB and IWVB, these are distinguished from more classical criteria such as Arrow's independence of irrelevant alternatives, which no reasonable method satisfies.

2.4. Connections to other fairness criteria

In this section we briefly outline connections between our voter bloc criteria and other classical fairness criteria in the social choice literature. Because violations of our criteria involve removing ballots, arguably the most closely-related fairness criteria from the voting literature are those that address how electoral outcomes should change as the size of the electorate varies. Violations of such criteria are commonly referred to as *paradoxes of variable electorates* (Diss et al., 2024; Lepelley and Merlin, 2001), and include outcomes such as the no-show paradox and the positive abstention paradox.

The ILVB criterion, in particular, is most closely related to fairness criteria involving the notions of participation monotonicity and no-show paradoxes (Fishburn and Brams, 1983; Moulin, 1988). The term "no-show paradox" has different meanings in different strands of the literature, and almost all previous literature focuses on the single-winner setting. For our discussion, we adapt the definition of no-show paradox which appears in McCune and Graham-Squire (2024), and say that an election demonstrates a no-show paradox if there exists a set of ballots $\mathcal{B}$ with a losing candidate L ranked higher than a winning candidate W, and if we remove the ballots from $\mathcal{B}$ to create the modified profile P', then $L \in W(P', k)$ and $W \notin W(P', k)$. That is, an election demonstrates a no-show paradox if there exists a set of voters who obtain a preferable electoral outcome by not casting their ballots. Under this definition, some violations of ILVB are examples of no-show paradoxes. For example, suppose we remove ballots of the form $A > B$ where A and B are both losing candidates, and B wins a seat as a result; such an outcome is an ILVB violation as well as an example of a no-show paradox. These outcomes are very rare but they do occur. For example, in the 2022 election in the fifth ward of the City of Edinburgh council area candidate Malcolm Wood does not win a seat, but if we remove two bullet votes for Wood then he does win a seat (McCune and Graham-Squire, 2024).

Our voter bloc criteria are also related to various notions of the spoiler effect, such as Arrow's classical independence of irrelevant alternatives (IIA) criterion (Arrow, 1951). Our criteria could be interpreted as multiwinner analogues of the IIA criterion with respect to truncated ballots in the following sense: Suppose W is a winning candidate in the original profile P and L is a losing candidate, and there exists a set of ballots $\mathcal{B}$ ranking only losing candidates such that removal of these ballots creates a profile P' where L replaces W in the resulting winner set. In moving from P to P', the relationship between W and L either remains unchanged (if the ballots from $\mathcal{B}$ do not rank L) or W becomes stronger relative to L (if the ballots do rank L). In either case, the candidates in $\mathcal{B}$ are acting as "irrelevant alternatives" in some sense when it comes to deciding whether L or W is in the winner set. The aforementioned ILVB anomaly in the Alaska election is an example of such an IIA analogue, because removing some bullet votes for the "irrelevant" candidate (Palin) changes the winner from Peltola to Begich. Interestingly, the Alaska election also demonstrates another ILVB anomaly: removing thousands of Palin-Begich votes would also cause Begich to be the new winner, demonstrating a no-show paradox. In other words, the Alaska election demonstrates two different ILVB anomalies, one which is a variation of IIA and the other which is a no-show paradox.

The ILVB criterion is perhaps more connected to the spoiler effect than the winner bloc criteria, because typically an election is said to demonstrate a spoiler effect if there exists a losing candidate L (or perhaps a subset of losing candidates $\mathcal{L}$) such that when we remove L from all ballots (or remove all candidates from $\mathcal{L}$), the winner of the election changes. The ILVB criterion has a similar flavor, but focuses on removing only ballots on which losing candidates are ranked. Thus, it is straightforward for a method like Scottish STV to construct elections which demonstrate a spoiler effect but not an ILVB violation, and vice versa.

One final criterion from the literature related to our winner bloc criteria is Young's reinforcement axiom (Young, 1974, 1975), sometimes referred to as *consistency* (Elkind et al., 2017). Informally, a voting method is consistent if whenever P and P' are profiles such that $W(P, k) = W(P', k)$ then if we create a new profile P^* by combining the ballots from P and P', we have $W(P^*, k) = W(P, k)$. That is, if W is a winning candidate in a profile P and also in P', then W still wins when we combine the ballots from the two profiles. Our winner bloc criteria are similar in spirit to a "subtractive" version of this criterion, where if A wins in an overall profile and we remove some ballots which have nothing to do with A then A should still be in the resulting sub-profile.

Table 3
A profile illustrating a worst-case outcome for Scottish or Meek STV with respect to ILVB. If there are k seats, create k copies of the first five columns, with each copy containing a disjoint set of candidates.

7	9	12	13	2	7	…
A_1	A_1	B_1	C_1	B_1	A_2	…
B_1	C_1	C_1	A_1		B_2	…
C_1	B_1	A_1	B_1		C_2	…

3. Worst case analyses

In this section, we show that for each method except CC we can change the entire winning committee by removing partial ballots with only losing candidates ranked on them. The same is true for every method including CC when we remove partial ballots with only winning candidates ranked on them. Thus, the worst-case for both ILVB and IWVB violations is as bad as possible, with the exception of ILVB with CC. We also show that CC satisfies the IWVB* criterion, so the worst-case analysis does not apply to CC in this case. Let V denote the number of voters in an election. Throughout this section, we use a quota of $q = \frac{V}{k+1}$ for convenience. While some of the methods we described in Section 2 use slightly different quotas such as $\lfloor \frac{V}{k+1} \rfloor + 1$, the differences are vanishingly small for large V and have no impact on the theoretical results presented here.

A table (Table 16) summarizing our results from this section can be found in the appendix.

The preference profiles used to demonstrate worst-case outcomes are extreme, and unlikely to be observed in practice. In Section 4, we explore the outcomes found in real-world elections.

Proposition 3.1 (*ILVB - STV and EAR*)**.** *Let $k > 0$. For Scottish STV, Meek STV, and both versions of EAR, there exists a profile P (which depends on the voting method) and a set of ballots from P ranking only losing candidates such that when we remove these ballots, creating the modified profile P', $W(P, k) \cap W(P', k) = \emptyset$.*

Proof. Fix $k > 0$. For Scottish and Meek STV, consider the profile outlined in Table 3. (This profile is inspired by Example 3.5.2.1 in Felsenthal, 2012.) The profile contains $3k$ candidates which we label $A_1, \ldots, A_k, B_1, \ldots, B_k, C_1, \ldots, C_k$. For each $1 \le i \le k$, let there be 7 ballots of the form $A_i > B_i > C_i$, 9 ballots of the form $A_i > C_i > B_i$, 12 ballots of the form $B_i > C_i > A_i$, 13 ballots of the form $C_i > A_i > B_i$, and 2 bullet votes for B_i. No candidate achieves quota initially, and under either form of STV the C_i candidates are all eliminated, resulting in the winning committee $\{A_1, \ldots, A_k\}$. If we remove the 2 bullet votes for each B_i, then the B_i candidates are eliminated, resulting in the winning committee $\{C_1, \ldots, C_k\}$.

For EAR OM and PM, consider the profile containing $3k + 1$ candidates outlined in Table 4. First, when the bullet votes for C are present, the quota is $q = \frac{(10k+8)k+3k}{k+1} = \frac{10k^2+11k}{k+1}$. In the first round, the A_i candidates have 8 votes and the B_i candidates have $10k$ votes. Note $q = \frac{10k^2+11k}{k+1} > \frac{10k^2+10k}{k+1} = 10k$ and therefore no candidate initially achieves quota. In the second round, all A_i candidates have $10k+3k+8$ or $10k+8$ votes for the optimistic and pessimistic models of EAR, respectively. In both cases, all A_i candidates achieve quota and have more votes than the B_i candidates, so the winning committee is $\{A_1, \ldots, A_k\}$.

If we remove the $3k$ bullet votes for C, then the quota is $q = \frac{(10k+8)k}{k+1} = \frac{10k^2+8k}{k+1} < \frac{10k^2+10k}{k+1} = 10k$. Thus, the B_i candidates achieve quota in the first round and the winning committee is $\{B_1, \ldots, B_k\}$. ☐

While the worst-case ILVB violations for STV and EAR are as extreme as possible, by contrast the Chamberlin Courant rule satisfies the criterion.

Table 4
An example of an election where the winner set completely changes according to the Expanding Approvals Rule when throwing out the C loser ballots.

8	$10k$	…	8	$10k$	$3k$
A_1	B_1	…	A_k	B_k	C
D_1	A_1		D_k	A_k	

Proposition 3.2. *CC OM and CC PM satisfy ILVB.*

Proof. Let $W(P,k)$ be the winner set under either version of CC. Let $\mathcal{B}(\mathcal{L})$ be a set of ballots which rank only losing candidates. If we remove the ballots from $\mathcal{B}(\mathcal{L})$, then the CC score for the set $W(P,k)$ may decrease (in the OM model, e.g.) but any set of candidates not containing one of the losing candidates from $\mathcal{B}(\mathcal{L})$ will lose the same amount of points as $W(P,k)$. By definition of the CC rule, any set of candidates containing one of the losing candidates ranked on the ballots in $\mathcal{B}(\mathcal{L})$ will lose even more points than $W(P,k)$. Therefore, in the modified profile with the ballots in $\mathcal{B}(\mathcal{L})$ removed, the set $W(P,k)$ still has the maximal CC score and remains the winner set. □

We note that any CC rule which uses a reasonable model to process partial ballots also satisfies ILVB.

We now prove an analogous proposition for IWVB. In this case, CC also produces worst-case outcomes.

Proposition 3.3 (*IWVB*)**.** *Let $k > 0$. For all six election methods we examine, there exists a profile P (which depends on the voting method) and a set of ballots from P which rank only a single winning candidate such that when we remove these ballots, creating the modified profile P', $W(P,k) \cap W(P',k) = \emptyset$.*

Proof. Fix $k > 0$. For Scottish and Meek STV, consider the profile with $2k$ candidates outlined in Table 5. The quota is $20k-12$, which also equals the number of first-place votes for A_1. No other candidate achieves quota initially, and thus A_1 earns the first seat but has no surplus to transfer. B_1 has the fewest first place votes and is eliminated, transferring two votes to A_i for $2 \le i \le k$. As a result, each of the remaining B_i candidates have fewer votes than any of the A_j candidates, all B_i candidates are eliminated, and the winning committee is $\{A_1,\dots,A_k\}$.

However, if we remove the $14k-12$ bullet votes for A_1, then A_1 has only $6k$ votes and is eliminated in the first round. Each B_i candidate receives a positive vote transfer from A_1 while the remaining A_i candidates do not. Thus each of the A_i candidates are eliminated and the winning committee is $\{B_1,\dots,B_k\}$.

For both EAR models, consider the profile in Table 6. The quota is $q = 20k+10$, so candidate A is the only candidate to earn quota initially and is given the first seat. Furthermore, because quota is so large, no other candidates make quota.

When considering the second place votes, the B_i candidates have at least $10+20k$ votes and make quota. They also have more votes than any other candidates, so the B_i candidates win the remaining $k-1$ seats and the winning committee is $\{A, B_1,\dots,B_{k-1}\}$. (In the optimistic model, the C_i candidates make quota with the extra support from the A bullet votes, but the B_i candidates also get that support and have more votes.)

If we remove the bullet votes for A, the quota is $q = \frac{10(2k+1)(k-1)}{k+1}$, which is less than $20k$. Thus, the C_i candidates win the k seats and the winning committee changes to $\{C_1,\dots,C_k\}$.

For CC, consider the profile in Table 7. For $1 < i < k$, there are two ballots of the form $A_i > B_i$, two of the form $B_i > A_i$, two of the form $A_i > B_{i+1}$, and two of the form $B_{i+1} > A_i$. For $i \in \{1,k\}$ the ballots are as shown in the table. Let $\mathcal{A} = \{A_1,\dots,A_k\}$ and $\mathcal{B} = \{B_1,\dots,B_k\}$. For CC OM, the score for $\mathcal{A}$ is

$$(2k-1)(5+4(k-1)) + (2k-2)(4k) = 16k^2 - 10k - 1$$

and the score for $\mathcal{B}$ is

$$(2k-1)(4k) + (2k-2)(3+2+4(k-1)) = 16k^2 - 10k - 2.$$

Thus, $\mathcal{A}$ will have more CC OM points than $\mathcal{B}$ in this case. In the pessimistic model, the score for $\mathcal{B}$ goes down by $3(2k-2)$ to $16k^2-16k+4$ so $\mathcal{A}$ still beats $\mathcal{B}$ in that situation. Since $\mathcal{A}$ has the property that every voter has a candidate in the subset ranked in their top two rankings, swapping out any A_i with a B_j will result in a lower CC score. Specifically, exchanging an A_i with a B_j, will result in some columns of votes that have A_i now receiving no points (or only $2k-3$ points in OM), leading to a reduced CC score. Exchanging any B_i for A_j in $\mathcal{B}$ similarly results in a reduced CC score. It follows that $\mathcal{A}$ wins both CC OM and CC PM for the profile in Table 7.

Now consider the profile in Table 7 with the first column (3 bullet votes for A_1) removed, and call this profile P'. For P', the score for $\mathcal{A}$ is $(2k-1)(2+4(k-1)) + (2k-2)(4k) = 16k^2-16k+2$ and the score for $\mathcal{B}$ is $16k^2-16k+4$ (equal to the PM score with bullet votes), so $\mathcal{B}$ has more points than $\mathcal{A}$. As with the original profile, swapping out some candidates from $\mathcal{A}$ or $\mathcal{B}$ will similarly result in lower scores for the modified subsets, thus $\mathcal{B}$ wins both CC OM and CC PM for profile P'. □

To conclude this section, we provide the worst-case analysis for violations of IWVB* for the STV and EAR methods, and show that both forms of CC satisfy the criterion.

Proposition 3.4 (*IWVB* - STV and EAR*)**.** *Let $k > 0$. For Scottish STV, Meek STV, and both models of EAR, there exists a profile P (which depends on the voting method) and a set of ballots from P which rank only a single winning candidate A such that when we remove these ballots, creating the modified profile P', $W(P,k) \cap W(P',k) = \{A\}$.*

Proof. For Scottish and Meek STV, we first give the proof in full generality, then give an illustrative example at the end. Fix $k > 0$ and consider the profile P in Table 8.

For large enough b (relative to c), A will clearly win the first seat. After A is elected, either the B_i candidates have more votes and the C_i candidates are removed, or vice versa, depending on the relative values of a, b and c. We must determine how many votes each B_i candidate receives. The quota is $\frac{a+(k-1)(b+c)}{k+1}$, so each B_i candidate receives

$$\frac{a+(k-1)b - \frac{a+(k-1)(b+c)}{k+1}}{a+(k-1)b}\, b.$$

When $a = 0$, the B_i candidates will receive the fewest votes possible with $\frac{kb-c}{k+1}$ votes. On the other hand, if we take the limit as $a \to \infty$, the share of B votes approaches $\frac{kb}{k+1}$.

Therefore, if $\frac{kb}{k+1} > c > \frac{kb-c}{k+1}$, there is a value of a where A transfers enough ballots to the B_i candidates that they also win seats, but if the a ballots are removed, A consumes almost all the votes and the C_i candidates win. One example that works for Scottish and Meek STV is in Table 9 with $k = 3$, $a = 1000$, $b = 20$, and $c = 13$.

For EAR PM, consider Table 10. When there are no bullet votes for A ($a = 0$), quota is $q = \frac{(10k+20)(k-1)}{k+1} = 10\frac{(k+2)(k-1)}{k+1}$. A makes quota, so A wins a seat and the ballots with A ranked first have their weights adjusted from 1 to $\epsilon = \frac{10k(k-1)-q}{10k(k-1)} = \frac{k^2-2}{k(k+1)}$. No other candidates make quota so the second place votes are considered. The B_i candidates have 20 votes, which is short of quota, but the C_i candidates have $10+10\epsilon k$. $10+10\epsilon k = 10 + 10\frac{k^2-2}{k+1} = 10\frac{k+1+k^2-2}{k+1} = 10\frac{k^2+k-1}{k+1} > 10\frac{(k+2)(k-1)}{k+1} = q$, so the C_i candidates all make quota and are elected, filling all k seats.

Now consider what happens as a approaches infinity. As $a \to \infty$, $q \to \infty$, but A still makes quota in the first round. Once A is removed and votes transferred, though, the ballots with A ranked first have weights that approach $\frac{k}{k+1}$, and no other candidates ever make quota. After considering the second and then third place votes, the B_i candidates have $20+10k\frac{k}{k+1}$ votes, which is greater than the $10+10k\frac{k}{k+1}$ votes the C_i candidates have. Therefore, the B candidates win seats. □

Table 5
An example of a profile where the winner set completely changes when throwing out the A_1 bullet votes using Scottish STV or Meek STV.

$14k-12$	$4k+2$	2	…	2	$6k+2$	2	…	2	$10k$	$10k$	…	$10k$	$10k$
A_1	A_1	A_1	…	A_1	B_1	B_1	…	B_1	A_2	B_2	…	A_k	B_k
	B_1	B_2	…	B_k		A_2	…	A_k					

Table 6
An example of a profile where the winner set completely changes under EAR when throwing out the A bullet ballots.

$20k+20$	10	…	10	$20k$	…	$20k$	$20k$
A	B_1	…	B_{k-1}	C_1	…	C_{k-1}	C_k
	D_1		D_{k-1}	B_1	…	B_{k-1}	E

Finally, for EAR OM, consider Table 11. When $a = 0$, there are $80(k-1)$ votes, so quota is $q = 80\frac{k-1}{k+1}$. A makes quota immediately, because $80(k-1) > q$, and those 30 votes are all reduced to $50 - \frac{80}{k+1}$ votes. However, some algebra shows that $50 - \frac{80}{k+1} = \frac{50k-30}{k+1} < \frac{80k-80}{k+1} = q$ for $k > 1$, so the B_i candidates do not make quota in the second round. In the third round, the C_i candidates now have the support from all remaining ballots and win the last $k-1$ seats.

Again, as a approaches infinity, the quota goes to infinity. A makes quota in the first round, and the 50 votes in columns 2 through k are reduced to $50\frac{k}{k+1}$. The A bullet votes now support all candidates and all candidates make quota. But the B_i candidates have an additional $50\frac{k}{k+1}$ votes while the C_i (and D_i) candidates only have 30 extra votes. For $k \geq 2$, $50\frac{k}{k+1} > 30$, so the B_i candidates win seats.

Proposition 3.5. *CC OM and CC PM satisfy IWVB*.*

Proof. Let W be the winner set under a given version of CC with the ballot profile P. Let $\mathcal{B}$ be a set of ballots which rank only candidates in $\mathcal{W} \subset W$. We claim that if we remove the ballots in $\mathcal{B}$ from P, creating a modified profile P' which has the (potentially different) winner set W' with $\mathcal{W} \subset W'$, then $W = W'$.

Suppose not, so $W \neq W'$. Because W is the winning committee under P, the CC score for W is greater than the score for W' under P. When we remove the ballots in $\mathcal{B}$, this subtracts the same number of points from both W and W' because $\mathcal{W}$ is a subset of both W and W'. Therefore, under P', the score for W must still be greater than the score for W'. Therefore, W' cannot be the winning committee under P', a contradiction, so we conclude that $W = W'$. □

4. Empirical results

In this section we provide empirical results using the Scottish elections dataset. For each multiwinner election we performed a heuristic search for a set of ballots which would demonstrate a violation of ILVB, IWVB, or IWVB* under each voting method. One methodological limitation of any such empirical analysis is that we cannot know how the ballot data might change if a method other than Scottish STV were used. If a method like CC OM were used instead then it is possible that candidates, parties, and voters would behave differently (for example, more moderate candidates run for election or voters cast longer ballots), producing different ballot data. Since all six of our voting methods attempt to achieve proportional representation, we hope the ballot data would look similar across voting methods.

Before presenting the results, we describe the methodology used to find violations of the criteria.

4.1. Methodology for finding violations

As the number of ballots cast in an election grows, the number of different sets of ballots that could be removed to potentially cause a violation grows exponentially, and thus any kind of brute force algorithm is not feasible. Therefore, we use a heuristic search that removes only a small subset of all possible ballot combinations to test for violations of our voter bloc criteria. The procedure for searching for losing voter bloc violations is slightly different than winning voter bloc violations, but they both begin by finding the winner set W and the loser set L of the full election. We give high level descriptions of these algorithms below. The code we use, along with all violations found under each method (including a record of the ballots removed) can be found at https://github.com/MattJonesMath/Irrelevant_Voters_Project. We note that some of the code we use for the EAR method was provided by Tuva Bardal and Jannik Peters.

Because we use heuristic algorithms, we cannot guarantee that we found all fairness violations for any of the voting methods. With that said, we implemented code at different levels of granularity and stopped when we found only marginal changes in the number of violations. Due to this, we believe we have found almost all of the ILVB, ILWB, and ILWB* violations present in the Scottish data. Ideally we could find necessary and sufficient conditions that a preference profile must satisfy to produce a violation for a given method, and our code could simply check each preference profile against such conditions. Given the challenge of finding such conditions for various fairness criteria in the three-candidate single-winner case (Felsenthal, 2019; Graham-Squire, 2024; Miller, 2017) and the absence of any known conditions in the multiwinner case, we doubt such conditions are forthcoming in our setting.

Additionally, these algorithms are agnostic to the election method used, and all use the same basic principle of identifying all eligible ballots, removing some fraction of those ballots, and then running the election again to see if the winner set changed in a desirable way. Therefore, we have no reason to believe that our results are biased in favor of finding violations with one election method over another, although we cannot rule out this possibility completely.

4.1.1. ILVB violation search

Select a candidate $B \in L$. We try to remove loser ballots so that B wins a seat. It is very unlikely (although not impossible) that removing a ballot that ranks B will help B win a seat, so identify $\mathcal{B}$, the set of all ballots that rank only losing candidates that are not B.

Removing loser ballots has no impact on the votes that winning candidates receive. Instead, removing these ballots can influence the winner set by changing the quota. We want to test many different quotas, because violations are not monotone with respect to the removal of ballots. That is, it is possible that removing some fraction of the ballots in $\mathcal{B}$ causes a violation but if we remove all the ballots in $\mathcal{B}$, we do not observe a violation. Therefore, we try removing σ_ℓ different fractions $(\frac{1}{\sigma_\ell}, \frac{2}{\sigma_\ell}, \dots, 1)$ of $\mathcal{B}$, each time rerunning the election to see if B has won a seat. For neatness, if this requires removing a fractional ballot, we round down the number of ballots removed to the next largest integer.

Repeat this process for all $B \in L$.

4.1.2. IWVB and IWVB* violation search

Select a pair of candidates $A \in W$ and $B \in L$. We try to remove winner ballots so B takes A's seat. Because we are removing winner ballots now, these ballots can influence the election by changing the quota, but also by preventing surplus votes from being transferred from one candidate to another. Therefore, in this search, we consider how many votes each winning candidate who is not A may send to A and B if they are elected. For each candidate $C_* \in W \setminus \{A\}$, count the number

Table 7
An example of a profile where the winner set completely changes according to both CC rules (OM and PM) when throwing out winner ballots. The middle block of ballots are copied for $i = 2, 3, \dots k-1$.

3	1	2	1	2	2	2	2	2	2	2	2	2
A_1	A_1	B_1	A_1	B_2	A_i	B_i	A_i	B_{i+1}	A_k	B_k	A_k	B_1
	B_1	A_1	B_2	A_1	B_i	A_i	B_{i+1}	A_i	B_k	A_k	B_1	A_k

Table 8
A generic set of ballots that demonstrate the worst-case scenario for STV. A is always a winner. For sufficiently large a, when we throw out the A bullet votes, the rest of the winner set will change from Bs to Cs if $\frac{kb}{k+1} > c > \frac{kb-c}{k+1}$.

a	b	…	b	c	…	c
A	A	…	A	C_1	…	C_{k-1}
	B_1	…	B_{k-1}			

Table 9
An election demonstrating a worst-case scenario for both versions of STV. When the 1000 bullet votes for A are removed, the winner set changes from $\{A, B_1, B_2\}$ to $\{A, C_1, C_2\}$.

1000	20	20	13	13
A	A	A	C_1	C_2
	B_1	B_2		

Table 10
An example of a set of ballots for EAR PM. A is always a winner. For sufficiently large a, when we throw out the A bullet votes, the rest of the winner set will change from Cs to Bs.

a	$10k$	…	$10k$	10	…	10	10	…	10
A	A	…	A	B_1	…	B_{k-1}	C_1	…	C_{k-1}
	C_1	…	C_{k-1}				B_1	…	B_{k-1}
	B_1	…	B_{k-1}						

Table 11
An example of a set of ballots for EAR OM. A is always a winner. For sufficiently large a, when we throw out the A bullet votes, the rest of the winner set will change from Bs to Cs.

a	50	…	50	30	…	30
A	A	…	A	C_1	…	C_{k-1}
	B_1	…	B_{k-1}	D_1	…	D_{k-1}
	C_1	…	C_{k-1}	E_1	…	E_{k-1}

of ballots that rank C_* above both A and B but rank A above B. These are ballots that could potentially be transferred to A if C_* wins a seat. Then also count the number of ballots that rank C_* above both A and B but rank B above A, which are ballots that could potentially transfer to B. Rank the candidates in $W \setminus \{A\}$ by the difference of these two counts so we have $C_1, C_2, \dots, C_{k-1}$. C_1 will have many votes that could transfer to A and C_{k-1} will have many votes that could transfer to B.

We begin by taking $\mathcal{B}$ to be only bullet votes for C_1. Like above, we want to try many different quota values, so remove σ_w different fractions of $\mathcal{B}$ and try running the election.

Continue by letting $\mathcal{B}$ be the ballots that only rank C_1 and C_2, then C_1 through C_3, and so on until $\mathcal{B}$ is all ballots that rank all the winning candidates except A, each time removing σ_w fractions of $\mathcal{B}$.

Repeat this process for all pairs of $A \in W$ and $B \in L$.

When searching for violations for the alternate definition IWVB* (Def Definition 3), we also confirm that all candidates ranked by a ballot in $\mathcal{B}$ still win a seat.

4.1.3. Party dynamics

When searching for violations that result in a change in party control, we modify the above algorithms as follows. Select a pair of candidates $A \in W$ and $B \in L$ (even for the ILVB search). When considering which ballots to remove, restrict to only ballots that do not rank any candidates from the parties of A and B. If a potential violation has been found where A loses their seat and B gains a seat, confirm that A's party has one fewer seat than before, and B's party has one additional seat. This is to avoid the case where many winning ballots are removed and many seats change winners, which can make it appear that a seat has changed parties even if it has not.

4.2. Heuristic algorithm validation

It is important to emphasize that this search is not guaranteed to find all possible violations in the data. Rather, the values presented in Tables 12–14 are lower bounds on the true numbers of violations that are present. Because of this limitation, we ran a number of different tests to confirm that we missed no easily detectable violations.

First, on a set of 76 small elections, we exhaustively searched every possible set of removed ballots, using brute force to establish the "ground truth", and confirmed that our ILVB algorithm found every violation in those elections for several election methods. Unfortunately, repeating this for the winner bloc criteria, other election methods, or larger elections is computationally infeasible.

Second, we developed an earlier heuristic search that was very fine-grained, removing many fractions of bullet votes, and the violation searches we describe above found every single one of the violations found by the earlier heuristic.

Finally, increasing the σ_ℓ and σ_w parameters increases the computation time but can also increase the number of violations found. We gradually increased these parameters until the number of new violations did not justify the increased computation time. This method of identifying diminishing returns is similar to techniques used in other research areas, such as Markov chain Monte Carlo (MCMC) (Roy, 2020), in which heuristic methods are necessary. For example, in Herschlag et al. (2020) MCMC is used to create a representative ensemble of districting plans for the state of North Carolina. It is necessary to create a large number of districting plans to accurately represent the underlying space, but it is not clear how large that number must be. One way to validate that the MCMC-generated distribution matches the empirical reality is to consider changes in the distribution as the number of elections in the ensemble increases. Once the distribution ceases to change with increased numbers of elections, that indicates the ensemble has converged to likely be a fair representation of the underlying space. Similarly, if changing initial parameters of the MCMC algorithm result in a virtually identical ensemble of elections, one can be confident in the quality of the heuristic algorithm used to create the distribution.

We also note that heuristic methods are common in research on real-world monotonicity anomalies (see Graham-Squire and McCune (2025), Graham-Squire and Zayatz (2021), McCune and Graham-Squire (2024), for example), since brute-force methods are computationally intractable. To date, prior heuristic results of similar published empirical voting research have not been found to have missed any monotonicity anomalies.

Heuristic algorithms have been employed to investigate other features of STV. For example, Blom et al. (2020, 2016), Cary (2011), Magrino et al. (2011) develop algorithms for estimating margins of victory in STV elections. The use of heuristics is necessary because of the complexity of the STV algorithm.

For all these reasons, we believe that the heuristic methods we employed are appropriate and the values reported in Tables 12–14 are fairly close to the true values.

Table 12
The number of violations found in the Scottish election data using parameters $\sigma_\ell = 10$ and $\sigma_w = 3$ in our violation search. The first row shows how many elections we found that violate ILVB (Definition 1), the second row counts the elections that violate IWVB (Definition 2), and the third row counts elections that violate the alternative IWVB* (Definition 3).

	Scottish	Meek	EAR OM	EAR PM	CC OM	CC PM
ILVB Violations	40	19	71	54	0	0
IWVB Violations	108	103	225	198	23	28
IWVB* Violations	103	102	219	181	0	0

4.3. Violations in Scottish elections

We searched for violations in the Scottish election data with the five methods from Section 2. We found 573 different violations (not counting the IWVB* violations, which are special cases of IWVB violations) in 430 elections. Our algorithms could find no violations under any of our voting methods in 640 elections (≈ 60%). CC OM was least susceptible to violations, followed closely by CC PM. As expected, we found no ILVB or IWVB* violations in the Scottish data using a CC rule. Meek STV appears to be the sequential process with the fewest violations. The expanding approvals rule had the most of every type of violation, specifically the optimistic model. The full results for $\sigma_\ell = 10$ and $\sigma_w = 3$ are shown in Table 12. Larger values of σ_ℓ and σ_w yield only marginal improvements for large increases in computational cost.

By examining which elections have violations for specific methods, we confirm that there are no violations with one method that guarantee a violation with another method. It seems likely that there are no relationships of this kind, although we cannot say that for certain since we did not conduct an exhaustive search for all possible violations. The closest we come in this data is that IWVB violations under CC OM only occur in elections in which there is also an anomaly with either CC PM or EAR PM.

Another interesting observation concerns EAR. The original method presented in Aziz and Lee (2020) is the optimistic model, in which truncated ballots support all candidates, while the pessimistic model we implement is suggested in Remark 3 of Aziz and Lee (2020). Note that EAR PM outperforms EAR OM for all three violation types in Table 12. We suggest that this is because the real election data from Scotland is highly truncated and EAR OM does not handle extreme truncation well. To be more specific, extreme truncation "washes out" EAR OM in a way it does not with the other methods. Consider an election with 4 seats where 20% of the ballots are bullet votes. If no one makes quota in the first round then second place rankings count and, in the optimistic model, those 20% of voters approve all candidates. Every single candidate makes quota. This is not an immediate problem because the candidate with the most votes (call them A) probably deserves a seat. But then every bullet voter helps "buy" that candidate's seat, even if they cast a bullet vote for someone else. Thus, most of the burden of electing A goes to voters who did not truly support A, and A voters pay a reduced price to elect their candidate. This issue compounds, because now the A bullet voters (who paid a relatively low price to elect A and now have more remaining electoral power than we would like) are supporting every candidate and lowering the cost to elect everyone after that, further "washing out" the election. The pessimistic model of EAR avoids this concern, and we see fewer anomalies as a result.

Finally, for about 14% of found anomalies, the violation requires throwing out only a fraction of all possible ballots. When $\sigma_\ell = \sigma_w = 1$ (where we throw out all of $\mathcal{B}$ and do not try different quota values), we find about 10% fewer violations, as can be seen by comparing Tables 12 and 13. EAR seems to produce the bulk of the non-monotone violations.

4.4. Party dynamics

In slightly more than half of elections with a violation, a seat can change parties without discarding any ballots that support either party. Those results are shown in Table 14.

Table 13
The number of violations found in the Scottish election data when $\sigma_\ell = \sigma_w = 1$. Compare to Table 12 to see how many violations are non-monotone and require keeping some ballots and removing others.

	Scottish	Meek	EAR OM	EAR PM	CC OM	CC PM
ILVB Violations	38	19	66	48	0	0
IWVB Violations	93	92	196	163	21	25
IWVB* Violations	88	91	188	142	0	0

Table 14
The number of violations found in the Scottish election data where at least one seat changes party when $\sigma_\ell = 10$ and $\sigma_w = 3$. Compare to Table 12 to see that slightly more than half of violations can result in changes to the party composition of the final committee.

	Scottish	Meek	EAR OM	EAR PM	CC OM	CC PM
ILVB Party Swaps	32	13	56	42	0	0
IWVB Party Swaps	64	60	145	126	11	15
IWVB* Party Swaps	62	60	142	113	0	0

For the purposes of these results, we treat all independent candidates as members of one party. Therefore, these results represent a conservative estimate of how often these violations can affect the political composition of a council. Two independent candidates could have completely opposing views, but we would not notice that shift with the data we have.

5. Proportionality and voter bloc criteria

Throughout this article we focus on methods that aim to provide proportional representation. As we have seen, with the exception of CC the proportional methods we study fail all of our proposed criteria. Perhaps these results point to a conflict between achieving proportionality and satisfying the irrelevant voter bloc criteria. In this section, we explore such potential tension.

First, we briefly review what is commonly meant by "proportionality". In the social choice literature, proportionality is commonly understood in terms of Dummett's proportionality for solid coalitions (PSC) criterion (Dummett, 1984) and its variants (Woodall, 1994). A set of voters form a *solid coalition* if the candidates can be partitioned into two sets C_1 and C_2 so that every voter in the coalition prefers every candidate in C_1 to every candidate in C_2. Put simply, PSC-type criteria establish a quota $q \in (\frac{V}{k+1}, \frac{V}{k}]$ and then use that quota to determine how much representation is deserved by a solid coalition of voters. The q-PSC criterion states that if a solid coalition achieves a size of at least jq for some $j \in \mathbb{N}$, then at least j candidates supported by the coalition should earn seats, assuming the coalition supports at least j candidates. If the coalition supports fewer than j candidates, then q-PSC requires that all candidates supported by the coalition earn seats. PSC has been referred to as the most important requirement for proportional representation (Tideman, 1995), and is often used as the justification for using some form of STV. (All forms of STV used in this article, as well as EAR, satisfy q-PSC for the Hare quota $q = \frac{V}{k}$. CC does not satisfy q-PSC for any q, but in the Scottish election data, there are no violations for $q = \frac{V}{k}$.) There is a sizable literature which studies PSC and proportionality more broadly, and we do not attempt to survey it here. We refer the reader to Aziz and Lee (2020), Dummett (1984), Brill and Peters (2023), Elkind et al. (2017), Woodall (1994) for deeper discussions of PSC and related criteria.

Practically speaking, in a given election the q-PSC criterion establishes a set of acceptable winning committees from which a q-PSC-compatible voting rule must choose. If a candidate earns more than q first-place votes then any PSC-acceptable winning committee includes that candidate, for example. This requirement illustrates why there is tension between satisfying q-PSC and a criterion involving irrelevant voter blocs. If we remove enough ballots from an election, then q, which

Table 15
Preference profiles illustrating why a q-PSC scoring rule fails ILVB.

333	1	333	332
A	B	C	D
		D	C

1	666	332
A	C	B
	D	

is a function of the number of voters, decreases. Consequently, some solid coalitions which were too small to make claims on seats with the original quota may now be large enough to earn seats with the smaller quota. That is, if we reduce the quota enough, then the resulting set of PSC-acceptable winning committees is a proper subset of the PSC-acceptable committees in the original election. As a result, satisfying PSC may require a new winning committee if the original winning committee is no longer in the set of PSC-acceptable committees after quota is decreased.

To demonstrate the challenge of trying to satisfy both PSC and irrelevant voter bloc criteria, and to illustrate the observations of the previous paragraph, we introduce a new family of PSC-compatible voting methods based on positional scoring. As mentioned in Section 2, all positional scoring rules satisfy our independence of irrelevant voter bloc criteria, and therefore adapting such rules to the PSC setting is a natural test case for trying to reconcile PSC and irrelevant voter bloc criteria. Furthermore, these new voting methods do not occur sequentially in rounds like most classical methods which satisfy PSC, and thus we can analyze our criteria independent of the sequential nature of STV-type voting rules.

A *scoring rule* is a method which uses a scoring vector $(s_1, \dots, s_m)$ with $s_i \geq s_{i+1}$, $s_1 > 0$, and $s_i \geq 0$ to determine the number of points that a ballot contributes to a candidate's score, based on where that candidate is ranked on the ballot. If the ballot ranks a candidate in the ith ranking, then the candidate receives s_i points from the ballot. The k candidates with the largest total scores, summed across all ballots, are the winners. Consider a modification of a scoring rule which is designed to satisfy q-PSC.

Definition 4. Fix $q \in (\frac{V}{k+1}, \frac{V}{k}]$ and a scoring vector $(s_1, \dots, s_m)$. The *q-PSC scoring rule* works as follows. For a given preference profile, determine all candidate subsets of size k which are compatible with q-PSC (this is computationally straightforward; see Aziz and Lee (2022), for example). Assign each of these subsets a score by adding the scores, as determined by $(s_1, \dots, s_m)$, for each candidate in the subset. The subset with the highest score is the winning committee. For simplicity, we assume partial ballots are processed under the pessimistic model, so that candidates left off a ballot receive no points from it.

That is, a q-PSC scoring rule is simply a positional scoring rule which can consider only winning committees compatible with q-PSC. We note that the following analysis can be easily adapted to other models of processing partial ballots.

Suppose we set q just larger than $V/(k+1)$, as in STV. We prove that any q-PSC scoring rule violates ILVB (it is straightforward to show violation of IWVB as well).

Proposition 5.1. *Let $q = \lfloor V/(k+1) \rfloor + 1$. Then any q-PSC scoring rule violates ILVB.*

Proof. Consider the left preference profile in Table 15. Note the profile contains 999 voters and thus $q = 334$. Let $k = 2$. Then q-PSC requires only that C or D earn a seat in the winning committee, since the solid coalition supporting C and D exceeds 1/3 of the voters but is less than 2/3. Thus, we calculate the aggregate scores of any committee of size two containing C or D. If $s_2 > s_1/333$, then the winning committee under a q-PSC scoring rule is $\{C, D\}$. However, if we remove one of the bullet votes for B, the quota changes to $q = 333$ and A is required to earn a seat by q-PSC since the only q-PSC compatible committees in this case are $\{A, C\}$ and $\{A, D\}$. Of these two committees, the q-PSC scoring rule chooses $\{A, C\}$, and we see a violation of ILVB.

If $s_2 \leq s_1/333$, then consider the right preference profile of Table 15. As in the previous election, q-PSC requires only that the winning committee contain C or D. Because s_2 is so small relative to s_1, the winning committee is $\{B, C\}$. However, after we remove the one bullet vote for A the quota decreases to 333, and q-PSC requires that the winner set is $\{C, D\}$. □

Thus, even though positional scoring rules satisfy our criteria, if we attempt to combine these rules with a PSC-type constraint then the resulting methods fail, despite the fact that the q-PSC scoring rules do not occur sequentially like STV. We think this points to a deep tension between proportionality axioms like PSC and the voter bloc criteria. An earlier version of this paper conjectured that there is no voting rule which satisfies q-PSC and also satisfies ILVB or IWVB. It was subsequently shown in Aziz et al. (2024) that there exists a PSC-compliant rule which also satisfies the ILVB criterion; see the authors' discussion of their *Solid Coalition Refinement* rule. Future work could explore which other proportionality axioms are compatible with the voter bloc criteria, and if there exists a PSC-compliant voting rule which satisfies IWVB or IWVB*.

6. Conclusion

The ILVB and IWVB criteria represent an attempt to articulate reasonable fairness criteria regarding truncated ballots in the real-world multiwinner setting where such ballots are very common. We argue that if we remove a set of truncated ballots, only the candidates who are ranked on those ballots should be negatively affected. Our analysis shows that the voting rule of Chamberlin-Courant performs well with respect to these criteria while the expanding approvals rule performs relatively poorly, with STV-type rules falling in between.

Future research could investigate how well other voting methods perform with respect to these criteria. In particular, it would be interesting to investigate the compatibility of ILVB and IWVB with proportionality axioms other than PSC. Similar fairness criteria, both strengthening or weakening our studied criteria, could give greater clarity to the impact of ILVB and IWVB in real-world elections. Other work could provide better algorithms to search the real-world dataset for violations, perhaps increasing the amount of violations found. Finally, the introduction of new criteria invites a discussion about how important such criteria are, especially as compared to previously studied axioms. For example, are ILVB and IWVB as normatively desirable as candidate monotonicity or consistency (Elkind et al., 2017)? We argue the answer is Yes, but others may feel that our criteria are less important.

CRediT authorship contribution statement

Adam Graham-Squire: Writing – review & editing, Validation, Supervision, Project administration, Conceptualization. **Matthew I. Jones:** Writing – review & editing, Writing – original draft, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **David McCune:** Writing – review & editing, Writing – original draft, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization.

Funding sources

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Appendix

In this appendix we provide formal definitions of the voting rules we use, a simple example to illustrate differences between the voting rules, and a table organizing our theoretical results (Table 16).

First, we note that all methods described below can theoretically have ties, which are often resolved through a random choice. While the frequency of ties is extremely low, to avoid the possibility of randomness affecting our results, we do not include any anomalies where a tie occurs, either in the original election (one case) or in an anomaly after ballots are removed (five cases).

Scottish STV. Scottish single transferable vote (Scottish STV) is a multistage election procedure which works as follows. In a given stage, if a candidate has a number of first-place votes in excess of the quota $\lfloor \frac{V}{k+1} \rfloor + 1$, then this candidate is given a seat. Their surplus votes above quota are transferred proportionally to candidates ranked next on this candidate's ballots, where we assume that any candidate who has been previously elected or eliminated no longer appears on any ballot. If two or more candidates achieve quota in a given round, then the candidate with the higher vote total has their surplus transferred first (that is, Scottish STV does not use simultaneous surplus transfers). If no candidate has enough first-place votes to achieve quota, then the candidate with the fewest votes is eliminated from all ballots and their first-place votes are transferred to the candidates ranked second on such ballots. One of the distinguishing features of Scottish STV, in contrast to some other forms of STV, is that once a candidate is elected they cannot receive any future vote transfers. If a tied outcome occurs in any stage, the tie is broken at random. Complete details of the method can be found at https://www.legislation.gov.uk/sdsi/2007/0110714245.

Meek STV. The original description of Meek STV can be found at Meek (1994). The precise algorithmic implementation of Meek's method is extensive and can be found at https://era.ed.ac.uk/bitstream/handle/1842/40863/PalmerJE_2023.pdf?sequence=1&isAllowed=y. We give here a brief overview with emphasis on how Meek and Scottish STV differ, and end with a full algorithmic description of a Meek STV implementation. The initial steps of Meek's implementation are identical to Scottish STV, and they diverge in two ways once the first candidate achieves quota and is elected. The first difference is that under Scottish STV, once a candidate C is elected, C no longer appears on any ballots. With Meek STV, C remains on the ballot and any ballots transferred to C are then redistributed proportionally according to *all* ballots that prefer C, not just the remaining ballots. Thus all future transferred votes flow through C, which changes the distribution of votes as later candidates are elected and eliminated. The second difference is that with Scottish STV, the quota is fixed and does not change in later rounds. With Meek STV, as candidates are elected and eliminated the quota can decrease if truncated ballots are not transferred to any candidate.

A full implementation of Meek STV follows these steps, taken from (Hill et al., 1987):

2.1 At each stage, each candidate is in one of three states, designated as 'elected', 'excluded' and 'hopeful'. At the start every candidate is in the hopeful state.

2.2 At each stage the votes are scanned, and the one vote allowed to each voter may be split into parts that are assigned to the various candidates according to the voter's choices. At the first stage the whole of the vote goes to the first choice — this follows automatically from the operation of rules 2.1 and 2.3.

2.3 Each candidate, x, has an associated weight, w_x, and keeps a proportion w_x of each vote or part of a vote received, while passing on to another candidate (as specified by the voter's choices) a proportion $1 - w_x$. Every hopeful candidate has weight 1, and therefore keeps everything received and passes nothing on. Every excluded candidate has weight 0, and therefore keeps nothing and passes everything on. Elected candidates have weights between 0 and 1, to be calculated by rule 2.5.

2.4 Thus if someone has voted for candidate a as first choice, b as second, c as third, and no more:

- a receives from that voter w_a of a vote
- b receives from that voter $(1 - w_a)w_b$ of a vote
- c receives from that voter $(1 - w_a)(1 - w_b)w_c$ of a vote

A fraction $(1 - w_a)(1 - w_b)(1 - w_c)$ remains and this goes to 'excess'. (Note that if a hopeful candidate appears in the list, all the fractions beyond that point automatically become 0).

2.5 The quota is defined as (total votes − total excess)/(number of seats + 1), and the weights for elected candidates are found such that the total vote remaining with each of them equals the quota. This is done by the convergent iterative scheme specified in rule 2.9.

2.6 The weights having been found, the resulting total votes for each hopeful candidate are examined, and any candidate whose total votes equal or exceed the quota changes state from hopeful to elected (except in the special case where all the hopeful candidates either have zero votes or exactly equal the quota. In this case all those with zero votes are excluded, one other is excluded by a pseudo-random choice and the others are elected).

2.7 If no candidate were elected under rule 2.6, then the hopeful candidate with the fewest votes changes state from hopeful to excluded. Any tie is resolved by a pseudo-random choice.

2.8 If the total number of elected candidates is equal to the number of seats, the election is complete. Otherwise the process is repeated from rule 2.2.

2.9 The convergent iterative scheme is as follows: set w_j equal to 0 for excluded candidates, 1 for hopeful candidates, and their last calculated values w_j^0 for elected candidates. (Immediately after election of any candidate the last calculated value is 1 initially.) Applying rule 2.3, using these weights, let v_j be the total value of votes received by candidate j and let e be the total excess. Using this value for e, calculate the new quota q using rule 2.5. Finally update the weights for elected candidates to values $w_j^1 = w_j^0 / v_j$. Repeat the process of successively updating v_j, e, q and w_j until every fraction q/v_j, for elected candidates, lies within some error value of 1 (we used 0.001 in our code).

Expanding Approvals Rule (EAR) (Aziz and Lee, 2020). This method, proposed by Aziz and Lee, is designed so that voters can express weak linear orders over the candidates, but easily translates to our setting in which a voter's ordering is strict (with the possible exception that candidates left off the ballot are considered tied for the last ranking). Briefly, EAR can be thought of as a multiwinner extension of the single-winner Bucklin method. A more general description can be found in Aziz and Lee (2020), but for this paper we implemented the following algorithm: Set quota equal to $q = \frac{V}{k+1}$ and let the vote total for each candidate c be v_c. Initialize v_c as the first-place vote totals for each candidate. If $v_c > q$, candidate c is elected. If k candidates are not elected initially, all second-place votes are added to v_c, and the process continues, electing candidates as v_c exceeds quota. Similar to our explanation of the Chamberlin-Courant rule below, since the Scottish data is highly truncated we use two implementations to deal with truncated ballots once rankings have been exhausted, an "optimistic model" (OM) and a "pessimistic model" (PM). Under OM, exhausted ballots count for all candidates in the election, and under PM exhausted ballots count for only the candidates listed on the ballot. For example, suppose an election has 4 candidates, A, B, C and D, and a voter casts a ballot with A in first-place and B is second, with C and D not ranked. Under either model, first-round votes only count for A and second-round votes count for both A and B. Under OM, third-round votes would count for A, B, C, and D, whereas under PM, third-round votes would count for only A and B.

Chamberlin-Courant Rule (CC) (Chamberlin and Courant, 1983). There are a number of variants of this rule. In each, voters are "assigned" a member of the winning set. This assignment gives each voter a measure of individual satisfaction; these measures are combined to create a measure of social satisfaction and the winning committee is defined to be the set of candidates that maximizes this social satisfaction. The version of CC we consider is the original method proposed

Table 16
Table showing which elections methods are not guaranteed to satisfy the various fairness criteria. An X indicates that violations are possible.

	Scottish	Meek	EAR OM	EAR PM	CC OM	CC PM
$ILVB$	X	X	X	X		
$IWVB$	X	X	X	X	X	X
$IWVB^*$	X	X	X	X		

in Chamberlin and Courant (1983) which uses a candidate's Borda score and its sum to determine the level of individual satisfaction and social satisfaction, respectively. Formally, let r_{ic} denote the rank of candidate c on voter i's ballot, so that this voter gives $m - r_{ic}$ points to c. For a fixed committee X of size k, let $V_c(X)$ denote the set of voters for whom candidate c is the most preferred candidate in the committee X. CC selects the committee X of size k which maximizes the value

$$\sum_{c\in X}\sum_{i\in V_c(X)} m - r_{ic}.$$

Because CC relies on Borda scores, we must decide how many points a ballot contributes to a candidate subset if none of the candidates in that subset are ranked on the ballot. Following Baumeister et al. (2012), we use two implementations of CC, an "optimistic model" and a "pessimistic model". Baumeister et al. study only single-winner Borda count in Baumeister et al. (2012), but their models of processing partial ballots and their terminology translate to the multiwinner CC setting. There are other ways of adapting CC to partial ballots, but we feel that including these two variants is adequate for the purposes of this article.

Chamberlin-Courant Rule, Optimistic Model (CC OM). Suppose a ballot does not rank any candidates from a candidate subset currently under consideration and the ballots ranks t candidates. Then this ballot gives $m - t - 1$ points to the subset.

Chamberlin-Courant Rule, Pessimistic Model (CC PM). If a ballot does not rank any candidates from a candidate subset currently under consideration, then this ballot gives no points to the subset.

Example election, modified from Hill et al. (1987).

To illustrate the differences between the methods, consider the following election with 4 candidates and 2 seats:

3	4	2	1	2	2
A	A	D	B	B	C
C	C	A		D	D
D	B	C		C	B
				A	

STV (and Meek STV): The quota is 5 and A is elected in the first round. The excess votes from A accrue to C, who wins in the second round. The winning committee is $\{A, C\}$.

EAR: The quota is 5 and A is elected in the first round. Once second-place votes are added in, in PM C now has 9 votes and is elected (next highest is D with 6 votes). The winning committee is $\{A, C\}$. Note that under OM the result is the same but the numbers are different: in the second round C will have 10 votes (1 extra vote from the voter who only voted for B) and D will have 7 votes.

CC: Under PM, the points for committee $\{A, B\}$ accrue as follows: the committee receives 3 points each from ballots of the form ACD, ACB, B, and BDCA, 2 points for each DAC ballot, and 1 point for each CDB ballot, for a total of $3 \times (3 + 4 + 1 + 2) + 2 \times 2 + 1 \times 2 = 36$ points. Similar calculations show that other committees receive the following points: $\{A, C\} = 29$, $\{A, D\} = 35$, $\{C, B\} = 31$, $\{D, B\} = 26$, and $\{C, D\} = 30$. The winning committee under CC PM is $\{A, B\}$. Under OM, the scores are virtually identical except that committees $\{A, C\}$, $\{A, D\}$, and $\{D, C\}$ each receive an extra two points from the voter who bullet-voted B. Thus $\{A, B\}$ still has a score of 36 whereas $\{A, D\}$ now has 37, and the winning committee under CC OM is $\{A, D\}$.

Data availability

All data (and code) used is included as links at relevant points in the article.

References

Arrow, K., 1951. Social Choice and Individual Values, first ed. John Wiley & Sons, Inc., New York.
Aziz, H., Lederer, P., Peters, D., Peters, J., Ritossa, A., 2024. Committee monotonicity and proportional representation for ranked preferences. Preprint: arXiv:2406.19689.
Aziz, H., Lee, B., 2020. The expanding approvals rule: improving proportional representation and monotonicity. Soc. Choice Welf. 54, 1–45.
Aziz, H., Lee, B., 2022. TA characterization of proportionally representative committees. Games Econom. Behav. 133, 248–255.
Baumeister, D., Faliszewski, P., Lang, J., Rothe, J., 2012. Campaigns for lazy voters: Truncated ballots. In: Proceedings of the International Conference on Autonomous Agents and Multiagent Systems. AAMAS'12, pp. 57–584.
Blom, M., Conway, A., Stuckey, P.J., Teague, V.J., 2020. Shifting the balance- of-power in stv elections. In: Electronic Voting: 5th International Joint Conference, E-Vote-ID 2020, Bregenz, Austria, October (2020) 6–9, Proceedings, vol. 5, Springer, pp. 1–18.
Blom, M., Teague, V., Stuckey, P.J., Tidhar, R., 2016. Efficient computation of exact irv margins. In: ECAI 2016. IOS Press, pp. 480–488.
Börgers, C., 2010. Mathematics of Social Choice. SIAM, Philadelphia.
Brams, S., 1982. The AMS nominating system is vulnerable to truncation of preferences. Notices Amer. Math. Soc. 29, 136–138.
Brill, M., Peters, J., 2023. Robust and verifiable proportionality axioms for multiwinner voting. arXiv:2302.01989.
Cary, D., 2011. Estimating the margin of victory for instant-runoff voting. In: Electronic Voting Technology Workshop/Workshop on Trustworthy Elections. EVT/WOTE, USENIX Associa-tion.
Chamberlin, J., Courant, P., 1983. Representative deliberations and representative decisions: Proportional representation and the Borda rule. Am. Political Sci. Rev. 77 (3), 718–733.
Clelland, J., 2026. Ranked choice voting and condorcet failure in the Alaska 2022 special election: How might other voting systems compare? Math. Enthus. 23 (3), 185–196.
Diss, M., Kamwa, E., Tlidi, A., 2024. The effect of close elections on the likelihood of voting paradoxes: Further results in three-candidate elections. Preprint: hal-04230359v1.
Dummett, M., 1984. Voting Procedures. Oxford University Press.
Elkind, E., Faliszewski, P., Skowron, P., Slinko, A., 2017. Properties of multiwinner voting rules. Soc. Choice Welf. 28, 599–632. http://dx.doi.org/10.1007/s00355-017-1026-z.
Felsenthal, D., 2012. Review of Paradoxes Afflicting Procedures for Electing a Single Candidate. In: Studies in Choice and Welfare, Springer Berlin, Heidelberg.
Felsenthal, D.S., 2019. On paradoxes afflicting voting procedures: Needed knowledge regarding necessary and/or sufficient condition(s) for their occurrence. In: Laslier, J.F., Moulin, H., Sanver, M., Zwicker, W. (Eds.), The Future of Economic Design. In: Studies in Economic Design, Springer, Cham, http://dx.doi.org/10.1007/978-3-030-18050-8_13.
Felsenthal, D.S., Moaz, Z., 1992. Normative properties of four single-stage multi-winner electoral procedures. Behav. Sci. 37, 109–127.
Fishburn, P.C., Brams, S., 1983. MParadoxes of preferential voting. Math. Mag. 56 (4), 207–214.
Fishburn, P.C., Brams, S., 1984. Manipulability of voting by sincere truncation of preferences. Public Choice 44, 397–410.
Graham-Squire, A., 2024. Conditions for fairness anomalies in instant-runoff voting. In: Jones, M., McCune, D., Wilson, J. (Eds.), Mathematical Analyses of Decisions, Voting and Games. In: Contemporary Mathematics, vol. 795, American Mathematical Society.
Graham-Squire, A., McCune, D., 2024. Ranked choice wackiness in Alaska. Math Horizons 31 (1), 24–27.
Graham-Squire, A., McCune, D., 2025. An examination of ranked-choice voting in the United States, 2004–2022. Representation 61 (1), 1–19.
Graham-Squire, A., Zayatz, N., 2021. Lack of monotonicity anomalies in empirical data of instant-runoff elections. Representation 57 (4), 565–573.
Heckelman, J.C., Chen, F.H., 2013. Strategy proof scoring rule lotteries for multiple winners. J. Public Econ. Theory 15 (1), 103–123.
Herschlag, G., et al., 2020. Quantifying gerrymandering in North Carolina. Stat. Public Policy 7 (1), 30–38.
Hill, I.D., Wichmann, B.A., Woodall, D., 1987. Algorithm 123: Single transferable vote by meek's method. Comput. J. 30, 277–281.
Lepelley, D., Merlin, V., 2001. Scoring runoff-paradoxes for variable electorates. Econom. Theory 17, 53–80.
Magrino, T., Rivest, R., Shen, E., 2011. Computing the margin of victory in {IRV} elections. In: 2011 Electronic Voting Technology Workshop/Workshop on Trustworthy Elections. EVT/WOTE 11.

McCune, D., Graham-Squire, A., 2024. Monotonicity violations in scottish local government elections. Soc. Choice Welf. http://dx.doi.org/10.1007/s00355-024-01522-5.

McCune, D., Wilson, J., 2024. The spoiler effect in multiwinner ranked-choice elections, arXiv:2403.03228.

Meek, B.L., 1994. Voting matters. 1, pp. 1–11.

Miller, N., 2017. Closeness matters: monotonicity failure in IRV elections with three candidates. Public Choice 173, 91–108. http://dx.doi.org/10.1007/s11127-017-0465-5.

Miller, N.R., 2019. Reflections on Arrow's theorem and voting rules. Public Choice 179, 113–124.

Monroe, B., 1995. Fully proportional representation. Am. Political Sci. Rev. 89 (4), 925–940.

Moulin, H., 1988. Condorcet's principle implies the no show paradox. J. Econom. Theory 45 (1), 53–64.

Nurmi, H., 1999. Voting Paradoxes and how To Deal with Them. Springer Berlin, Heidelberg.

Roy, V., 2020. Convergence diagnostics for Markov chain Monte Carlo. Annu. Rev. Stat. Appl. 7 (1), 387–412.

Sáchez-Fernández, L., Elikind, E., Lackner, M., Fernández, N., Fisteus, J., Basanta Val, P., Skowron, P., 2017. Proportional justified representation. Proc. the AAAI Conf. Artif. Intell. 31 (1).

Tideman, N., 1995. The single transferable vote. J. Econ. Perspect. 9 (1), 27–38.

Woodall, D., 1994. Properties of preferential election rules. Voting Matters 3, 8–15.

Young, H.P., 1974. An axiomatization of Borda's rule. J. Econom. Theory 9, 43–52.

Young, H.P., 1975. Social choice scoring functions. SIAM J. Appl. Math. 28 (4), 824–838.