A diagram for analyzing ordinal voting systems

Andrew C. Eggers*

January 16, 2020

Abstract

I introduce a diagram for describing and analyzing single-winner elections in which voters rank the candidates – a class of voting systems including positional methods (e.g. plurality, Borda count, anti-plurality), Condorcet methods, and instant-runoff voting (i.e. ranked-choice voting or the alternative vote). The diagram shows how the outcome of an election depends on each candidate's share of top rankings as a function of the voting system and the pattern of lower rankings. Using as examples two Brexit polls, a mayoral election in San Francisco, and the US's first instant-runoff congressional election (all from 2018), I show how the diagram can concisely present preferences and results under different voting systems, identify Condorcet cycles, highlight system properties such as join-inconsistency and the no-show paradox, and illuminate strategic voting incentives.

*Nuffield College, University of Oxford; andrew.eggers@nuffield.ox.ac.uk. Thanks to Tobias Nowacki, who has puzzled over the IRV ternary diagram with me as part of our joint work on strategic voting in that system, and to Bernie Grofman, Jack Santucci, and Alex Kuo for useful discussions; thanks also to Ben Lauderdale for helpful comments and especially for sharing the June 2019 Brexit poll analysis.

1 Introduction

In November, 2018, as the initial deadline for the United Kingdom's departure from the European Union ("Brexit") approached, YouGov conducted a survey in which over 20,000 British respondents were asked to rank three alternatives: remain in the EU (Remain), accept the withdrawal agreement promoted by Prime Minister Theresa May (Deal), or leave the EU without a deal (No Deal). Table 1 shows the estimated proportion of the UK electorate with each possible ordering of the three alternatives.[1] The poll was highly salient as opponents of Brexit clamored for a second referendum to allow voters to choose among these alternatives.[2]

Table 1: Brexit preferences, November 2018

		Second choice		
		Remain	Deal	No Deal
First choice	Remain (R)	-	.375	.087
	Deal (D)	.059	-	.212
	No Deal (N)	.038	.228	-

NOTE: Proportions estimated from YouGov poll via MRP by Ben Lauderdale. See footnote 1.

If voters voted according to the preferences in Table 1, which alternative would have won in a hypothetical Brexit referendum in November 2018? Does the winner depend on the method used (e.g. plurality vs. Borda count)? Which voting method would produce the most favorable result for each alternative? With a bit of computation we can certainly answer these questions using the numbers in Table 1 and present the result in additional tables.[3] But tabular presentations are often less intuitive than graphical ones, including in the case of elections (Saari, 1994; Tufte, 2001; Kastellec and Leoni, 2007). A good diagram would quickly indicate, in a consistent way for several relevant election methods, not only who wins under a given method but also how close each other candidate was to

[1] See `https://yougov.co.uk/topics/politics/articles-reports/2018/12/06/mays-brexit-deal-leads-just-two-constituencies-it-` for the YouGov summary and a link to Ben Lauderdale's report, which describes the survey and the MRP model used to analyze it.

[2] See for example, Michael Savage, "Remainers gear up for second referendum, with new NHS pledge", *The Guardian*, 8 Dec. 2018 (`https://bit.ly/2CKZDUC`).

[3] For example, summing the rows indicates that Remain would win a plurality contest (with Deal in a distant second), while giving each candidate two points for each first-place vote and one point for each second-place vote indicates that Deal would win under Borda count (with Remain fairly close behind).

winning. Beyond representing the outcomes for a particular set of preferences like those in Table 1, a good diagram might also illustrate some of the desirable and undesirable properties of each voting system, as seen in the work of Donald Saari 1994; 2011.

The fundamental reason why we need such a diagram is that in systems where voters' second choices matter (as in e.g. Borda count or instant-runoff) the election outcome depends on more numbers than can be comfortably digested in a table or visualized in a straightforward way. The preferences in Table 1 can be summarized by five numbers,[4] but humans find it difficult to visualize more than three dimensions (and even to make sense of two-dimensional representations of three dimensions). What we need is a digestible and consistent way to illustrate the results of a wide range of ordinal voting methods. As I discuss in the next section, diagrams exist for visualizing some classes of ordinal methods (most notably Saari (1994)'s triangle for representing positional methods), but as far as I know there is no other diagram that can depict a result under such disparate methods as the Borda count, the Kemeny (Condorcet) method, and instant-runoff voting (IRV) in a consistent way.

This paper introduces a consistent approach to diagramming the result of any anonymous single-winner election method in which voters submit ranked preferences (i.e. ordinal voting systems). It can be used not just to analyze a specific result for a given set of ordinal ballots and one voting method but also to compare results across voting methods and analyze some of the properties of these methods. To address the dimensionality problem, the diagram shows which alternative wins as a function of two parameters (the proportion of voters ranking each candidate first) conditional on the voting rule and three or more other parameters (the proportion of each candidate's supporters who rank each other candidate second). Put differently, the diagram emphasizes how the winner in a given system depends on the row sums of Table 1 as a function of the column proportions within each row. This approach easily accommodates incomplete rankings, which pose a problem for some previous graphical approaches,[5] and it can be useful in IRV and

[4]There are six numbers but they sum to one, so one can be omitted.

[5]Essentially, incomplete rankings complicate any graphical approach where the simplex vertices are unique preference orderings, such as Saari and Valognes (1999).

Condorcet elections with more than three alternatives.

The diagram has several uses. Most directly, it tells us how the election result would change in a given system if the distribution of voters' first choices were to shift, holding fixed the pattern of second preferences (as might happen if turnout increased or decreased among the supporters of one of the alternatives); it can also do the opposite, telling us how the result would change if the distribution of voters' first choices remained fixed but the pattern of second choices shifted. By comparing two diagrams representing different elections under the same voting rule, we can see how a shift in preferences involving both first and second preferences might affect the outcome. By comparing diagrams from two systems for a given preference profile, we can also see how the result would change holding expressed preferences fixed if we moved from one system to another. The diagram can make intuitive some well-known properties of voting systems, such as join-inconsistency or the no-show paradox. Because the diagram summarizes five preference parameters in two dimensions,[6] it can be viewed not just as a means of analyzing election methods but also as a concise summary of preference profiles.[7]

2 The plurality ternary diagram and previous extensions to it

The diagram introduced this paper is an extension of the ternary diagram that has been used to represent vote shares in three-candidate plurality elections since at least Ibbetson (1965). In this section I review the ternary diagram and previous extensions to it.

Consider a plurality election or referendum in which voters choose among three alternatives, a, b, and c. Let v_i denote the proportion of voters choosing some alternative i. Then a plurality election result can be summarized by the vote shares v_a, v_b, v_c and represented on a *ternary diagram*, an example of which appears in Figure 1.a. A ternary

[6]For example, we can infer all five numbers in Table 1 from the Condorcet diagram in Figure 3.d: the location of the dot gives us the row sums and the angle of the majority tie lines gives us the column proportions in each row.

[7]I plan to release the code behind the figures before this paper is published. In the meantime, readers may produce figures interactively (assuming only complete ballots) at `https://andyeggers.shinyapps.io/voting_viz/`.

diagram represents three shares that must sum to one (and thus that can be summarized by two numbers) in a way that maintains symmetry among the three alternatives.[8] The ternary diagram in Figure 1.a represents the proportion of voters supporting Remain, Deal, and No Deal in an election mirroring the Nov. 2018 Brexit poll. Each vertex indicates a result where the named alternative receives all votes (e.g. a result at the bottom-left vertex indicates Remain receiving a vote share of 1). The gray lines indicate results where a given alternative receives a given share; moving away from the Remain vertex, for example, we encounter results where Remain wins shares of .75, .5, and .25 (indicated by solid gray lines); dashed and dotted lines show corresponding levels for No Deal and Deal, respectively. The gray dot in the center shows a result where the shares are equal. The black dot below and to the left of the center indicates the estimated proportion of voters listing each alternative first in the November 2018 poll (i.e. the row sums in Table 1).

To better understand a plurality result, we can divide the ternary diagram into regions in which each candidate would win. Note that $v_i = v_j$ describes the line where alternatives i and j receive the same vote share, which we will call the *ij-plurality tie line*. As an example, the Deal-No Deal plurality tie line is drawn in Figure 1.b. Figure 1.c adds the plurality tie lines for the other two pairs of alternatives; jointly, these three lines partition the ternary diagram into six sub-triangles, each corresponding to a unique order of finish among the three alternatives. Figure 1.d then divides the ternary diagram into three *win regions* identifying results that would lead to each alternative being chosen; a black dot indicates the actual plurality result given voting ballots reflecting Table 1. Remain is the plurality winner.

Other researchers have extended the ternary diagram for plurality elections beyond the three-alternative case by having two of the vertices represent votes for specific candidates and the third represent votes for everyone else. Grofman et al. (2004) apply this approach

[8]The same information can be plotted on standard coordinates (with the horizontal axis indicating e.g. No Deal's vote share and the vertical axis Deal's), but the symmetry of the problem would then be lost: the regions where No Deal and Deal win a plurality election would have a different shape than Remain's win region. (Katz and King (1999) and Grofman et al. (2004) provide excellent and more detailed explanations of the ternary diagram.

Figure 1: Ternary diagram, tie lines, and win regions for plurality

1.a: Ternary diagram

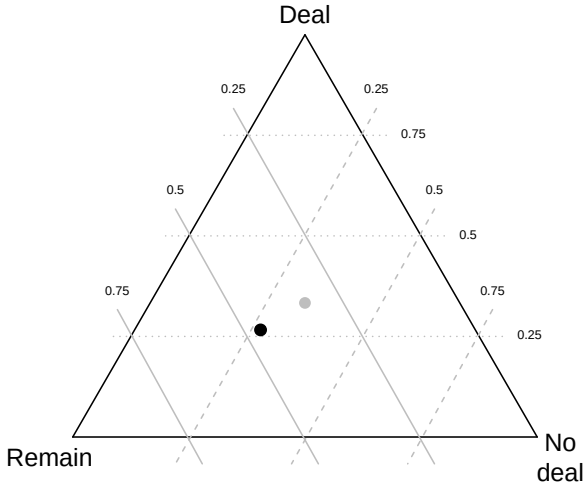

1.b: One plurality tie line

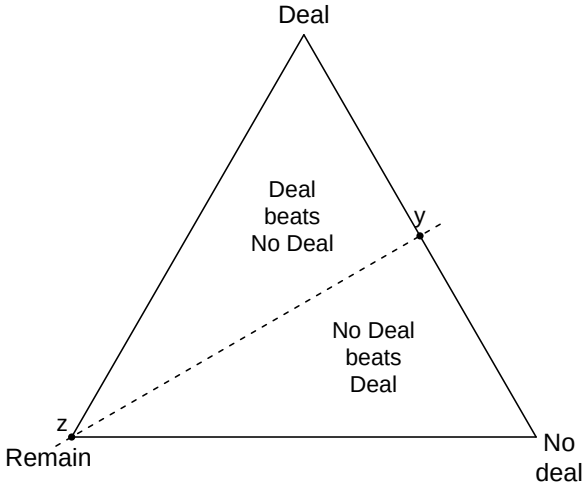

1.c: All plurality tie lines

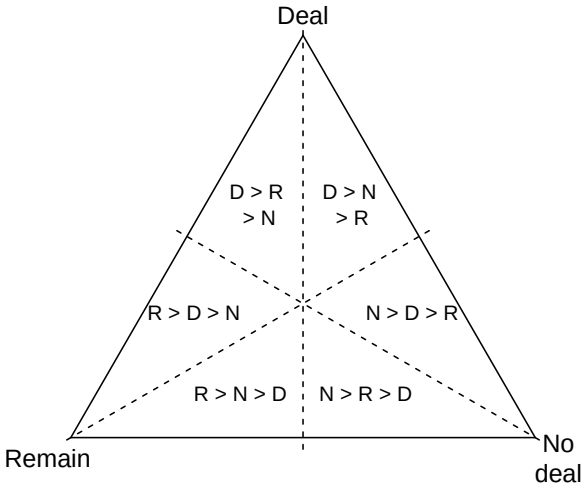

1.d: Win regions and plurality result

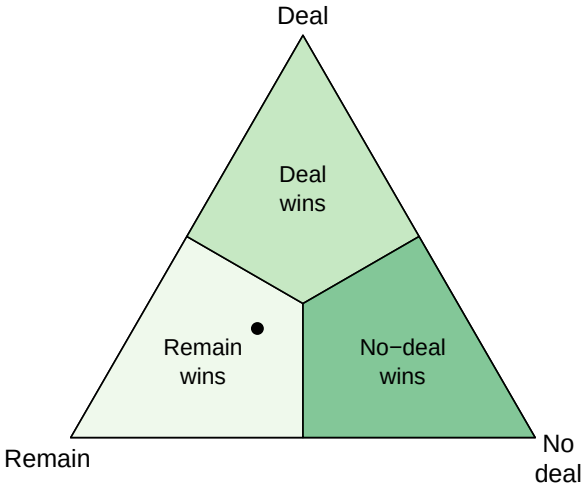

NOTE: A tie line indicates first-preference shares at which two alternatives tie according to a given rule and given a pattern of second preferences. Figure 1.b shows only the Deal-No deal tie line; Figure 1.c shows all three tie lines, which jointly divide the ternary diagram into a different sub-triangle for each order of finish. In Figure 1.d divides the diagram according to which candidate wins; the black dot shows the actual first-preference result given the preferences in Table 1.

to district-level Italian election results, showing a consolidation of support toward the two main party blocs over three elections. Grofman et al. (2004) also review the Nagayama diagram, which similarly plots the vote share for the candidate receiving the largest and second-largest number of votes and can be used to track competitiveness and party fragmentation (see also Dunleavy and Diwakar, 2013). These approaches could also be used to depict vote shares in PR elections or first-preference shares in IRV elections (e.g. in Australia), although clearly when the support for some alternatives is aggregated (and lower preferences are ignored, in IRV) we do not retain enough information to determine the election outcome.

Saari (1994) generalized the plurality ternary diagram to voting methods in which the alternatives are assigned points or scores, a class that includes positional methods (e.g. plurality, Borda count, anti-plurality) and approval voting. The basic idea is simple: whereas the plurality ternary diagram represents the share of votes won by each alternative, Saari's diagram represents the share of *points* won by each alternative. Figure 2a illustrates Saari's approach (which I will call the *point-share triangle*[9]) using the Nov. 2018 Brexit preferences from Table 1. The point shares for each alternative are shown for plurality (where points for first, second, and third rankings are $(1, 0, 0)$), Borda count (where points are $(1, \frac{1}{2}, 0)$), and anti-plurality (where points are $(1, 1, 0)$). The line connecting these results, which Saari (1994) calls the *procedure line*, indicates the point shares under every possible positional system between the two extremes of plurality and anti-plurality. In this case, it shows that the winner changes from Remain to Deal as we assign a greater weight to second rankings; the order of finish under plurality is Remain-Deal-No Deal, but as we give more weight to second rankings Remain drops below Deal and ultimately to last place behind No Deal. The point-share triangle is an impressively compact way to summarize the possible outcomes for a given ballot profile under various positional methods, but it cannot be easily extended to handle voting methods (such as Condorcet or IRV) in which the winner is not determined by a point tally.[10]

[9]Saari (1994) refers to it as the representation triangle. Others (e.g.) have referred to it as the Saari triangle, though Nurmi (1999, pg. 34) also uses that term to refer to what I call profile triangles.

[10]Saari (1994, pg. 118) uses a "fat triangle" to connect the point-share triangle to a pairwise majority relation. This shows how positional and pairwise procedures may disagree, but it is not a method for

Figure 2: Saari's representations of Nov. 2018 Brexit preferences

2.a: Point-share triangle **2.b: Profile triangle**

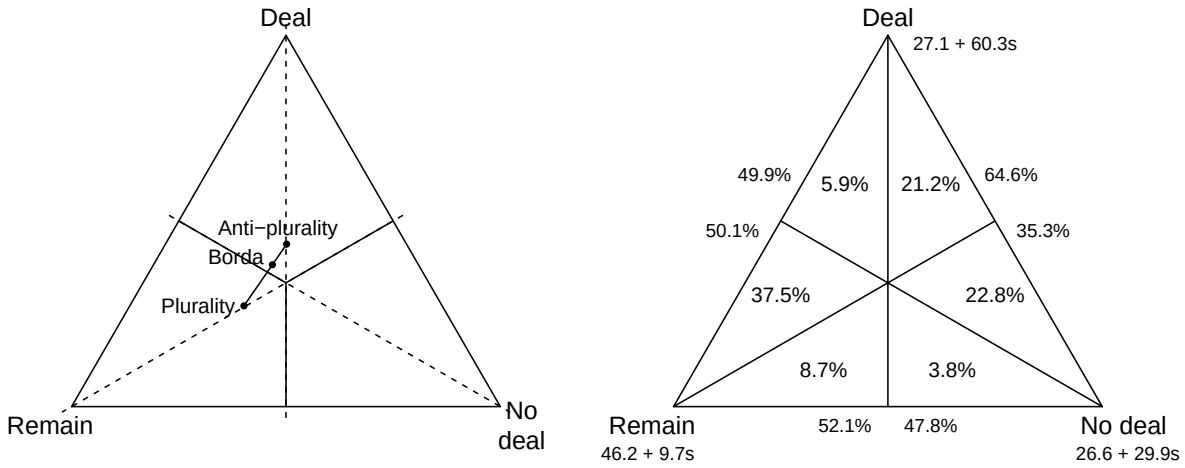

Saari (1994) introduced another triangle designed to summarize a profile of ballots or preferences; we can call this the *profile triangle*. The profile triangle reorganizes preference data like that in Table 1 and adds summary numbers relevant to positional and pairwise voting procedures. Figure 2.b illustrates the approach with the Brexit data from Table 1. Each of the six numbers inside the triangle indicates the proportion of ballots with a given preference ordering; e.g. 37.5% is the percentage of ballots ranking the alternatives Remain-Deal-No Deal. At the middle of the edge connecting two alternatives (e.g. the edge between the Remain vertex and the Deal vertex), we write the proportion favoring each of those alternatives in a pairwise vote (e.g. 50.1% for Remain and 49.9% for Deal).[11] At the vertex for an alternative (e.g. the Remain vertex), we write the positional score for that alternative, where a first ranking is worth one point and a second ranking is worth s (e.g. $46.2 + 9.7s$ for Remain).[12] Saari (2011) uses the profile triangle to show how one can generate disagreements between the plurality result, other positional results, and pairwise results; he also extends this representation to develop his profile coordinate approach to analyzing voting rules. The profile triangle can be viewed as a specially

representing a specific election result.

[11] These come from summing the values in the three triangles on either side of the Remain-Deal plurality tie line.

[12] For a given vertex, the first number comes from summing the numbers in the two closest triangles to that vertex (where that alternative is ranked first) and the second comes from summing the two next closest triangles (where that alternative is ranked second).

formatted reorganization of the numbers in Table 1 that makes it easier to determine the winner under various methods (including Condorcet and IRV).[13]

Another way to extend the plurality ternary diagram to three-candidate ordinal elections is to replace its three vertices (one for each alternative) with six vertices (one for each possible ranking), but in these higher-dimensional spaces the tradeoff between geometric intuition and generality becomes stark. For example, Saari (2011) uses an "opened tetrahedron" to represent restricted profiles in which only four of the six possible rankings are chosen; to allow for five rankings to be chosen, he recommends similarly opening a four-dimensional tetrahedron into five three-dimensional tetrahedrons, a prospect that some analysts find daunting (e.g. Dunleavy and Diwakar, 2013). (The problem only becomes worse if we allow incomplete rankings, such that there are nine vertices instead of six.) Geometric representations that allow a dimension for each ordering seem better suited for proving voting system properties than for communicating election results or exploring the properties of naturally occurring profiles, which are my objectives.

This paper introduces a diagram that draws on these previous approaches but offers several innovations. My diagram maintains the basic setup of the classic ternary diagram and Saari's point-share triangle, in that each vertex of the triangle corresponds to a distinct alternative, and a result closer to a vertex indicates more support for the associated alternative. Like Saari's point-share triangle, my diagram is not restricted to plurality elections and can depict the election result for a fixed ordinal profile under a variety of rules. Unlike Saari's point-share triangle, my diagram can show results for any ordinal method (including Condorcet methods and IRV), not just point-based methods.[14] My approach can be seen as a new extension of the plurality ternary diagram that more fully represents ordinal preference data and, unlike previous diagrams, accommodates the full range of ordinal voting methods (including when some ballots do not rank all candidates).

[13]For example, in an IRV election No Deal would finish last in first-preference votes (based on the numbers at the vertices, taking $s = 0$), after which Remain would defeat Deal 50.1% to 49.9% (based on the numbers along the edge between the Remain and Deal vertices).

[14]In its favor, Saari's point-share triangle can compare two or more positional methods on the same diagram, while in my approach such comparisons work best with separate diagrams. Saari's point-share triangle can also show the convex hull for approval voting, while my diagram is limited to ordinal methods.

As an added benefit, my diagram illustrates properties of voting systems that, although known, can otherwise be hard to understand.

3 The diagram

Consider an ordinal election or referendum in which voters rank three alternatives. Let v_{ij} denote the proportion of voters ranking alternative i first and j second, for $i, j \in \{a, b, c\}$. (A ballot ranking i first and j second implicitly ranks k last.) Let v_{ix} denote the proportion of ballots listing $i \in \{a, b, c\}$ first and no candidate second, which is a relevant possibility if incomplete rankings are permissible.[15] Then the winner of any election using an anonymous ordinal voting method depends on the vector

$$\mathbf{v} \equiv (v_{ab}, v_{ac}, v_{ax}, v_{ba}, v_{bc}, v_{bx}, v_{ca}, v_{cb}, v_{cx}).$$

Because its components must sum to one, $\mathbf{v}$ lives on the 8-dimensional simplex. In principle, that 8-dimensional space could be divided into win regions in the same way the 2-dimensional ternary diagram was above. The problem is that visualizing win regions in eight dimensions is essentially impossible.

My solution is to depict win regions as a function of *first-preference shares* (the proportion of voters ranking each candidate first) given *conditional second-preference shares* (the proportion of voters ranking each candidate second conditional on ranking a given candidate first). Let $v_a \equiv v_{ab} + v_{ac} + v_{ax}$ denote a's first-preference share, with v_b and v_c defined equivalently; let $p_{ab} \equiv \frac{v_{ab}}{v_a}$ indicate the proportion of ballots listing b second conditional on listing a first, with p_{ac}, p_{ax}, p_{ba}, etc. defined similarly and the set of these proportions collectively referred to as conditional second-preference shares. Noting that $v_c = 1 - v_a - v_b$ and $p_{ax} = 1 - p_{ab} - p_{ac}$ (and similarly for p_{bx} and p_{cx}), we can express each component of $\mathbf{v}$ in terms of two first-preference shares (v_a and v_b) and six conditional

[15] Incomplete rankings are permitted in IRV elections in several Australian states (Farrell and McAllister, 2006). IRV mayoral elections in London and San Francsico set an upper limit on the number of rankings.

second-preference shares $(p_{ab}, p_{ac}, p_{ba}, p_{bc}, p_{ca}, p_{cb})$.[16] We can then use the ternary diagram to represent the proportion of voters ranking each alternative first (v_a, v_b, v_c), just as we do (implicitly) in the plurality ternary diagram, while dividing the diagram into regions in which each alternative would win, given the conditional second-preference shares and the voting system. I call these regions *first-preference win regions*.

Note that by definition the winner of any anonymous ordinal voting method depends only on $\mathbf{v}$, and note also that (given the conditional second-preference shares) $\mathbf{v}$ depends only on v_a, v_b, and v_c. It follows that, again given the conditional second-preference shares, the winner of any anonymous ordinal voting system depends only on v_a, v_b, and v_c and can therefore be identified on the ternary diagram.[17]

I now apply this approach to prominent ordinal voting rules.

3.1 Condorcet methods

A Condorcet method selects as winner a candidate who wins a pairwise contest against every other candidate (i.e. a Condorcet winner). Condorcet methods differ in how to proceed if there is no Condorcet winner. I begin by showing how to identify first-preference regions in which there is a Condorcet winner; I then show how to fill in the remaining cyclic region depending on the specific Condorcet method.

Let a *majority tie line* for two candidates indicate first-preference shares where, given conditional second-preference shares, the two candidates would tie in a pairwise majority vote. This is true for i and j where

$$v_i + v_{ki} = v_j + v_{kj},$$

which can be written in terms of first-preference shares and conditional second-preference shares as

$$v_i = v_j + v_k(p_{kj} - p_{ki}). \tag{1}$$

[16] For example, $v_{ac} = v_a p_{ac}$ and $v_{bx} = v_b(1 - p_{ba} - p_{bc})$ and $v_{cb} = (1 - v_a - v_b)p_{cb}$.

[17] The mapping will not necessarily be simple, e.g. for the estimated centrality rule (Tideman, 2017, pp. 215-217).

Equation 1 describes the ij majority tie line.

To construct a majority tie line we identify the two points where the line intersects an edge of the ternary diagram. One intersection point, which we call y, is located where $v_k = 0$ and $v_i = v_j = \frac{1}{2}$; this is where half of the voters rank i first and the other half rank j first, which obviously produces a majority tie. Point y is labeled for the Deal-No Deal majority tie line on Figure 3.a and Figure 3.b. The location of the other intersection point of the ij-majority tie line, which we call z, depends on the conditional second-preference shares for ballots ranking k first. If $p_{kj} > p_{ki}$ (i.e. if ballots ranking k first more often rank j second than i second), then z is located where $v_i = v_k(p_{kj} - p_{ki})$, $v_j = 0$, and $v_k = 1 - v_i$; otherwise, z is located where $v_i = 0$, $v_j = v_k(p_{ki} - p_{kj})$, and $v_k = 1 - v_j$. Figure 3.a shows extreme cases for the Deal-No Deal majority tie line: if Remain voters (i.e. those who submit ballots ranking Remain first) always rank Deal second, then the Deal-No Deal majority tie line is the one connecting points y and z_1; if Remain voters always rank No Deal second, then the Deal-No Deal majority tie line is the one connecting point y and z_2.

Figure 3.b shows the actual Deal-No Deal majority tie line given the Brexit preferences from the Nov. 2018 poll. Remain voters are much more likely to rank Deal second than No Deal second, such that the Deal-No Deal majority tie line is much closer to the No Deal vertex than the Deal vertex, and Deal beats No Deal in a pairwise majority vote over much more than half of the ternary diagram. The key general point is that the "tilt" of the ij-majority tie line is a measure of how k's supporters (i.e. voters who rank k first) tend to rank i and j: the more k's supporters favor i as their second choice, the more the ij-majority tie line tilts away from i's vertex, indicating that i can tie j in a pairwise vote when receiving a lower first-preference share.

Figure 3.c shows all three majority tie lines for the Brexit case. The angle of the majority tie lines reflect Remain voters' preference for Deal over No Deal (noted immediately above), Deal voters' preference for No Deal over Remain (indicated by the counter-clockwise tilt of the Remain-No Deal tie line), and No Deal voters' preference for Deal over Remain (indicated by the counter-clockwise tilt of the Remain-Deal tie line).

Figure 3: Tie lines and first-preference win regions for Condorcet methods

3.a: Range of Deal-No Deal tie lines

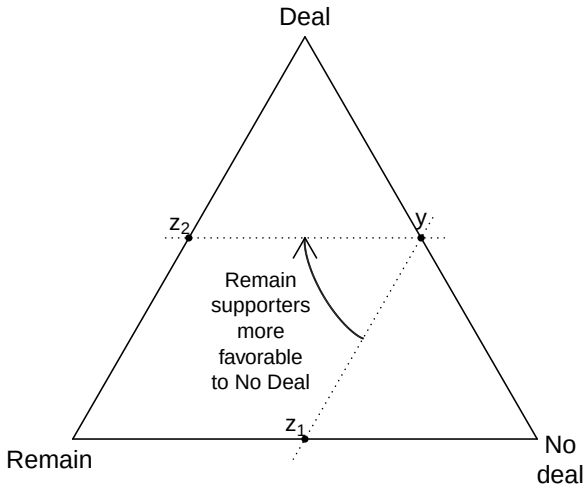

3.b: Actual Deal-No Deal tie line

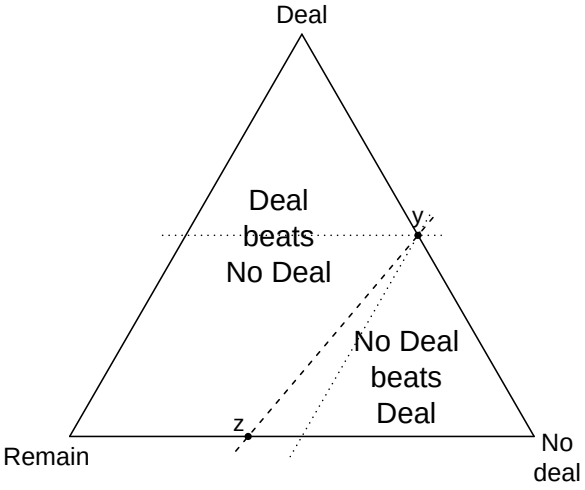

3.c: All tie lines

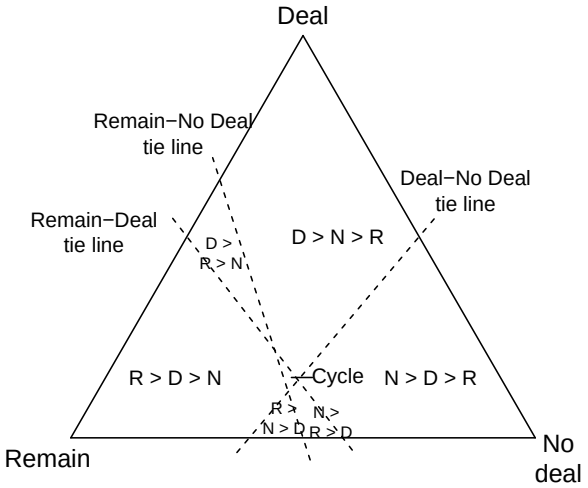

3.d: First-preference win regions

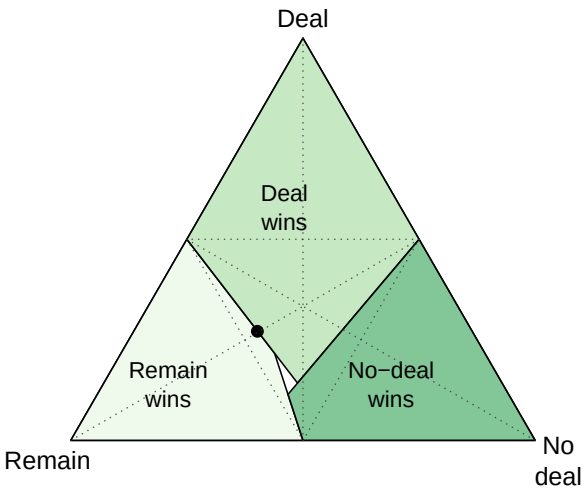

NOTE: A majority tie line indicates first-preference shares such that two candidates tie in a pairwise vote, given conditional second-preference shares. The Deal-No Deal majority tie line starts at the midpoint of the edge between the Deal vertex and the No Deal vertex; its angle depends on the proportion of Remain ballots on which each alternative is ranked second (Figure 3.a). The actual Deal-No Deal majority tie line (based on the Nov. 2018 Brexit poll) indicates Remain supporters' strong preference for Deal over No Deal (Figure 3.b). The three majority tie lines partition the ternary diagram, indicating first-preference shares producing each transitive order and a cycle (Figure 3.c) and first-preference shares at which each alternative is the Condorcet winner (Figure 3.d).

Generically, majority tie lines do not intersect at a point; the three majority tie lines divide the ternary into six sub-triangles in which a different transitive majority preference order obtains, plus one sub-triangle (labeled "Cycle" in Figure 3.c) in which there is a majority cycle: for first-preference shares in this triangle (given the pattern of second preferences), a majority prefers Deal over No Deal (because the result is above the Deal-No Deal majority tie line), a majority prefers No deal to Remain (because the result is to the right of the Remain-No deal majority tie line), and a majority prefers Remain to a Deal (because the result is below the Remain-Deal majority tie line).

Figure 3.d shows the implied first-preference win regions given the Nov. 2018 pattern of conditional second preferences, with a black dot indicating the first-preference result. (To aid comparison across diagrams, I include faint plurality tie lines and lines where each alternative receives half of first-preference votes; first-preference win region boundaries for Condorcet methods must be within the triangle formed by the latter set of lines.) Deal and Remain essentially tie for first, with the first-preference result lying along the Deal-Remain majority tie line.

In Figure 3.d I leave the cyclic triangle unfilled, but there are various methods for determining the winner in the event of a cycle among three candidates. On the diagram, the cyclic triangle can be filled in according to the specific method:

- Black's method (or Condorcet-Borda) decides the winner by Borda count, so the cyclic triangle can be filled in using the method shown in the next section.

- Condorcet-IRV (equivalent in the three-candidate case to what Tideman (2017, pg. 232) calls Alternative Smith) decides the winner by instant runoff voting, so the cyclic triangle can be filled in using the IRV method shown further below.

- The Kemeny, Kemeny-Young, or minimax/maximin method[18] chooses the candidate with the smallest losing margin. To fill in the cyclic triangle, we must define

[18]Kemeny/Kemeny-Young is not generally equivalent to minimax/maximin, but they are equivalent for choosing a winner among three candidates. To see this, note that the decision rule for maximin at pg. 212 in Tideman (2017) is the same as the decision rule using matrix of majorities for Kemeny-Young at pg. 182, because in a three-candidate cycle each row of the matrix of majorities has only one negative entry.

and draw an additional set of tie lines that identify first-preference results where (given the pattern of second preferences) two losing margins are the same. These tie lines bisect the vertices of the cyclic triangle and meet at a point within it. I give an example for the Brexit poll in Figures 6 and 7.

3.2 Borda count and other positional methods

Given three candidates, a positional method gives each alternative 1 point for each ballot on which it is ranked first and $s \in [0, 1]$ points for every ballot on which it is ranked second; the winner is the candidate with the most points. Special cases include plurality ($s = 0$), Borda count ($s = 1/2$), and anti-plurality ($s = 1$).

A *positional tie line* for two candidates indicates first-preference shares where the two candidates have the same number of points, given s and the conditional second-preference shares. Assuming that an ix ballot (i.e. one that lists i only) awards 1 point to i and 0 to j and k,[19] candidates i and j tie when

$$v_i + s(v_{ji} + v_{ki}) = v_j + s(v_{ij} + v_{kj}),$$

which can be written in terms of first-preference shares, conditional second-preference shares, and s as

$$v_i(1 - sp_{ij}) = v_j(1 - sp_{ji}) + sv_k(p_{kj} - p_{ki}). \tag{2}$$

This describes the ij positional tie line. The first-preference win regions are defined by the intersection of positional tie lines.

To construct the diagram, we begin by identifying the point (y) where the positional tie line intersects the edge connecting the i vertex and the j vertex; at point y, i and j tie while k receives no first-preference support (i.e. $v_k = 0$). Substituting $v_k = 0$ and

[19] Alternatively one could assign s points to i and 0 to j and k (as suggested for Borda count by Saari (2008, pg. 197, fn. 6) and discussed in Emerson (2013)). Then the ij-positional tie line is where $v_{ij} + v_{ik} + s(v_{ji} + v_{ki} + v_{ix}) = v_{ji} + v_{jk} + s(v_{ij} + v_{kj} + v_{jx})$, which can be plotted in the same manner.

$v_j = 1 - v_i$ into Equation 2 and rearranging, we obtain

$$v_i = \frac{1 - sp_{ji}}{2 - s(p_{ij} + p_{ji})}. \tag{3}$$

The location of point y depends on how often j's supporters rank i second and vice versa. If ballots ranking j first never rank i second while ballots ranking i first always rank j second (i.e. $p_{ji} = 0$ and $p_{ij} = 1$), then $v_i = \frac{1}{2-s}$ at point y; thus for Borda count ($s = 1/2$) point y would be located where $v_i = 2/3$ and $v_j = 1/3$, which (assuming i is No Deal and j is Deal) is the point labeled y_1 on Figure 4.a. If by contrast ballots ranking j always rank i second while ballots ranking i first never rank j second (i.e. $p_{ji} = 1$ and $p_{ij} = 0$), then $v_i = \frac{s}{2-s}$ at point y; thus for Borda count ($s = 1/2$), point y would be located where $v_i = 1/3$ and $v_j = 2/3$, which is the point labeled y_2 on Figure 4.a.

Next we seek to identify point z, where the ij-positional tie intersects another edge of the ternary diagram. The location of this point depends on several conditional second-preference shares, but whether it lies on the ik edge or the ij edge depends on the sign of $p_{ki} - p_{kj}$: if k's supporters more often rank j second, then z lies on the edge between the i and k vertices; otherwise, z lies on the edge between the j and k vertices. More specifically, by substituting $v_j = 0$ and $v_k = 1 - v_i$ into Equation 2, we obtain

$$v_i = \frac{s(p_{kj} - p_{ki})}{1 - s(p_{ij} - p_{kj} + p_{ki})}, \tag{4}$$

which locates z in the case where k's supporters favor j; by substituting $v_i = 0$ and $v_k = 1 - v_j$ into Equation 2), we obtain

$$v_j = \frac{s(p_{ki} - p_{kj})}{1 - s(p_{ji} - p_{ki} + p_{kj})}, \tag{5}$$

which locates z in the case where k's supporters favor i. On Figure 4.a, z_1 indicates the extreme case for the Deal-No Deal tie line (under Borda count) where all Remain supporters rank Deal second, and z_2 indicates the other extreme case where all Remain

Figure 4: Tie lines and first-preference win regions for Borda count

4.a: Range of Deal-No Deal tie lines

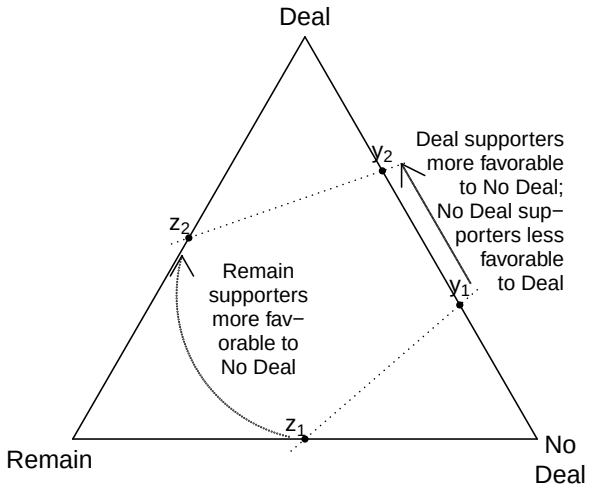

4.b: Actual Deal-No Deal tie line

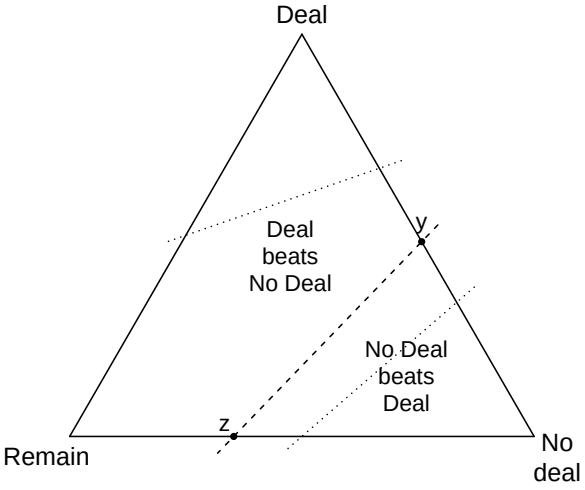

4.c: All tie lines

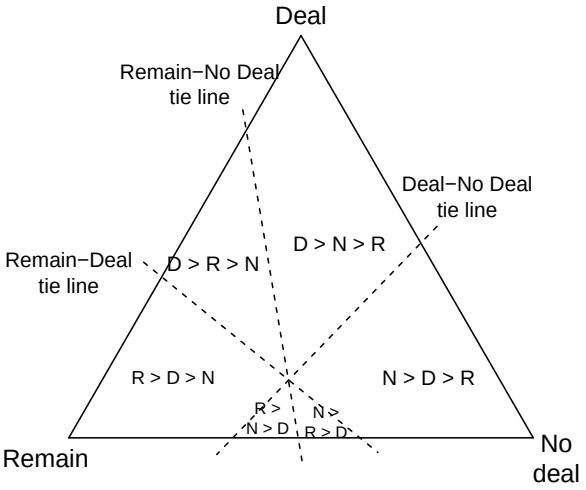

4.d: First-preference win regions

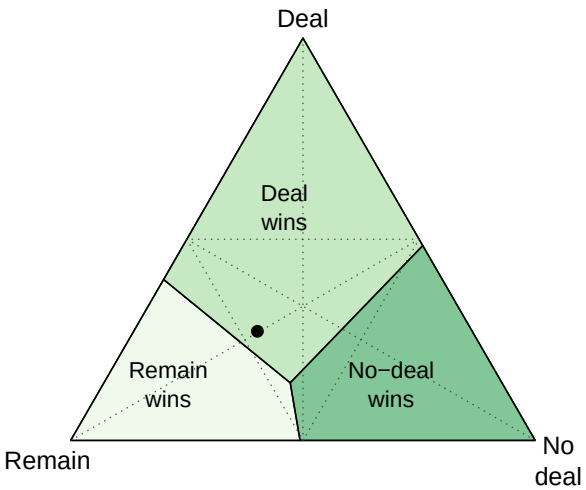

NOTE: A positional tie line indicates first-preference shares such that two candidates tie in a positional vote, given conditional second-preference shares and a relative value of second preferences. The location of the Deal-No Deal positional tie line for Borda count depends on several second-preference shares, but like the majority tie line it swings closer to the Deal vertex if Remain supporters are more favorable to No Deal (Figure 4.a). The actual Deal-No Deal tie line indicates Remain supporters' strong preference for Deal over No Deal (Figure 4.b). The three positional tie lines partition the ternary diagram, indicating first-preference shares producing each transitive order (Figure 4.c) and first-preference shares at which each alternative is the positional winner (Figure 4.d).

supporters rank No Deal second.[20]

Broadly, then, the angle of the ij-positional tie line can (like the ij-majority tie line) be thought of as a measure of how k's supporters feel about i and j: the more k's supporters favor i as their second choice, the more the ij-positional tie line tilts away from i's vertex. The connection is not as clear as in the majority tie line case, though, because the location and angle of the line also depends on the second preferences of i's and j's supporters.

Figure 4.b shows the actual Deal-No Deal tie line for Borda count (given conditional second preferences from the Nov. 2018 Brexit poll). Point y is located near the midpoint of the Deal-No Deal edge, indicating that supporters of Deal and No Deal are about equally likely to list each other's first choice second; point z is located close to the midpoint of the Remain-No Deal edge, indicating that supporters of Remain are far more likely to list Deal second than No Deal second.

Figure 4.c shows all three Borda tie lines. Unlike majority tie lines, positional tie lines always intersect at a point,[21] dividing the ternary diagram into six subtriangles (each corresponding to an order of finish).

Finally we obtain the first-preference win regions, shown in Figure 4.d. Because Deal attracts most second preferences from voters whose first preference is Remain or No Deal, Deal's first-preference win region is substantially larger than the other two. The first-preference result is located comfortably inside Deal's win region, indicating that Deal would win by Borda count.

3.3 Instant-runoff voting

In an instant-runoff voting election (e.g. for Australian lower-house seats),[22] voters submit ranked ballots and the lowest-scoring candidates are successively eliminated until only

[20] In the first case we also assume that No Deal supporters all rank Deal second, while in the second we assume that all Deal supporters rank No Deal second.

[21] The ij positional tie line intersects with the jk positional tie line where i, j, and k all have the same score, so the ik positional tie line must run through the same point.

[22] The system we describe is also referred to as ranked-choice voting (RCV) in the US, the alternative vote system (AV) in the UK, preferential voting in Australia, single-winner STV in Ireland, and occasionally the Hare system.

one remains. In the three-candidate case where elimination is based on first preferences (i.e. plurality votes), this simplifies to a procedure in which the candidate who receives the lowest first-preference share is eliminated and, of the remaining candidates, the winner is the one who is ranked higher on a larger share of ballots.

The IRV ternary diagram is built from a combination of positional tie lines (determining who is eliminated) and majority tie lines (determining which of the remaining candidates is the winner). We first use plurality tie lines to identify regions in which each candidate receives the lowest share of first-preference votes and is eliminated; these elimination regions are shown in Figure 5.a. Next, we superimpose the majority tie lines introduced above. In the region where a given candidate is eliminated, the majority tie line for the other two candidates determines which candidate wins; for example, in the region where Remain is eliminated, the Deal-No Deal majority tie line (singled out in Figure 3.b above) determines which candidate wins. To reflect this, in Figure 5.b I draw each majority tie line with a solid stroke in the region where it determines the winner (and a dashed stroke elsewhere), and I label each portion of the resulting figure according to how the non-eliminated candidates fare. The resulting first-preference win regions are shown in Figure 5.c. The irregular shape of these win regions reflects the combination of positional and majority tie lines that determine a winner in IRV, and will play a role in later illustrations of join-inconsistency and other voting system properties.

To develop intuition about the IRV diagram, consider Figure 5.d, which shows IRV first-preference win regions under a possible shift in conditional second preferences. (We will shortly see that a similar shift actually happened.) Recall that the majority tie line for alternatives i and j reflects the second-preference pattern of voters whose first preference is k. Suppose Deal voters become more favorable to Remain compared to No Deal (i.e. more likely to rank Remain second rather than No Deal); then we would see a clockwise tilt in the Remain-No Deal majority tie line, indicated by the rightward-pointing arrow at the boundary between the Remain-No Deal win regions in Figure 5.d. The other two arrows highlight a hypothetical shift (compared to Figure 5.c) of Remain voters toward Deal (indicated by the counterclockwise movement of the boundary between the Deal

Figure 5: IRV: elimination regions, results of final pairings, and first-preference win regions

5.a: Elimination regions

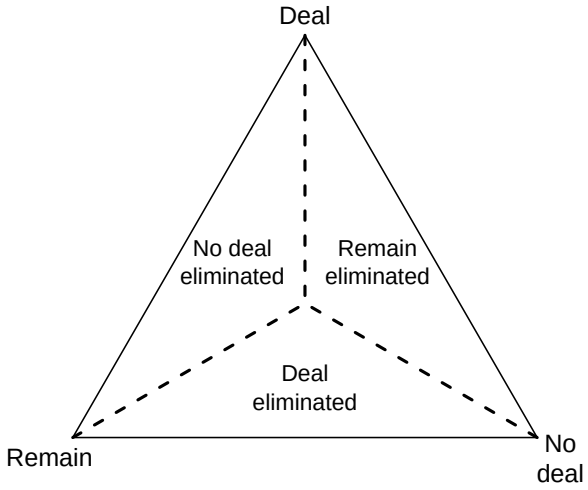

5.b: Results of final pairings

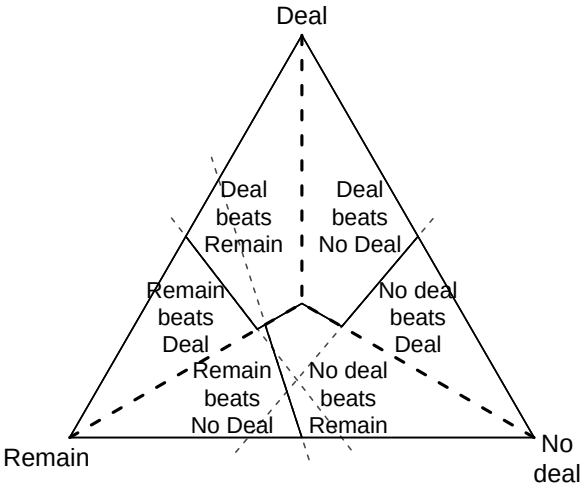

5.c: First-preference win regions

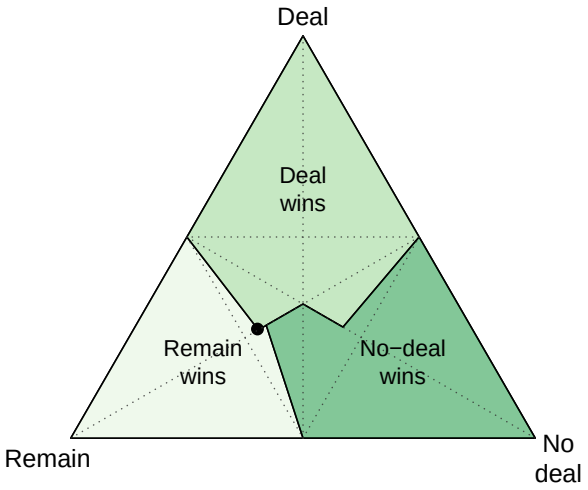

5.d: Possible change in second preferences

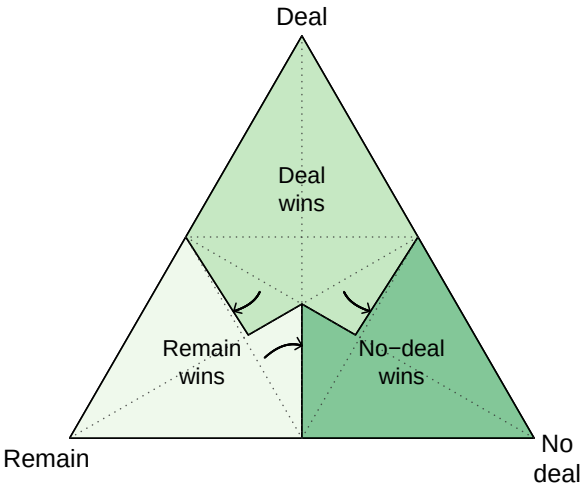

NOTE: In three-candidate IRV, one alternative is initially eliminated based on a positional method (in Figure 5.a, plurality). Then majority tie lines determine which of the remaining alternatives wins (Figure 5.b). Together, the positional tie lines and majority tie lines imply first-preference win regions with irregular shapes (Figure 5.c). Changes in conditional second-preference shares cause the majority tie lines to tilt and thus shift the first-preference win regions (Figure 5.d).

and No Deal win regions at the right of the figure) and of No Deal voters toward Deal (indicated by the clockwise movement of the boundary between the Deal and Remain win regions at the left of the figure).

Figure 5.c indicates that, given November 2018 Brexit preferences, the IRV winner would be either Remain or Deal. Deal and No Deal are effectively tied for second in first preferences, such that the first-preference result sits at the edge between the region where Deal is eliminated and the region where No Deal is eliminated. If Deal is eliminated, then Remain would defeat No Deal in the "runoff": the first-preference result is to the left of the Remain-No Deal majority tie line. If No Deal is eliminated, then Remain and Deal would effectively tie in a pairwise majority vote: the first-preference result essentially lies on the Remain-Deal majority tie line. Thus Remain could be helped by a hypothetical shift in first-preference shares that strengthened No Deal, leading to Deal being eliminated and Remain winning. In fact, such a shift in preferences occurred, as we will shortly see.

3.4 Other voting rules

We can depict further ordinal voting systems by combining and/or extending elements already introduced. Although I do not depict these below for reasons of space, consider:[23]

- In the Bucklin system, a candidate who wins a majority of top rankings is the winner; if there is no such candidate, the candidate who wins the largest share of first or second rankings is the winner. For the Bucklin diagram, begin by identifying regions where each candidate wins a majority of first preferences (each defined by a line connecting the midpoint of two edges) and fill in the remaining central area using anti-plurality win regions.

- The Coombs system is a variation of IRV in which elimination is based on the number of bottom rankings rather than top rankings. Specifically, a candidate who wins a majority of top rankings is the winner, but if there is no such candidate, the candidate who is ranked last on the most ballots is eliminated and the winner

[23] In each case I assume three candidates.

chosen based on rankings of the remaining two. To make the Coombs diagram, start with the regions where each candidate wins a majority of first preferences (as in Bucklin); to fill in the remaining central area, draw anti-plurality tie lines (positional tie lines with $s = 1$) to identify elimination regions and use majority tie lines to identify the winner within each elimination region.

- In the Nanson system the winner is determined by successive elimination of candidates who receive a below-average Borda score. For the Nanson diagram we must define a new tie line that identifies first-preference shares where, given second-preference shares, a candidate receives the average Borda score;[24] these tie lines partition the diagram into regions where either one or two candidates is eliminated. In regions where a single candidate k is eliminated, use the ij-majority tie line to determine which of the remaining candidates is the winner.

3.5 An update to the Brexit poll

A new Brexit poll was conducted in June, 2019, producing another set of estimated preference shares.[25] Table 2 reports the June, 2019, shares alongside the November, 2018, shares reported and illustrated above. Between November and June, first-preference support dropped for Deal and first-preference support for No Deal rose; also, preferences appear to have become more single-peaked, with Deal becoming the second choice for a larger proportion of voters ranking Remain or No Deal first.

Figure 6 displays the ternary diagrams for plurality, Borda, anti-plurality, Kemeny-Young, and IRV for both the November 2018 poll (left column) and the June 2019 update (right column). The plurality diagram shows Remain winning in both polls, with the black dot moving southeast between the two polls to reflect the shift in first-preference support from Deal to No Deal. In the Borda and anti-plurality diagram, the first-preference win region for Deal expands considerably from November 2018 to June 2019 (reflecting the

[24] The line where i receives the average positional score is given by $v_i + s(v_{ji} + v_{ki}) = \frac{1}{3}\left(v_i + s(v_{ji} + v_{ki}) + v_j + s(v_{ij} + v_{kj}) + v_k + s(v_{ik} + v_{jk})\right)$.

[25] Estimates provided in personal communication with Ben Lauderdale in September 2019. The client for whom he conducted the research approved the release of these aggregate estimates.

Table 2: Brexit preferences, November 2018 and June 2019

		Second choice					
		Nov, 2018			June, 2019		
		Remain	Deal	No Deal	Remain	Deal	No Deal
First choice	Remain (R)	-	.375	.087	-	.415	.049
	Deal (D)	.059	-	.212	.060	-	.162
	No Deal (N)	.038	.228	-	.002	.313	-

NOTE: Both estimates due to Ben Lauderdale. See footnotes 1 and 26.

increase in second-preference support for Deal as preferences became more single-peaked), such that even as Deal's first-preference support drops it wins more comfortably in June 2019 under both voting rules. Whereas Deal and Remain were essentially tied in the November 2018 poll in the Kemeny-Young (Condorcet) system, Deal wins in June 2019 (despite securing less first-preference support) because of the greater second-preference support from No Deal. In IRV, by contrast, the change in preferences was favorable to Remain. Recall that in November 2018 Deal and No Deal were essentially tied for second in first-preference shares; Remain would win if Deal were eliminated, but Remain and Deal would essentially tie if No Deal were eliminated. By June 2019, Deal finishes securely in last place in first-preference shares, after which Remain defeats No Deal in the IRV system. Thus not only do the different systems produce different winners with a given preference profile, but the same change in preferences can be beneficial to a given alternative in one system and harmful in another.

In Figure 7 we zoom in on the Kemeny-Young diagram for the Nov. 2018 Brexit poll from Figure 6. Where the majority tie lines form the boundary between two win regions they are drawn with a solid stroke; I use a dashed stroke to show the continuation of these lines where they do not form win region boundaries. Within the triangle defined by these dashed lines, there is a majority cycle where Remain beats Deal, Deal beats No Deal, and No Deal beats Remain. The Kemeny-Young rule says that the winner is the candidate who is defeated by the smallest margin in pairwise comparisons. To partition the majority cycle region according to this rule, we draw a new set of tie lines based on the margin of candidates' defeats. For example, the Remain-Deal tie line (which bisects

Figure 6: Outcomes given Nov. 2018 and June 2019 Brexit preferences

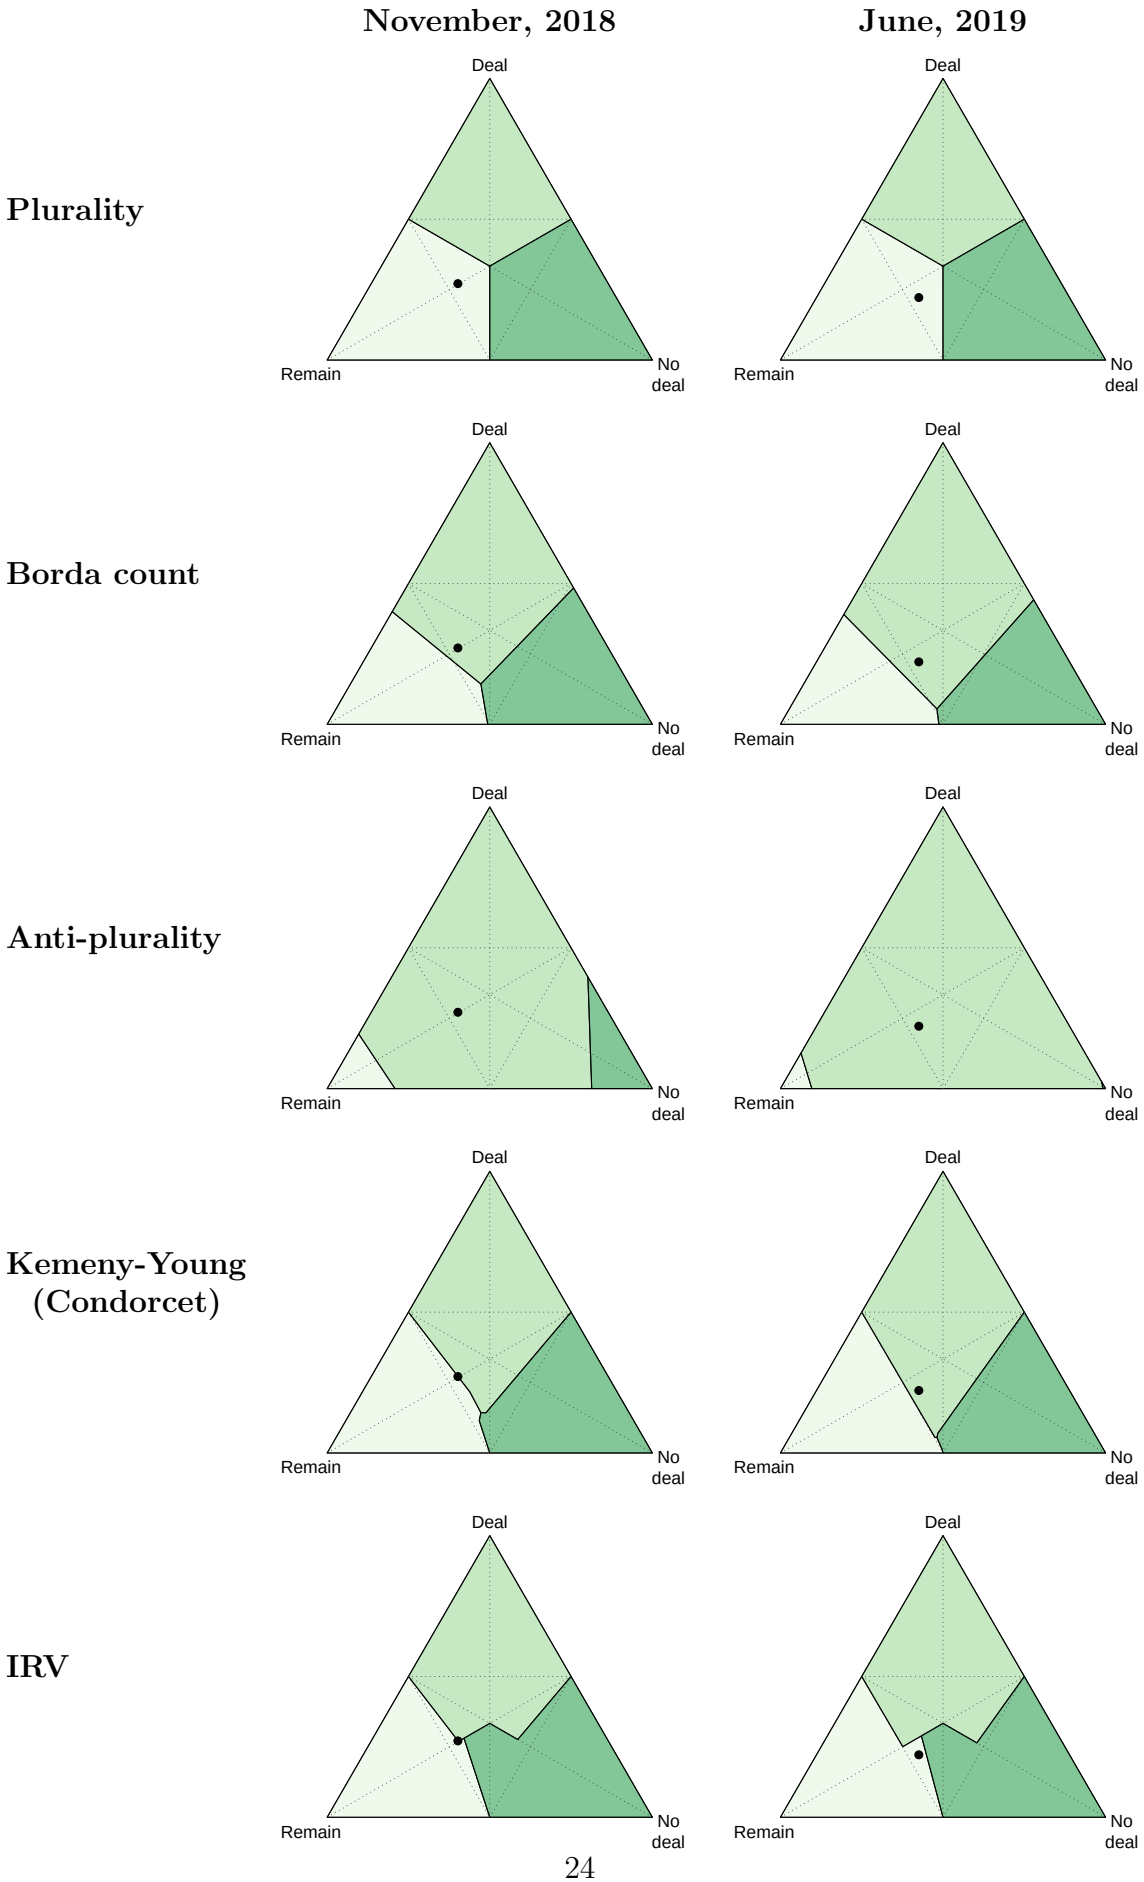

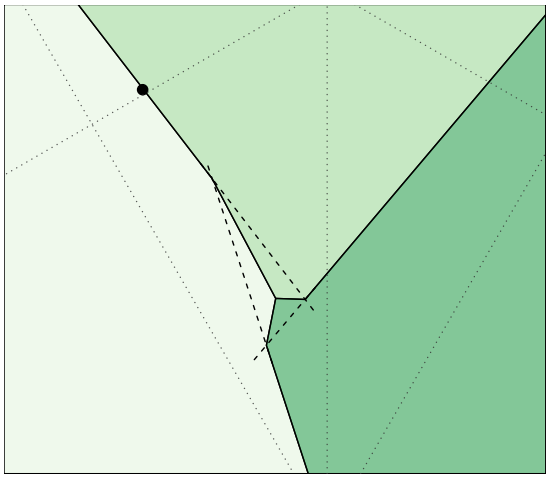

Figure 7: Closeup of Kemeny-Young win regions within the cyclic triangle for Nov. 2018 Brexit poll preferences.

(The full-size diagram is second from the bottom on the left in Figure 6.)

the uppermost vertex of the cyclic triangle) indicates first-preference shares where the margin by which Remain is defeated by No Deal is the same as the margin by which Deal is defeated by Remain. Each such tie line bisects an angle of the cyclic triangle, and the three lines meet at a point where the three candidates have the same margin of defeat.

3.6 Handling more than three alternatives

So far we have considered cases with only three alternatives, as is common in previous geometrical approaches to voting (e.g. Saari, 1994). The diagram can still be useful for many voting methods when there are more candidates. In particular, if there are three candidates a, b, c who would beat any other candidate in pairwise rankings, then any Condorcet method that resolves a cycle through an ordinal method applied to the Smith set can be diagrammed by considering only those three candidates, i.e. by eliminating candidates $d, e, \ldots$ from all ballots and considering only orderings of a, b, and c. Similarly, if there are three candidates a, b, and c who would be the last candidates standing in IRV or a similar elimination procedure, then we can eliminate all other candidates from the ballots and use the diagram to represent the final rounds of the competition. (Unfortunately the same approach does not work for positional methods: crucial information is lost if we reduce e.g. nine ranks to three.)

Table 3: Rankings in 2018 San Francisco mayoral election

Candidate	Number of times ranked first	second	third
London Breed	91,918	40,806	37,647
Mark Leno	61,276	75,927	33,063
Jane Kim	60,644	56,589	32,112
Angela Alioto	17,447	28,953	34,088
Ellen Lee Zhou	9,521	8,154	11,450
Richie Greenberg	7,016	6,949	8,434
Others	2,554	8,521	33,693

NOTE: Table reports the number of times each candidate was ranked first, second, and third in the 2018 SF mayoral election. Each of the top three candidates easily defeat all others in pairwise comparisons and qualify for the third-to-last IRV elimination round by a large margin, so it makes sense to focus on these candidates.

3.6.1 Example: 2018 mayoral election in San Francisco

The June 5, 2018, special mayoral election in San Francisco used a version of IRV in which voters could rank up to three candidates. Table 3 reports the number of times each candidate was ranked first, second, and third,[26] showing that three candidates (London Breed, Mark Leno, and Jane Kim) received the vast majority of first and second preferences. Assuming these three candidates are the last to be eliminated, we can ignore the other candidates and use the diagram to depict the final stages of the counting procedure.

The joint distribution of preferences over the three frontrunners is shown in Table 4. Note that the table has the same format as Table 1 except that a column is added to indicate ballots that rank only one candidate, which arises because voters were not required (indeed, not permitted) to rank all candidates. In fact, around a quarter of all ballots ranked *none* of the frontrunners and are thus eliminated entirely; of those that ranked at least one frontrunner, around a quarter ranked all three frontrunners, around half ranked ranked exactly two, and around a quarter ranked exactly one.

Using the proportions in Table 4 we can make a diagram exactly as described above. As shown in Figure 8, the results for Black's method (Condorcet with Borda to break cycles) and IRV happen to be very similar to the Brexit case. Breed narrowly wins both

[26] Analysis based on the final set of ballot images, which was made available June 28 2018 by the City of San Francisco Department of Elections website (link).

Table 4: First- and second-preferences within top three candidates in San Francisco mayoral election, June 2018

		Second choice			
		Breed	Leno	Kim	None
First choice	Breed	-	.149	.108	.176
	Leno	.083	-	.147	.059
	Kim	.056	.188	-	.034

NOTE: Table reports preferences over the top three candidates once all others have been eliminated from ballots. Because voters were permitted to rank only three candidates, many ballots did not rank any of the frontrunners (and are thus omitted from the table) and about a quarter of those remaining only ranked one.

Table 5: First- and second-preferences among top three candidates in Maine 2nd district congressional election, Nov 2018

		Second choice			
		Golden	Poliquin	Bond	None
First choice	Golden	-	.035	.236	.192
	Poliquin	.043	-	.104	.323
	Bond	.032	.014	-	.021

NOTE: Table reports preferences over the top three candidates in the Maine congressional election once the fourth candidate had been eliminated from ballots, based on ballot data made available on the Department of the Secretary of State's website on Nov. 15 2018 (link).

under Black's method and IRV.

3.6.2 Example: 2018 election in Maine's 2nd congressional district

The 2018 congressional election in Maine was the first use of IRV in a U.S. federal election. In Maine's 2nd district, two candidates (incumbent Bruce Poliquin and Jared Golden) won over 90% of first-preference votes (with Poliquin winning slightly over 46% and Golden winning slightly under 46%). Table 5 shows the preferences over the top three candidates once the fourth-placed candidate had been eliminated. Notably, over half of all ballots ranked only one candidate, including almost a third of ballots ranking Bond first.

The corresponding diagrams for Black's method and IRV appear in Figure 9. In both systems the result was essentially decided by the Golden-Poliquin majority tie line, which depends on Bond's second-preference shares; given Poliquin's narrow first-preference ad-

Figure 8: Condorcet and IRV outcomes in San Francisco mayoral election (top three candidates)

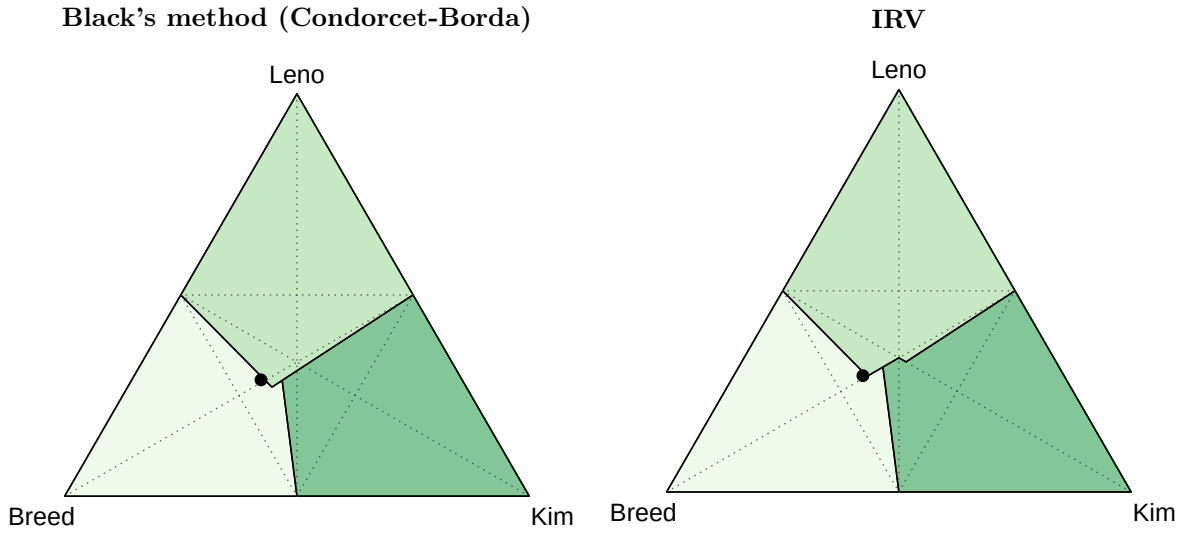

NOTE: The ternary diagrams depict the result of the San Francisco mayoral election using Black's method (left) and IRV (right, actually used). The black dot indicates the estimated first-preference shares; the shaded areas indicate the first-preference win regions given the estimated pattern of second preferences (conditional second-preference shares) and the voting system. Breed narrowly beats Leno in both systems.

vantage, Golden won because voters who ranked Bond first favored Golden as their second choice (by a two-to-one margin).

4 Illustrating voting system properties

To this point we have used the ternary diagram to represent the result of an election under different voting rules given a set of preferences over three (or more in some cases) candidates. In this section we show how the same diagram can be used to understand properties of voting rules.

4.1 Condorcet cycles

In the classic illustration of a Condorcet cycle (e.g. Arrow, 1951, pg. 3), three voters have the sincere preference orderings *abc*, *bca*, and *cab*, such that a beats b in a majority vote, b beats c, and c beats a (all by a 2-1 margin). The left panel of Figure 10 shows the Condorcet ternary diagram given the implied conditional second-preference shares in this

Figure 9: Condorcet and IRV outcomes in Maine's 2nd congressional district race (top three candidates)

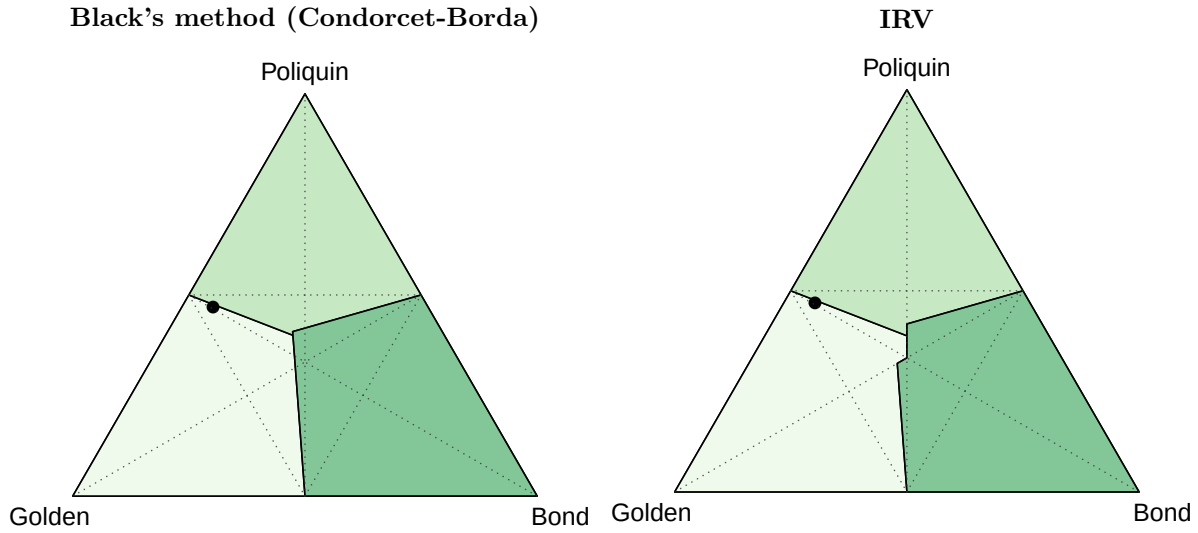

NOTE: The ternary diagrams depict the result of the Maine congressional election using Black's method (left) and IRV (right, actually used). The black dot indicates the estimated first-preference shares; the shaded areas indicate the first-preference win regions given the estimated pattern of second preferences (conditional second-preference shares) and the voting system. Golden narrowly beats Poliquin in both systems.

classic case ($p_{ab} = 1, p_{ba} = 0, p_{ca} = 1$).[27] With an equal number of voters supporting *abc*, *bca*, and *cab* (e.g. one each, as in the classic illustration), the first-preference result lies in the center of the diagram, producing a Condorcet cycle. As the diagram shows, there is actually a Condorcet cycle for *any* first-preference result given these conditional second-preference shares (i.e. any number of *abc*, *bca*, and *cab* voters) as long as no preference ordering has a majority of support.[28]

Our diagram can be used to identify conditions for a Condorcet cycle given any vector of conditional second-preference shares. For example, the right panel of Figure 10 shows the Condorcet winner given the conditional second-preference shares from the Nov. 2018 Brexit poll. The white triangle indicates first-preference shares where a Condorcet cycle occurs, given these conditional second-preference shares. Such a triangle exists unless the three majority tie lines happen to intersect at a single point which, given a large

[27] The diagram looks the same if the conditional second-preference shares are $p_{ab} = 0, p_{ba} = 1, p_{ca} = 0$, which produces the "reverse cycle".

[28] The left panel of Figure 10 resembles Figure 2.5.11 in Saari (1994, p. 105). In Saari's case, the vertices represent three out of the six preference orderings (chosen such that, if no other preference ordering has any weight and no single ordering has a majority, a pairwise cycle occurs).

Figure 10: Condorcet cycles

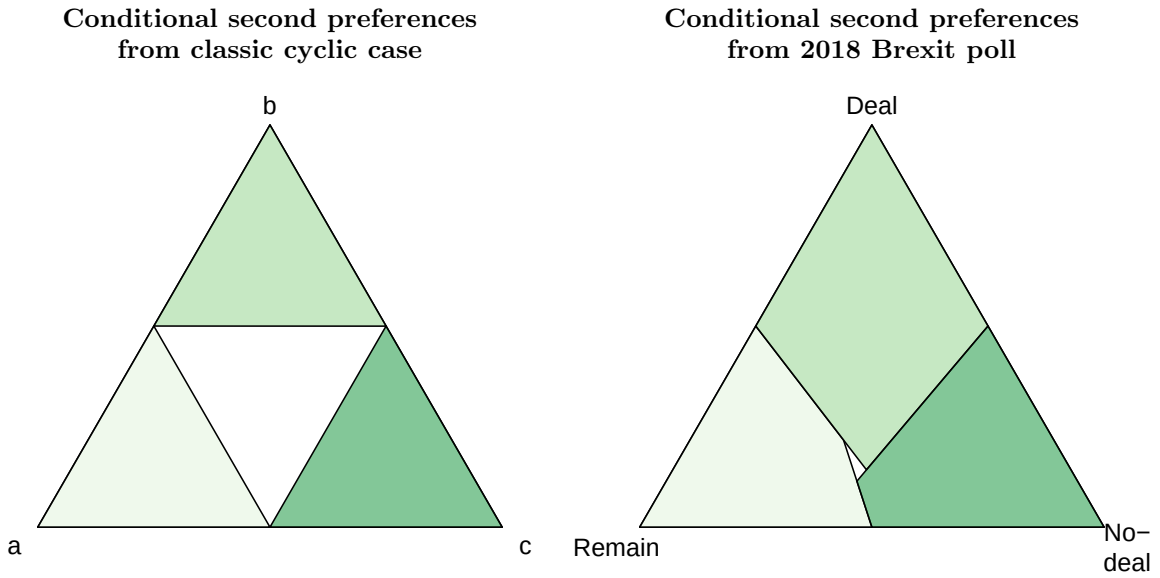

NOTE: When all voters have the preference ordering *abc*, *bac*, or *cab*, the first-preference win regions under Condorcet are indicated by the ternary diagram at left: there is a cycle (white region) unless one ordering has a majority. When conditional second-preference shares reflect the Nov. 2018 Brexit poll, the cyclic region is considerably smaller (right).

electorate and some uncertainty, is very unlikely to happen (although the cyclic triangle may be quite small).

4.2 Join-inconsistency, no-show paradox, and non-monotonicity

A voting method is *join-consistent* if, whenever an alternative i wins in both a subset of an electorate and the complement of that subset, i also wins in the entire electorate.[29] It follows that a voting method is join-consistent if, whenever i wins at vector $\mathbf{v}$ and vector $\mathbf{v}'$, it also wins at $\mathbf{v}'' = \alpha\mathbf{v} + (1-\alpha)\mathbf{v}''$ for all $\mathbf{v}$ and $\mathbf{v}'$ and all $\alpha \in [0,1]$. Now, suppose we plot the first-preference win regions for a particular set of conditional second-preference shares and a particular voting rule. It follows from the above that, if the voting rule is join-consistent, the first-preference win regions must be convex: any two points on the ternary diagram can be viewed as vectors $\mathbf{v}$ and $\mathbf{v}'$ (which happen to have the same conditional second-preference shares), so if a first-preference win region is not convex then there is a $\mathbf{v}$, $\mathbf{v}'$, and α such that one candidate wins at both $\mathbf{v}$ and $\mathbf{v}'$ but not at

[29] Saari (1994, Ch. 4) calls this criterion *weak consistency*.

$\mathbf{v}'' = \alpha\mathbf{v} + (1-\alpha)\mathbf{v}''$.[30] Figure 11 highlights the non-convexity of the first-preference win regions in IRV and Condorcet-IRV (i.e. the system where IRV is used in the event of a cycle) given the conditional second-preference shares from the Brexit survey. The winner is Deal at both x and y in the IRV diagram, but for part of the line connecting those points No Deal is the winner. Similarly, in the Condorcet-IRV case No Deal is the winner at points v (where there is a cycle) and w, but for part of the line connecting those points Deal is the winner. In both cases this shows join-inconsistency: a candidate could win in two sets of voters (which have identical conditional second-preference shares) but lose in the union of those sets.

Figure 11: Join-inconsistency given Nov. 2018 Brexit preferences

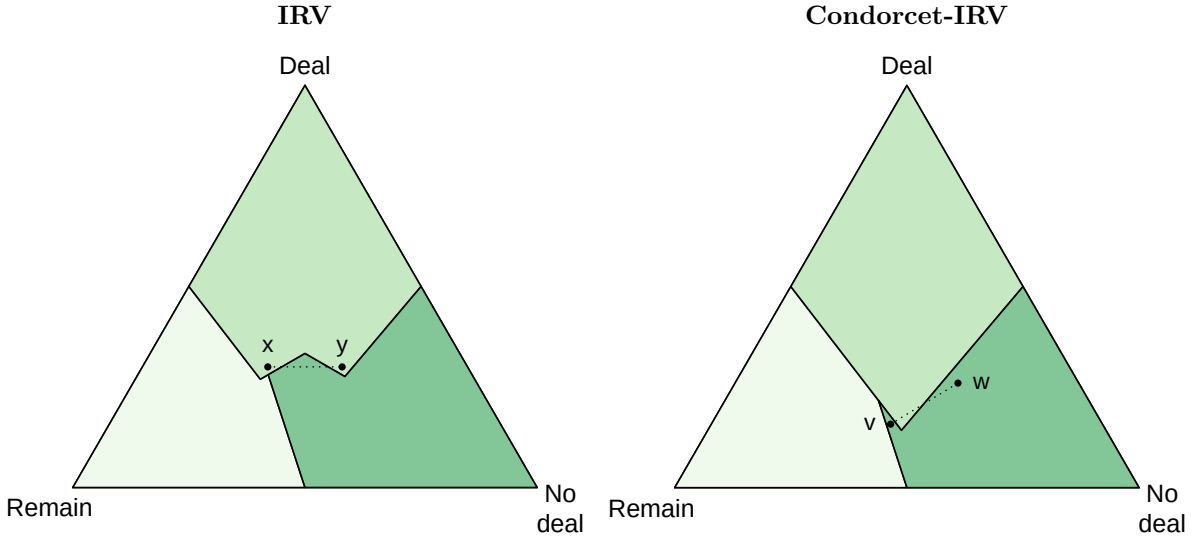

NOTE: If one candidate is the winner at two points on the ternary but not the winner somewhere along the line connecting those points, then the voting system is not join-consistent: an alternative could win in one part of the electorate and its complement but not in the whole electorate. In other words, non-convex first-preference win regions imply join-inconsistency.

It should be clear that while a single diagram can prove that a method is join-inconsistent, a single diagram cannot prove that a method is join-consistent (because it can only show one set of conditional second preferences), and it cannot show types of join-inconsistency in which the first-preference shares are the same in two profiles but the conditional second-preference shares differ (though one could try by superimposing two

[30] Saari (1994, Proposition 4.1.1 on p. 230) states this connection between join-consistency and convexity in profile space. He later proves the join-inconsistency of runoff methods with a diagram (Figure 4.1.6) that looks similar to mine, but has a different meaning: it shows the non-convexity of the region of point-share space in which a given alternative does not receive the lowest point share but could be the Condorcet winner.

Figure 12: No-show paradox for Remain supporters who rank No Deal last given Nov. 2018 Brexit preferences

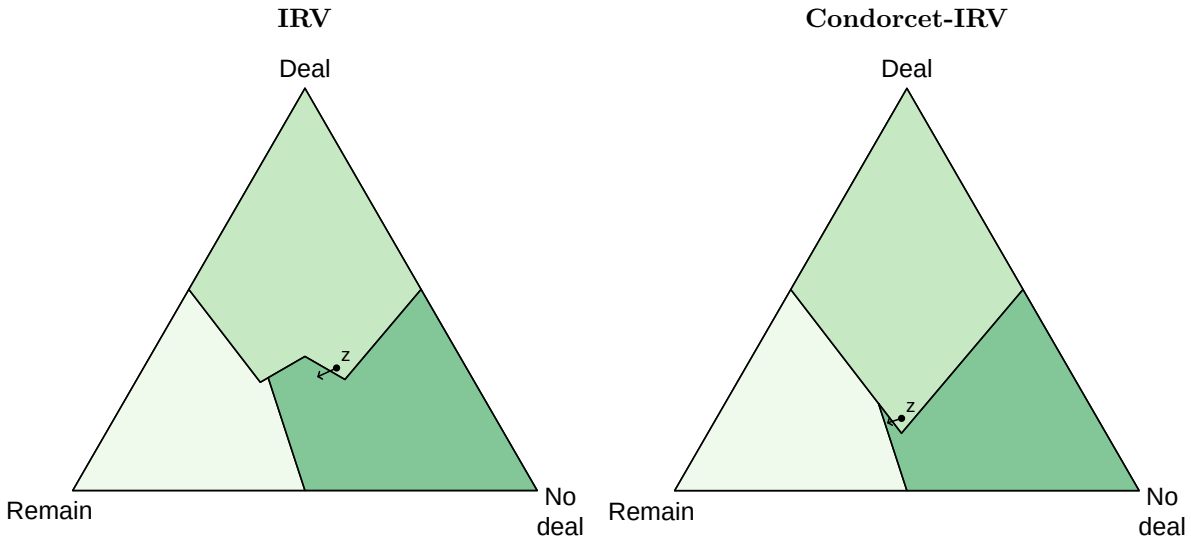

diagrams).

The no-show paradox (i.e. a violation of the "participation criterion" or "positive involvement" (Saari and Valognes, 1999)) occurs when a voter can obtain a better (expected) result by not voting than by voting sincerely. Figure 12 illustrates the no-show paradox using the second-preference shares from the November Brexit poll in IRV (left panel) and Condorcet-IRV (right panel). In both cases the no-show paradox occurs because, when the first-preference share is at point z, additional support for Remain moves the first-preference result down and to the left, changing the winner from Deal to No Deal;[31] thus a voter whose first choice is Remain and last choice is No Deal would be better off not voting than voting sincerely.[32] In IRV, additional support for Remain leads to Deal (the original winner) being eliminated rather than Remain, after which No Deal defeats Deal. In Condorcet-IRV, additional support for Remain leads to a cycle, which results in Deal being elected.

An ordinal voting system is non-monotonic if a winning candidate can become a losing candidate when one or more voters moves that candidate up in their ballot ranking.[33]

[31]The figure exaggerates the effect of one vote to make the direction clear. The analysis also ignores the complication of tie-breaking, but only to simplify.

[32]Adding an additional Remain-Deal-No Deal ballot also moves tie lines, but not the tie lines relevant to the no show paradox at point z.

[33]A voting system is also non-monotonic if a losing candidate can become a winner when one or more

Figure 13: Non-monotonicity in IRV given Nov. 2018 Brexit preferences

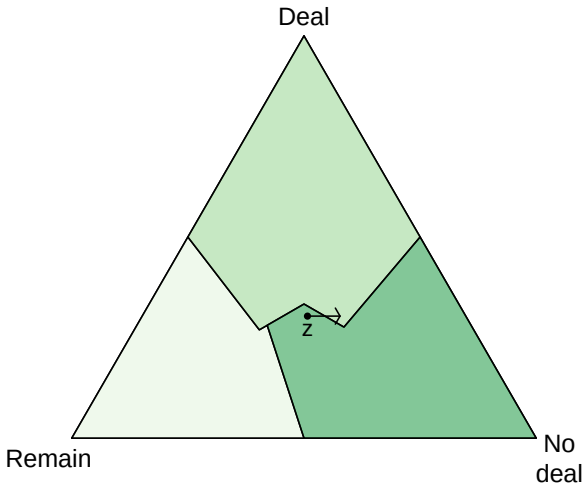

Figure 13 demonstrates the non-monotonicity of IRV. Suppose No Deal wins initially, at point z, but then a group of voters who originally ranked Remain first switch to ranking No Deal first and Remain second. To represent this shift on the diagram, we move the first-preference result rightward from point z.[34] But note that as we move rightward the winner shifts from No Deal to Deal, which shows that IRV is non-monotonic. The reason in this case is that while Deal was eliminated at z (clearing the way for No Deal to win), the shift in support from Remain to No Deal causes Remain to be eliminated instead, with the result that Deal defeats No Deal.

4.3 Susceptibility to strategic voting

A sincere vote in an ordinal system is one that reflects the voter's true preference ordering; an insincere vote deviates from that ordering. A voting system could be called susceptible to strategic voting to the extent that it rewards voters for submitting insincere votes, given some model of uncertainty. In terms of the ternary diagram, a single ballot can only make a difference when the election result is at or near the boundary between two first-preference win regions. To gain insight into the susceptibility of different systems

voters moves that candidate down.

[34]Majority tie lines would also shift, but if z is close to the Remain-Deal positional tie line these shifts will not affect the outcome.

to strategic voting, we ask which boundaries on the ternary diagram reward voters for submitting insincere ballots.

First, observe that when the result falls on a ij-positional tie line (where an increase in i's score or decrease in j score would elect i instead of j), there is always a type of voter who can elect their preferred candidate by switching to an insincere vote. For $s \in [0, 1)$, a voter whose sincere preference is $k \succ i \succ j$ could elect i by moving i up and submitting an ikj ballot, for example; for $s = 1$ (anti-plurality), a voter whose sincere preference is $i \succ j \succ k$ could elect i by moving j down and submitting an ikj ballot.

Next, observe that this is *not* true when the result falls on a ij-majority tie line (where an extra ballot ranking i above j would elect i instead of j). For a voter who prefers i over j, any ballot that ranks i over j has the same effect on the outcome (including a sincere ballot), so there is no incentive to submit an insincere ballot rather than a sincere one. If there is to be a strategic voting incentive at a majority tie line, it must be one where an ij-majority tie line separates win regions for alternatives other than i and j (e.g. due to a cycle).

Building on this insight, Figure 14 highlights first-preference win region boundaries (given the pattern of preferences from the November 2018 Brexit poll) where a voter could benefit from an insincere vote. All boundaries in the Borda count diagram are formed by positional tie lines (as they would be in plurality or anti-plurality), so they are all highlighted. By contrast, in IRV only the boundaries formed by positional tie lines are highlighted.

The right panel of Figure 14 shows the first-preference win regions for Condorcet-IRV. The only boundary where strategic voting may be rewarded lies along the upper-right edge of the cyclic triangle, where the Deal-Remain-majority tie line separates the win regions for Deal and No Deal. Along that edge, additional ballots ranking Deal above Remain make Deal the (Condorcet) winner, but additional ballots ranking Remain above Deal make No Deal the winner (by inducing a cycle); thus a voter whose sincere preference is Remain $\succ$ Deal $\succ$ No Deal would be better off insincerely ranking Deal above Remain to secure a victory for Deal, her second choice.

Figure 14: Results inducing insincere votes (given 2nd-pref. shares from Nov. 2018 Brexit poll)

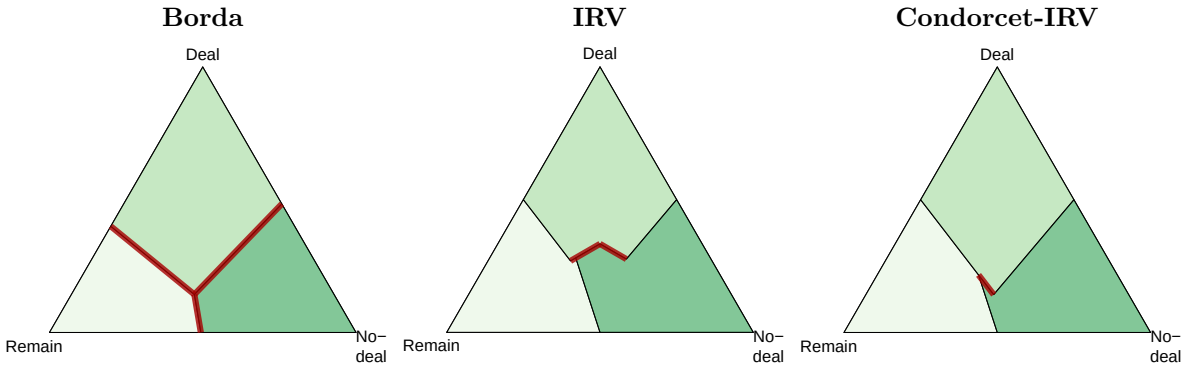

NOTE: In these diagrams (based on conditional second-preference shares from the Brexit poll), boundaries between first-preference win regions are highlighted if there is some voter who could benefit from an insincere vote. Because these boundaries are in different places in different systems, these systems' relative susceptibility to strategic voting depends on beliefs about likely outcomes. The fact that the extent of the highlighted boundaries goes down from left to right suggests that the systems' susceptibility to strategic voting may do so too.

The diagrams in Figure 14 point to two conclusions (one inarguable and one more speculative) about the susceptibility of voting methods to insincere voting.[35] The inarguable conclusion is that the susceptibility of these methods to strategic voting depends on beliefs about likely outcomes: Figure 14 shows that, conditional on a particular vector of conditional second-preference shares, there are first-preference results at which an insincere vote is rewarded in each method and neither of the others, which proves that one can specify beliefs about likely election outcomes such that each method is the most susceptible to insincere voting.[36]

The speculative conclusion from Figure 14 is that, for empirically relevant assumptions about voter preferences and beliefs about likely outcomes, the incentive to vote insincerely is higher in Borda (and more broadly, in positional methods) than in IRV and higher in IRV than in Condorcet methods. What evidence do we have for this conjecture? In Figure 14, the highlighted part of the ternary diagram (i.e. the measure of first-preference shares at which an insincere vote is rewarded) is largest in Borda and

[35]The diagrams look very similar for the other profiles examined in this paper; in each case the length of manipulable boundaries is longest in Borda and shortest in Condorcet-IRV.

[36]For example, if one expects Remain and Deal to receive almost all first-preference votes, then Borda rewards insincere voting from voters who rank No Deal second while IRV and Condorcet-IRV reward no insincere voting; if one expects a near three-way tie in first-preference shares then IRV induces insincere voting from voters who rank Remain first (especially those who rank No Deal last) while Borda and Condorcet-IRV induce no insincere voting.

smallest in Condorcet for one assumption about conditional second-preference shares. More generally, as shown above, all first-preference win region boundaries reward insincere votes in Borda but only some do in IRV and Condorcet; also, the highlighted regions of the Borda figure reward pushing one's second choice down, while the highlighted regions of the IRV figure do not. Finally, we have seen that an insincere vote may be rewarded in Condorcet methods only around (and possibly in) the triangle where a cycle takes place. Together, these observations point to our speculative conclusion but (as shown in the previous paragraph) such a conclusion necessarily depends on beliefs about likely outcomes (both about first-preference shares given conditional second-preference shares and about conditional second-preference shares). Testing this conjecture therefore requires specifying empirically relevant assumptions about voter preferences and beliefs and is the subject for another paper.

5 Conclusion

Because a diagram can sometimes make intuitive what otherwise is difficult to understand, the approach presented in this paper may help researchers and election analysts analyze elections held under a variety of single-winner voting systems. The diagrams may also be useful for communicating the results of elections and polls to the public. Given that the election systems studied in this paper depend on five or more preference frequencies, it should not be surprising that they are more difficult to explain and visualize than a plurality contest. This complexity may make it difficult for voters to understand and accept the results of unfamiliar systems, regardless of whatever advantages these systems may have.[37] It may also hamper voters' efforts to predict likely outcomes from polls and respond appropriately. The diagrams in this paper provide a compact summary of an election or poll result in a way that, while perhaps not widely accessible in the details, communicates the big picture rather simply: candidates with more top rankings generally

[37] In a possible attempt to benefit from this confusion, the campaign spokesman for the loser of Maine's first IRV election called it a "black-box voting system" reliant on "artificial intelligence" as his candidate appealed the results and demanded a recount. Emily Kopp, "Defeated Rep. Bruce Poliquin Calls for Lengthy Ranked Choice Recount", *Roll Call*, 27 Nov. 2018 (link accessed Jan. 3, 2019).

win, but the outcome also depends on lower rankings in a way that varies with the voting method. Perhaps these diagrams can help, then, not just in the analysis of voting systems but in their adoption and operation.

References

Arrow, Kenneth J. 1951. *Social choice and individual values*. Yale University Press.

Dunleavy, Patrick and Rekha Diwakar. 2013. "Analysing multiparty competition in plurality rule elections." *Party Politics* 19(6):855–886.

Emerson, Peter. 2013. "The original Borda count and partial voting." *Social Choice and Welfare* 40(2):353–358.

Farrell, David M and Ian McAllister. 2006. *The Australian electoral system: origins, variations, and consequences*. UNSW Press.

Grofman, Bernard, Alessandro Chiaramonte, Roberto D'Alimonte and Scott L Feld. 2004. "Comparing and contrasting the uses of two graphical tools for displaying patterns of multiparty competition: Nagayama diagrams and simplex representations." *Party politics* 10(3):273–299.

Ibbetson, D. 1965. "Comment on Berrington, 'The general election of 1964'." *Journal of the Royal Statistical Society. Series A (General)* 128(1):54–55.

Kastellec, Jonathan P and Eduardo L Leoni. 2007. "Using graphs instead of tables in political science." *Perspectives on politics* 5(4):755–771.

Katz, Jonathan N and Gary King. 1999. "A statistical model for multiparty electoral data." *American Political Science Review* 93(1):15–32.

Nurmi, Hannu. 1999. *Voting paradoxes and how to deal with them*. Springer Science & Business Media.

Saari, Donald G. 1994. *Geometry of Voting*. Springer-Verlag.

Saari, Donald G. 2008. *Disposing dictators, demystifying voting paradoxes*. Cambridge University Press.

Saari, Donald G. 2011. Geometry of voting. In *Handbook of social choice and welfare*, ed. Kenneth J Arrow, Amartya Sen and Kotaro Suzumura. Vol. 2 Elsevier chapter 27, pp. 897–945.

Saari, Donald G and Fabrice Valognes. 1999. "The geometry of Black's single peakedness and related conditions." *Journal of Mathematical Economics* 32(4):429–456.

Tideman, Nicolaus. 2017. *Collective decisions and voting: the potential for public choice*. Routledge.

Tufte, Edward. 2001. *The visual display of quantitative information*. 2 ed. Graphics Press USA.