Proportionality for ranked voting, in theory and practice

AUTHORS REDACTED

Classical social choice theory includes a long list of criteria, or fairness axioms, for elections where individuals rank their preferences. Famous impossibility theorems from the 1970s concern the properties of voting rules to convert profiles of ranked preferences to winner sets. But though public perceptions of fairness are strongly keyed to proportional representation, notions of proportionality are strikingly missing from the standard roster of fairness axioms. We design a framework to measure *the degree of proportionality of seats to voter preference* under a wide class of systems for electing legislative bodies, even when elections are conducted without party labels. We begin by building out a set of generative models for creating synthetic ranked preference profiles, with an emphasis on flexibility and realism; in particular, we can efficiently generate polarized elections with properties motivated by the extensive body of work on racially polarized voting in the United States. The models use notions of *blocs* of voters and their *slates* of preferred candidates, which need not be known to voters but could be implicit in their voting patterns. The models serve as a thought tool for building a new definition of proportional representation and provide a framework that allows researchers to compare systems of election in terms of their tendency to produce proportional outcomes. We illustrate this by giving both empirical and theoretical results for single transferable vote (STV) elections.

This work brings a statistical modeling toolkit to the questions around ranked choice voting and proportionality. At the same time, it builds a much-needed bridge from computational social choice theory to political science, where degrees of proportionality have been intensely studied for well over a century, and to the work of practitioners in current reform efforts around voting rights and democracy.

ACM Reference Format:
Authors Redacted. 2024. Proportionality for ranked voting, in theory and practice. 1, 1 (April 2024), 58 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

Author's address: Authors Redacted.

ACM XXXX-XXXX/2024/4-ART
https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

In this paper, we give what we believe to be the first definition of *the degree of proportionality of votes to seats* that is general enough for use with ranked preferences.[1] This fills a gap in the classical social choice literature. Ken Arrow's foundational work studied social choice functions that combine multiple input rankings into one output ranking; following this, a series of important results were conjectured and proved from the 1960s to the 1990s concerning the use of rankings to output winner sets. Impossibility theorems of Müller–Satterthwaite, Gibbard–Satterthwaite, and Duggan–Schwartz rule out the viability for single-winner or multi-winner elections of simultaneously securing multiple axioms of fairness (see, for instance, [Taylor, 2002]). Examples of fairness axioms from early social choice theory include strategy-proofness, monotonicity, and the Condorcet criterion. However, these simply do not rank high in the public discourse around democracy.

Another area of need in the computational social choice literature is in defining generative models of election using domain knowledge of real-world electoral dynamics. We construct novel generative models of ranking that are inspired by polarized elections in real-world settings; in particular, voting rights law in the United States has used notions of voting blocs and their degrees of cohesiveness for decades. With these models and data, we can test voting rules on both real and synthetic preference profiles, yielding information—some provable and analytic and some qualitative and simulation-based—on whether roughly proportional outcomes do indeed tend to result from so-called "semi-proportional" systems.

1.1 Contributions

New generative models. Generative models of voting use parameters and data—in our case, historical voting patterns, demographics, cohesion parameters, and candidate strength—to build a probability distribution from which ballots are sampled and elections can be simulated. In this paper we build and test generative models. These are the first mechanisms for producing ranked ballots that incorporate polarization according to candidate slates.

Rethinking proportionality. The proportionality of representation for a subgroup of voters could have a very simple interpretation in demographic terms (the group's seat share is in line with its share of the electorate). However, this fails to account for any complexity in the voting patterns of that group and the complementary voters. We define a framework that replaces demographic proportionality for a bloc of voters with *support proportionality* for a slate of candidates: the slate's seat share should be in line with the combined support for its candidates. We note that this kind of proportional representation is broader than that of PR systems such as party list voting, which secure support proportionality—on the basis of party only—by construction, so that the finding of proportional outcomes is vacuous in that setting. Here, we are measuring a kind of proportionality that is endogenous or emergent with respects to votes cast, and can be measured not only on the basis of party but with respect to any other cohesive preference.

Incorporating domain knowledge. This project engages domain knowledge in voting rights law and practice in multiple ways. First, we shift the definition of voter cohesion to match the ordinary and legal use of the term. In the social choice literature, definitions of *cohesive* groups of voters tend to revolve around overlapping approval ballots: for instance, Sánchez-Fernández et al. [2017] call a group of voters ℓ–cohesive, where n candidates are running for k seats, if they comprise at least $\ell n/k$ people and their preferences overlap in at least ℓ candidates. This nuances earlier notions in which "cohesion" requires only a non-empty overlap in approvals. By contrast, this paper introduces notions of cohesiveness keyed to the probability of members of a group to support

[1] All other notions we are aware of work by recourse to approval ballots, as we describe further below.

candidates from a certain slate. Compare this to, for instance, the landmark *Thornburg v. Gingles* decision of the U.S. Supreme Court, requiring Voting Rights Act plaintiffs to ascertain "whether members of a minority group constitute a politically cohesive unit" by measuring whether "a significant number of minority group members usually vote for the same candidates."[2] Expert work supporting a finding of cohesiveness revolves around "statistical evidence of voting patterns" using past elections, and polarization is typically summarized by using standard inference techniques to estimate the share of support for slates of candidates by blocs of voters [Hebert et al., 2010]. The authors of the present paper are drawing on just this kind of experience in voting rights expert work.[3]

Secondly, definitions related to justified representation are far from notions of proportionality in the political science literature and the popular vernacular: seat share in line with vote share. The relationship of seat share to vote share has been intensely studied at least since the late 19th century, and measurement of deviation from ideal seats/votes curves has generated a significant literature in the last fifty years especially in the work of Tufte, King, Grofman, and many more.

Finally, our use of ranked ballots rather than approval ballots is aligned with practice (and reform momentum) in the United States and internationally. Several U.S. states have recently debated adoption of ranked choice elections: Maine and Alaska now use ranked voting for statewide elections, with Nevada midway through the process of enacting a shift. Dozens of cities from San Francisco to Minneapolis use ranked choice for municipal elections, and New York City recently switched to ranked choice to elect city councillors and the mayor. Outside of the U.S., ranked choice voting is used for local or legislative elections in much of the Anglophone world—including Scotland, Ireland, New Zealand, and Australia—as well as for parliamentary elections in Malta and Papua New Guinea.

Illustrating with STV. While our notion of proportionality and the generative models we propose do not rely on a specific voting rule, we will use *single transferable vote* (STV) as a test case. STV is a family of voting rules within ranked choice voting, using a transfer mechanism for selection of multiple winners, where the number of seats to be filled in a single contest is called the *magnitude*. In STV elections, there is a threshold level of support needed to be elected—typically the threshold is about $1/(k+1)$ of the first-place votes, where k is the magnitude. The election is conducted in rounds. As candidates are either elected (by passing the threshold) or eliminated from contention, the (surplus) votes supporting those candidates are transferred to the next options on their respective ballots.[4] We note that *instant runoff voting* or IRV, an extremely popular alternative in practice, is the same voting rule as STV in the special case $k = 1$.

Though STV is the basis for the examples in this paper, the express goal of the work is to set up a framework suitable for the comparative study of any voting rules.[5]

Important note: links to data and code used to produce the examples and data visualization in this paper have been suppressed for anonymization purposes, and will be provided in a full replication repo after the paper is reviewed.

[2] *Thornburg v. Gingles* (1986), https://www.oyez.org/cases/1985/83-1968.

[3] For instance, consider recent expert work in Texas: minority racial groups were estimated to collectively support Democratic candidates in general elections from 2012–2020 at rates of 85-92%, while white voters supported Republican candidates at rates of 75-85% in the same contests. Expert report of REDACTED, *TX NAACP et al. v. Abbott*, Case No. 1:21-CV-00943-RP-JES-JVB.

[4] Specific mechanics vary; in this paper we have implemented the vote-tallying mechanism used by Cambridge, MA for its City Council elections, except as noted below.

[5] When single-winner rules like IRV are used to elect a representative body, as in the New York City Council, the framework here will be applicable.

1.2 Related work

Statistical ranking models, or models that assign a probability to permutations on a set of elements, have been studied at least since the early 20th century, going back to Thurstone [1927]. Subsequent models include those introduced by Bradley and Terry [1952] and Plackett [1975], which form the basis for the BT and PL models in this paper, respectively. Benter [2008] introduced a variation of the Plackett model with a dampening parameter to account for less careful deliberation of lower-ranked items. Johnson et al. [2002] proposed a model to combine rankings that were determined by several different sources—which could have used different methods and criteria—into an aggregate, or meta, ranking scheme.

Ranking models have been used in a variety of applications in the broader social science literature. Stern [1990] apply the methods to horse races, where the marginal probability of each horse finishing first is known in advance. Bradlow and Fader [2001] apply time series models to Billboard "Hot 100" list, to show how song rankings change over time. Graves et al. [2003] apply a combination of ranking models to racecar competition outcomes. In the area of election analysis, Upton and Brook [1975] fit a Plackett model to ranked ballots in London elections to determine the effect of candidate name ordering on the ballots, also known as positional bias. Gormley and Murphy [2008] fit a combination of Plackett-Luce and Benter models to polling data from Irish elections in 1997 and 2002. In particular, they find the models to be effective in identifying voting blocs (groups of voters with similar ranked preferences) within the electorate. In the same paper, the authors fit mixtures of Plackett-Luce models to cast vote records from Irish elections, with the main goal of identifying blocs within the electorate.[6] These analyses are descriptive, based on historical data. In a recent paper, Garg et al. [2022] model outcomes of elections in multi-member Congressional districts under a solid coalition assumption, which means that the ballots are effectively unranked (and do not differentiate candidates within each coalition).

Our work is related in several respects to the existing computational social choice literature. There is a large body of work on the axiomatic properties of voting rules in various settings, including notions with a family resemblance to proportionality. For example, defining (extended) justified representation (JR) [Aziz et al., 2017] allows certain guarantees in approval-based multi-winner voting: sufficiently large groups whose approvals have non-trivial overlap can't be shut out of the winner set. Refer to Lackner and Skowron [2022] for a more thorough discussion. Various papers have used proportionality language for functions that map approval ballots to ranked outcomes [Skowron et al., 2017] and, quite recently, for functions that carry ranked ballots to sets of approval ballots, and from there map to multi-winner outcomes [Brill and Peters, 2023]. While similar in spirit, it would be difficult to compare ideas invoking justified representation to ours directly because the JR family of axioms relies on a fundamentally different definition of cohesion.

In terms of generative models of election, numerical experiments in this literature traditionally rely on assumptions of *impartial culture* [Pritchard and Wilson, 2009], under which voters are independent and every permutation of candidates is equally likely, *impartial anonymous culture*, in which Lebesgue measure is used to set relative preferences, or use *spatial* or distance-based models [Elkind et al., 2017, Tideman and Plassmann, 2010]. See Szufa et al. [2022, 2020] for a comparison of common statistical cultures and recent discussion of how to sample approval elections.

Spatial models [Enelow and Hinich, 1984], in particular, which represent voters (and candidates) as ideal points in a metric space—in other words, using a space with a distance function as the latent space for voter preferences—are common across fields. Voters are presumed to vote either

[6] In the language that will be introduced below, this roughly corresponds to fitting a Name-PL model (see Remark 4) with unknown group sizes and no slate structure. That is, their method is designed to learn preferences for all candidates by each of two blocs. Fitting a mixture model in this way does not produce a canonical division of candidates into slates.

deterministically for their closest representatives or probabilistically (upweighting closer candidates) [Burden, 1997]. Two commonly used methods for estimating ideal points (typically from Congressional roll-call data) are NOMINATE [Poole and Rosenthal, 1985] and IDEAL [Clinton et al., 2004]. Ranked choice voting models can be built from spatial models. For example, Gormley and Murphy [2007] combine a spatial and Plackett-Luce model to analyze Irish STV elections (discussed further in §5), and Kilgour et al. [2020] use a spatial model (where voters rank by proximity) to measure the effect of ballot truncation on single-winner ranked choice outcomes. Garg et al. [2022] also use a spatial model in one section, with voter ideal points extracted from ideology ratings in a commercial voter file, to relate the "diversity" of elected officials to the sizes of multimember districts.

Spatial models on one hand, and approval votes on the other, are favored by the mathematically inclined because they lend themselves to provable theoretical properties of voting rules. For example, under the implicit utilitarian voting framework, ordinal votes are proxies for underlying utilities and the *distortion* of a voting rule captures its worst-case loss compared to having full information [Procaccia and Rosenschein, 2006]. Anshelevich et al. [2018] study the distortion of STV under metric preferences, and Gkatzelis et al. [2020] recently settled a well-known conjecture on the optimal metric distortion when aggregating rankings to elect a single winner.

Our goal is to strike out in a new direction, with definitions that enable new questions to surface.

2 BLOCS, SLATES, AND PROPORTIONALITY

2.1 Defining blocs, slates, and notions of preference

The concept of blocs and slates is straightforward: *slates* are disjoint sets of candidates, such that voter propensity to support the various slates can be measured. The idea that voters would exhibit a preference among slates makes sense for an electorate overall, or when split out into disjoint groups of voters we call *blocs*.

To make this precise, we must delineate what it means for the preference profile consisting of ranked votes from a group of voters to display an overall preference for one group of candidates over another. We list several notions of preference or propensity that can be measured in an observed vote profile—that is, these are measurements that can be made on any cast vote record that has been minimally cleaned so that each ballot is a partial ranking (a permutation of a subset of the candidates).

Definition 2.1. Suppose an election is conducted with bloc structure $(A, \mathcal{A}, B, \mathcal{B})$ consisting of sets of voters A, B and corresponding slates of candidates $\mathcal{A} = \{A_1, \dots, A_r\}$ and $\mathcal{B} = \{B_1, \dots, B_s\}$. We adopt the viewpoint of bloc B, which may be the whole electorate (the $A = \emptyset$ case). Suppose voters are allowed to rank up to $n \leq r + s$ candidates on their ballots—that is, ballots may be incomplete rankings of varying length, up to some maximum.

- Bloc B prefers slate $\mathcal{B}$ with *first-place preference* p_B if the share of first-place votes in the profile for $\mathcal{B}$ candidates is p_B.
- Bloc B prefers slate $\mathcal{B}$ with *positional preference* $P_B = (p_1, p_2, \dots, p_n)$ if the share of ballots placing an $\mathcal{B}$ candidate in position i (among those for which a vote is cast and neither slate was exhausted in the higher positions) is p_i. In particular, the special case of *consistent positional preference* p_B corresponds to $P_B = (p_B, p_B, \dots, p_B)$.
- Given a positional scoring rule with weights $(w_1, w_2, \dots, w_n)$, we say that B prefers slate $\mathcal{B}$ with *score preference* p_B if the share of their score for $\mathcal{B}$ candidates is p_B. The default option will be to give standard Borda weights to the top k ranks via the score vector $(k, k-1, \dots, 1, 0, \dots, 0)$ in a magnitude-k election; we will refer to this as (top-k) *Borda*

preference. For the purpose of Borda scoring, incomplete ballots are completed with an averaging convention (see A).

Preferences for the A bloc are defined analogously; the only difficulty in extending to *more* than two blocs is one of cumbersome notation.

We will interpret each of these preference parameters as an indication of how *cohesive* bloc B is, with higher preference parameters (closer to 1) indicating more strongly aligned blocs.

Example 2.2. Suppose an election has been conducted with $r = 3, s = 2, n = 5$ (i.e., complete rankings are allowed), and suppose the voters are labeled as A voters or B voters. Suppose that the summarized preference profile for the B bloc is given by

$\times 2$	$\times 3$	$\times 8$	$\times 1$	$\times 5$	$\times 3$	$\times 5$
B_1	B_1	B_1	A_1	B_2	B_2	B_1
B_2	A_2	B_2	B_1	B_1	A_3	B_2
A_1	B_2	A_2	B_2	A_1	A_1	
A_2	A_3	A_1		A_3	B_2	
A_3	A_2			A_2	A_2	

(by name)

i.e.,

$\times 7$	$\times 3$	$\times 8$	$\times 1$	$\times 3$	$\times 5$
B	B	B	A	B	B
B	A	B	B	A	B
A	B	A	B	A	
A	A	A		B	
A	A			A	

(by slate)

Then the first-place preference of the B bloc for $\mathcal{B}$ candidates is $26/27$, the positional preference is $(\frac{26}{27}, \frac{21}{27}, \frac{4}{7}, \frac{3}{3}, -)$, the Borda preference to all five places is $232/405$ with ballot completion, and the top-2 Borda preference is $73/81$. Note that the last few positional scores are $4/7$, $3/3$, and undefined—rather than $4/22$, $3/21$, and 0—because of only considering ballots which have not exhausted the B candidates.

2.2 Defining proportionality

If the electorate is undivided ($A = \emptyset$) and the voters support slate $\mathcal{B}$ with propensity π_B, then we interpret that as voters giving the slate π_B share of their support. In this case, the proportionality ideal is extraordinarily simple: seat share equals vote share, i.e.,

$$S_B = \pi_B.$$

When there are two distinct blocs with different voting behavior that partition the whole electorate, this extends by convex combination to a natural heuristic for a proportional outcome of an election. If π_B is the preference parameter for bloc B towards its candidates and likewise π_A for bloc A, then a natural target is to have the seat share S_B for the $\mathcal{B}$ slate satisfy

$$S_B = N_B \cdot \pi_B + (1 - N_B)(1 - \pi_A),$$

where N_B is the share of voters from the B bloc. That is, the combined support for $\mathcal{B}$ candidates is the size of the B bloc times its level of cohesion (the propensity to vote for $\mathcal{B}$ candidates) plus the size of the complementary bloc times its level of crossover voting (again, the propensity to vote for $\mathcal{B}$ candidates).[7]

This enables us to say, for instance, whether a particular election outcome was near-proportional (in a given bloc structure, if applicable) with respect to first-place preferences, or to Borda preferences, or any other notion of propensity. Proportionality is not a foregone conclusion for ranked choice voting even in the extremely simple case where the blocs are defined by first-place votes; lower-ranked choices may or may not track closely with first-place preference.

[7]One could consider alternative definitions of proportionality, for example, based on a bloc-weighted combination of the number of seats a slate wins in each of the hypothetical elections in which only one of the blocs participates. However, this requires fixing a voting rule. We deliberately propose a notion of proportionality that is agnostic to the choice of voting rule.

election	(r, s, k)	first place pref. π_B	first place pref. proportionality	top-k Borda share π_B	top-k Borda share proportionality	STV outcome
North Ayrshire 2022 Ward 1	$(8, 4, 5)$	0.17	0.87 seats	0.24	1.19 seats	0 seats
Angus 2012 Ward 8	$(4, 2, 4)$	0.24	0.96 seats	0.26	1.02 seats	1 seat
Clackmannanshire 2012 Ward 2	$(5, 3, 4)$	0.32	1.27 seats	0.31	1.25 seats	1 seat
Aberdeen 2022 Ward 12	$(7, 3, 4)$	0.31	1.26 seats	0.36	1.42 seats	1 seat
Aberdeen 2017 Ward 12	$(6, 4, 4)$	0.33	1.33 seats	0.41	1.63 seats	1 seat
Falkirk 2017 Ward 6	$(3, 3, 4)$	0.34	1.35 seats	0.43	1.71 seats	2 seats
Renfrewshire 2017 Ward 1	$(4, 4, 4)$	0.37	1.49 seats	0.46	1.84 seats	1 seat
Fife 2022 Ward 21	$(4, 4, 4)$	0.46	1.86 seats	0.51	2.02 seats	2 seats
Glasgow 2012 Ward 16	$(7, 5, 4)$	0.60	2.40 seats	0.58	2.32 seats	3 seats

Table 1. Here, s is the number of mainstream-left candidates (defined by membership in the Green, Liberal Democrat, and Labour parties), r is the number of candidates from all other parties, and k is the number of seats to be filled in the election. We treat the electorate as a single bloc (undivided) and measure π_B as the level of first-place support for the $\mathcal{B}$ slate, and the share of top-k Borda scores, respectively, for the $\mathcal{B}$ slate.

Example 2.3. We use a sample of nine real-world Scottish local government STV elections to illustrate how to use the definition in practice. We will use the simplified slate structure where three Scottish parties—Green, Liberal Democrat, and Labour—are defined as a "mainstream left" slate $\mathcal{B}$, and slate $\mathcal{A}$ combines all other parties.[8] In Table 1 we consider the level of proportionality in two ways. We first use first-place preference to define π_B, the propensity of voters to support slate $\mathcal{B}$. (Equivalently, we can think of this as defining blocs by first-place vote and adopting 100% cohesion.) This means that the number of seats needed to achieve (first-place) proportionality is $\pi_B \cdot k$, the proportional seat share times the number of seats. Applying an alternative choice of propensity, we can use π_B to measure the slate-$\mathcal{B}$ share of the top-k Borda scores to arrive at a different proportionality target.

3 GENERATIVE MODELS

3.1 Constructing the models

In this section, we set up generative models of election, including several variants derived from classical statistical ranking literature in the style of Plackett-Luce and Bradley-Terry models.[9] Though at first the by-name and by-slate versions may seem extremely similar, we find that Slate-PL and Slate-BT have several desirable properties compared to Name-PL and Name-BT. These, together with a model called the Cambridge sampler (Slate-CS), make up the generative models explored in the empirical work in this paper.

[8]The other parties include the centrist Scottish National Party, parties defined by their stance on independence from the UK, right-wing parties like the National Front, and some farther-left socialist parties. STV can be tested with respect to any slate, though securing proportionality by first-place preference will be most plausible when the slate has cohesive support from a subset of voters.

[9]Earlier versions of the Name-PL, Name-BT, and Slate-CS models have been discussed in unpublished work by an overlapping collection of authors. References are suppressed here for anonymization purposes.

Definition 3.1. For all of the models below, assume a fixed bloc structure $(A, \mathcal{A}, B, \mathcal{B})$ with $\mathcal{A} = (A_1, \dots, A_r)$ and $\mathcal{B} = (B_1, \dots, B_s)$, allowing the possibility that $A = \emptyset$ as before.

A *ballot* is a partial or complete ranking of the $r + s$ candidates and a *ballot type* is a partial or complete permutation of the symbols $A^r B^s$, i.e., a simplified ballot that treats the candidates of each slate as indistinguishable from each other.

The models below will use the following parameters to generate a profile for bloc B:

Cohesion Tendency of the bloc to support slate $\mathcal{B}$, given as a parameter $\pi_B \leq 1$ (typically required to be at least $1/2$ in the multi-bloc case).

Strength Tendency of bloc B to agree on preferred candidates *within* each slate. This consists of probability vectors $I_{BA} = (a_1, \dots, a_r)$ and $I_{BB} = (b_1, \dots, b_s)$.

We can combine the cohesion and strength data into a single probability vector

$$I_B = \big((1 - \pi_B)a_1, \dots, (1 - \pi_B)a_r, \quad (\pi_B)b_1, \dots, (\pi_B)b_s\big).$$

Using these components, we can define five generative models as follows. The first two work directly with ballots, while the latter three first construct ballot types. These are analogous to the profile by name and the profile by slate in Example 2.2.

Name-PL Plackett-Luce by name: Each B-bloc voter chooses candidate i to be ranked first with probability $I_B(i)$. They continue to select candidates for lower-ranked positions in order, at each stage selecting candidate j with probability proportional to $I_B(j)$. In other words, each voter samples their ballot without replacement from all candidates proportional to their weighting in I_B.

Name-BT Bradley-Terry by name: The probability that a B voter casts a ballot σ is proportional to

$$\prod_{i <_\sigma j} \frac{I_B(i)}{I_B(i) + I_B(j)},$$

where $i <_\sigma j$ means that i is ranked before (i.e., higher than) j in σ. In other words, for each pairwise comparison of candidates, we introduce a term for the likelihood of ranking one before the other coming from the relative weights in I_B.

Slate-PL Plackett-Luce by slate: Each B-bloc voter chooses between the symbol A and B in the ith position with probability π_B of choosing B, as long as both $\mathcal{A}$ candidates and $\mathcal{B}$ candidates remain available. Once a slate is exhausted, the rest of the complete ranking is filled in with the remaining symbol.

Slate-BT Bradley-Terry by slate: Suppose a ballot type σ is a permutation of $A^r B^s$, that is, an ordered list containing r A symbols and s B symbols. Suppose that out of the rs comparisons of the instances of A with the instances of B, the A occurs earlier than the B a total of $0 \leq i \leq rs$ times. The probability that a B voter casts this ballot is proportional to $(1 - \pi_B)^i (\pi_B)^{rs-i}$.

Slate-CS Cambridge sampler: We draw from a dataset consisting of ten years of ranked votes from city council elections in Cambridge, MA. Historical candidates have been labeled as white (W) or as people of color (C), with help from local organizers. To use this model, we make a choice to designate bloc $\mathcal{B}$ as corresponding to voters who put a W candidate first ($B = W$), or who put a C candidate first ($B = C$). We use the cohesion parameter π_B to decide probabilistically whether the voter chooses their own slate or the other slate in the first position. Then we complete the ballot type by drawing with weight proportional to frequency from the cast ballots with that header.

In all three Slate models, we must then assign candidate names to the symbols A and B. We do so by drawing without replacement (Plackett-Luce style) from I_{BA} and I_{BB} separately to order $\mathcal{A}$ and $\mathcal{B}$, then fill in names accordingly.

Remark 1 (Names versus slates). *It turns out to be an important distinction to work directly with the names or to create a type first, then add names. The reason for the divergence is that the Slate models handle I_{BA} and I_{BB} separately; concatenating them into I_B before making length comparisons yields unintended results, such as a highly cohesive bloc whose voters tend to put their strong candidate first and then immediately cross over to supporting the opposite slate. These effects can be explored in the supplementary plots (§B) which compare all five models.*

Remark 2 (About the Cambridge data). *Cambridge, Massachusetts uses STV for its city council and school board elections and has done so since 1941. Our source of Cambridge historical data is city council elections to fill $k = 9$ seats by STV from 2009 to 2017, coded by candidate race as described above; there are frequently 20 or more candidates who run in each contest. If a ballot type is selected from the historical frequency histogram that has more candidates from a given slate than the (r, s) chosen for a given simulation run allows, then we ignore further instances. For instance, a ballot type of AAABB in an election where $r = s = 2$ will be read as AABB.*

One valuable aspect of our use of Cambridge historical data in the present study is that it lets us incorporate realistic short-ballot voting behavior without a proliferation of extra parameters. For instance, Cambridge voters cast "bullet votes" (listing only one candidate and leaving other positions blank) 7501 times out of 87,914 ballots cast in our data set, and this will be reflected in the ballots generated by the CS model. However, a serious limitation is that we have coded the candidates by race, while Cambridge city council politics are likely more polarized by other candidate features—for instance, an explicit slate of affordable housing candidates is routinely advertised before election day and is highly salient to voter behavior. Nevertheless, race is a candidate feature often apparent to voters which allows us to observe naturalistic patterns of alternation in voting.

These give new generative models to study. As noted in the literature review (§1.2), many authors have considered only solid bloc voting (the "solid coalitions" assumption), in which every A voter casts a ballot of type $AA \dots ABB \dots B$. Others have used extremely stylized assumptions like Impartial Culture, Impartial Anonymous Culture, and spatial voting. These new models greatly expand on the generative models in the literature, and they do so in a manner that allows voting rights experts in the United States to plug in standard cohesion parameters for majority and minority groups as the π_A, π_B. We will give a brief demonstration of their flexibility below in §3.2.2.

Remark 3 (PL preferences). *Slate-PL with $(A, \mathcal{A}, B, \mathcal{B})$ and any cohesion and candidate strength parameters produces blocs with consistent positional preference π_B (respectively π_A) for their own slates, and therefore with first-place preference π_B (or π_A) as well.*

Remark 4 (Mixture models). *The definitions above are in terms of specified blocs of voters with different voting preferences. However, there is a strong connection to mixture models suggested by the structure here. In a mixture model, each voter is assigned independently to a class, and then randomly submits a ballot based on the parameters for that class. More precisely, if N_1 and N_2 are the weights for two different classes of voter with $N_1 + N_2 = 1$, and μ_1 and μ_2 are two distributions on ballots corresponding to the two classes, the probability of a ballot σ is*

$$\mu(\sigma) = N_1\mu_1(\sigma) + N_2\mu_2(\sigma).$$

As the number of voters increases, the fraction of voters assigned to each class converges to N_1 and N_2 respectively; for large numbers of voters we can therefore consider the size of each class to be predetermined and treat voters as if they belong to two blocs of fixed size.

In particular, since it considers pairwise probabilities, the BT model with two blocs resembles a mixture of Mallows models. It differs in allowing swaps to be weighted by preference between slates rather than by their position in the ranking.

3.2 Visualization

3.2.1 MDS plot of vote profiles. One difficulty in studying ranked choice elections is that, unlike oversimplified Example 2.2, real-world elections frequently have too many valid ballots possible to effectively see the full preference profile. For instance, an election with six candidates can be thought of as having 1236 possible ballots to cast—there are 6! complete rankings and a roughly equal number of partial rankings.[10] Thinking of profiles as distributions over valid ballots allows us to define natural notions of distance between profiles, such as the L^1 distance between profiles given by the sum over possible ballots of the absolute value of the difference of shares for that ballot. (Up to a constant factor, this is the same as the total variation distance of distributions.) With this notion we can visualize differences between the generative models as we vary parameters.

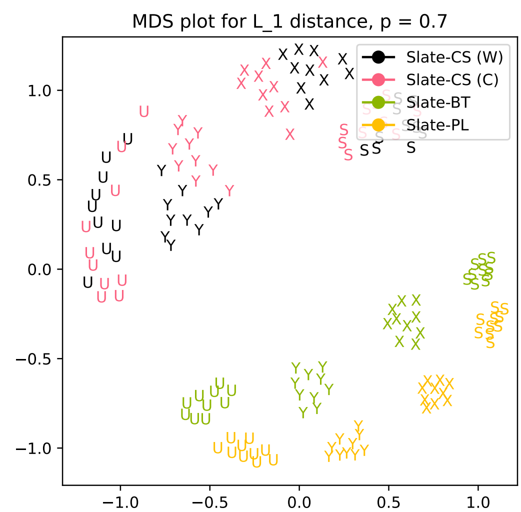

Fig. 1. Multi-dimensional scaling (MDS) plot for one-bloc profiles with $r = s = 3$ (3 candidates per slate), under a variety of generative models and candidate strength scenarios. Each model is designated by a different color, and the candidate strength scenarios are denoted U, S, X, Y, as described in the text. The pairwise distances between profiles are computed with L^1 distance on the distributions. Each preference profile has 1000 ballots, and we have generated 10 profiles by each of the 16 model/strength pairs. Note: it is not surprising that CS profiles, both when $B = W$ and $B = C$, fall far from PL and BT profiles, because PL and BT always generate complete rankings, while CS uses real historical data that includes many partial rankings. This observation can be used to give a sense of scale for the distances in the plot.

To illustrate the importance of candidate strength, we introduce four out of the infinitely many variations on I_B concerning the preferences of B-bloc voters.

- **U** (uniform-uniform): preferences are uniform over $\mathcal{A}$ candidates and uniform over $\mathcal{B}$ candidates, so within each slate any two candidates have a 1 : 1 ratio of support.
- **S** (strong-strong): preferences are strong over both slates, namely with one candidate receiving 10 times the support of all the others, who are equal.
- **X** (uniform-strong): uniform support for $\mathcal{A}$ candidates and 10:1:1 support for $\mathcal{B}$ candidates;
- **Y** (strong-uniform): the reverse.

[10]Here, we identify a ballot of length 5 with a complete ranking of length 6, since the last-place candidate is implicit.

In the multi-dimensional scaling (MDS) plot in Figure 1, the first-place preference for $\mathcal{B}$ candidates is $p_B = .7$; Supplemental Figure 38 shows how the outputs vary in p. In this plot, we can see some systematic differences and similarities.[11] For instance, strength scenarios Y and X interpolate between U and S, as we might have expected. Also, BT profiles resemble both kinds of Cambridge outputs more than PL profiles do, though the reason for this is far less clear.

3.2.2 Validation on Scottish elections. A benefit of using parameterized generative models is the possibility of fitting to real-world elections. Though we leave a full-bore fitting effort to future work, this section shows the potential of this approach to match the observed non-solidity of coalitions.

To this end, we define a *swap distance* between two ballot types, partial or complete. For complete ballots, this counts the smallest number of swaps of neighboring symbols necessary to transform one ballot type into the other; for instance, $\text{dist}(AABBB, ABBAB) = 2$. See §A for a discussion of efficiently measuring this distance, including an extension to partial or weakly ranked ballots.

Using swap distance, we can investigate the extent to which vote profiles deviate from the solid coalition assumption. Let us return to the nine Scottish elections and the mainstream-left slate $\mathcal{B}$ discussed above. For every ballot cast in the election we can compute its distance from the solid A-over-B ballot type A^sB^r. (Note that a solid vote of the opposite kind looks like B^rA^s, lying at distance rs from its reverse.) For the Aberdeen Ward 12 and Falkirk Ward 6 elections from 2017, these distances are summarized in Figure 2.

Next, we can attempt to match these histograms using the generative models in §3.1. We can accomplish interesting results even with an undivided electorate (one bloc). We choose our cohesion parameter by optimizing π_B to minimize d_{Wass} to the observed election.

The resulting distance distributions are visualized in Figure 2 (and see Supplement D for a full range of outputs). Note that the traditional assumption of solid coalitions produces distributions that are point masses at distances 0 and rs, which clearly have little in common with the real-world ballot distributions. Both visually and in terms of measured Wasserstein distance, the models do well at matching observed patterns of non-solidity of coalitions.[12]

[11]The reader should recall that MDS plots are simply low-distortion planar embeddings, which depend on a choice of random seed. The x and y axes have no meaning; only the relative pairwise distances are meaningful with respect to the data. We have verified that the structure of the plots stays the same for a few choices of random seed.

[12]To get a sense of scale, note that shifting the entire distribution by one bin would give $d_{\text{Wass}} = 1$.

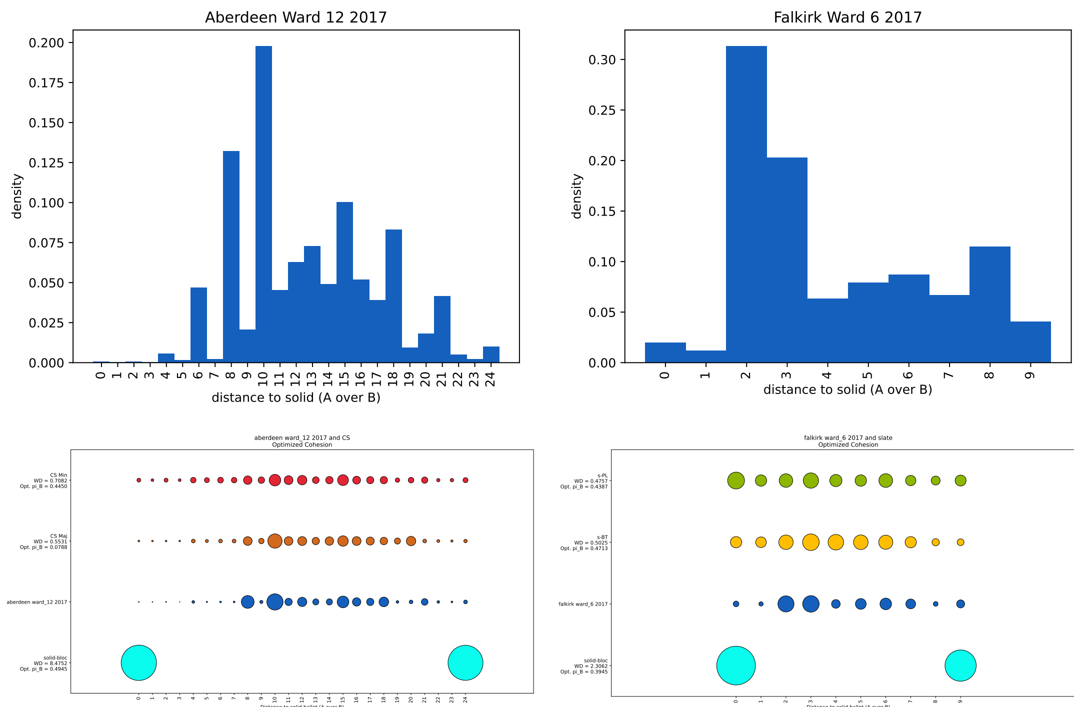

Fig. 2. Top: Histograms showing the distribution of swap distances to solid A-over-B type in Aberdeen Ward 12 and Falkirk Ward 6, 2017. Bottom: Bubble plots showing the distribution of swap distances for model outputs compared to the same two elections. (The dark blue row is the observed election, for which the data exactly repeats the conventional histograms.) The area of each circle is proportional to frequency.

3.2.3 Parameter interactions. Next, we leverage the generative models in combination with a voting rule to produce simulations that highlight complex interactive effects between model parameters.

We vary N_B over $\{.1, .2, .3, .4\}$ and we vary both π_A and π_B. We have selected four candidate strength scenarios for two blocs (compare the one-bloc scenarios in §3.2.1); these are chosen to give a small window on how powerfully candidate strength can interact with other factors.

- **UU** both blocs have uniformly random preference order over each slate;
- **UX**: I_{BB} has a strong (10:1:1) candidate while others are uniform;
- **XX-same**: A and B blocs strongly prefer the same $\mathcal{B}$ candidate;
- **XX-diff**: A and B blocs strongly prefer different $\mathcal{B}$ candidates.

In effect, this makes a 5-tuple of choices for each batch of runs: model, strength scenario, population share, cohesion for A voters, and cohesion for B voters. We then generate a batch of profiles from each tuple (100 for the 3-seat case and 25 for the 6-seat case) to place each symbol on the plot. The x-axis position is the combined support level for $\mathcal{B}$ candidates, given by $N_B \cdot \pi_B + (1 - N_B)(1 - \pi_A)$ as above—so a given support level can be achieved in many different ways. The y-axis position is the average number of seats won by $\mathcal{B}$ candidates when the batch of profiles is run through the STV voting rule.

If the proportionality ideal were hit exactly, the symbols would all fall on the main diagonal. The proportionality target rounded up and down to whole numbers of seats is shown with dotted lines

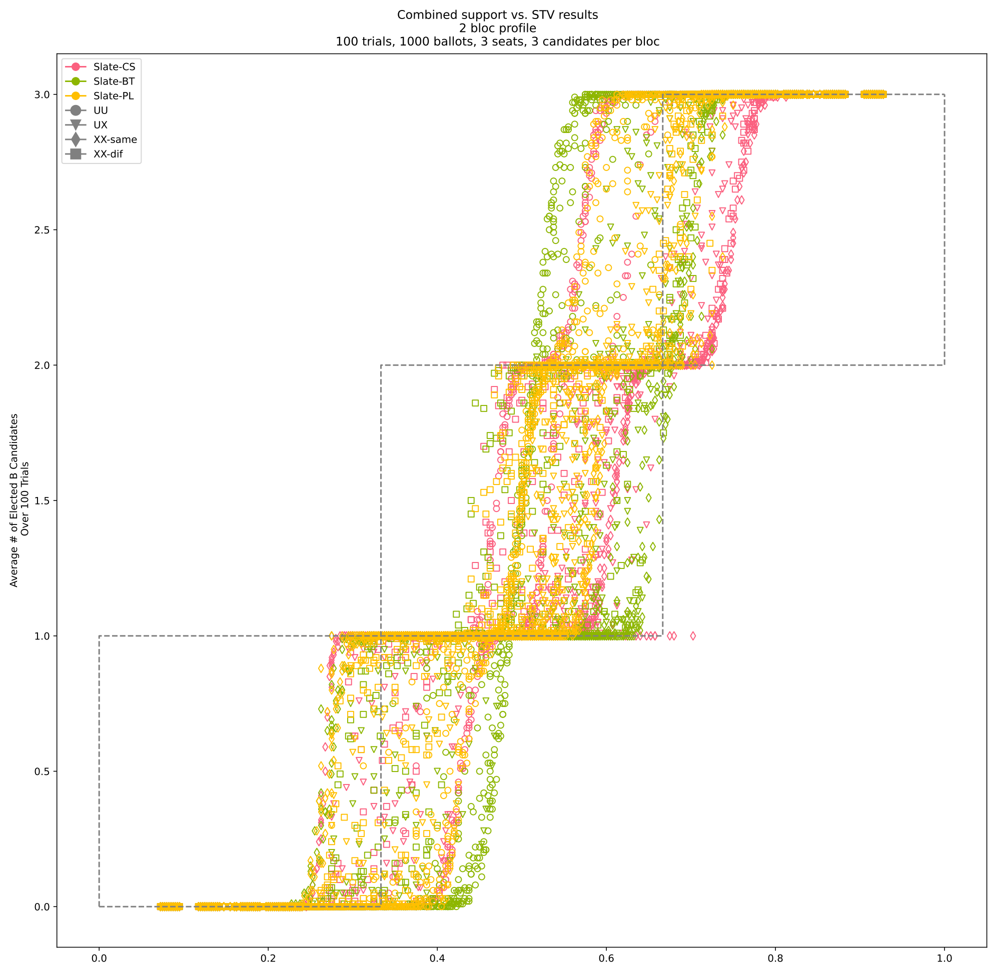

Fig. 3. Setting $(r, s, k) = (3, 3, 3)$, we independently vary the B proportion of the electorate, the generative model, the A and B cohesion, and the candidate strength settings. In this visualization, we have run 100 trials for each parameter tuple, recording the number of $\mathcal{B}$ candidates elected for each simulated profile. The x axis position is the combined support for $\mathcal{B}$ (with respect to first-place votes) and the y-axis position records the average number of seats over the trials with each tuple of parameters. The dashed lines show the proportionality target rounded up and down to the nearest whole number of seats.

in the plots. Instead of symbols falling squarely in these targets, Figures 3 and 4 show an intriguing "winner's bonus"—support shares away from 50% can get amplified seat shares through STV. And the effects are starkly different depending on candidate strength, with a very high slope observed when Bradley-Terry ("deliberative") voting combines with a lack of strong candidates within slates (scenario UU). In the presence of short ballots (CS model), having consensus strong candidates (XX-same / XX-diff) can create major representational shortfalls for the $\mathcal{B}$ slate, with one- and two-seat outcomes out of six persisting as support pushes past 50 and even 60%.

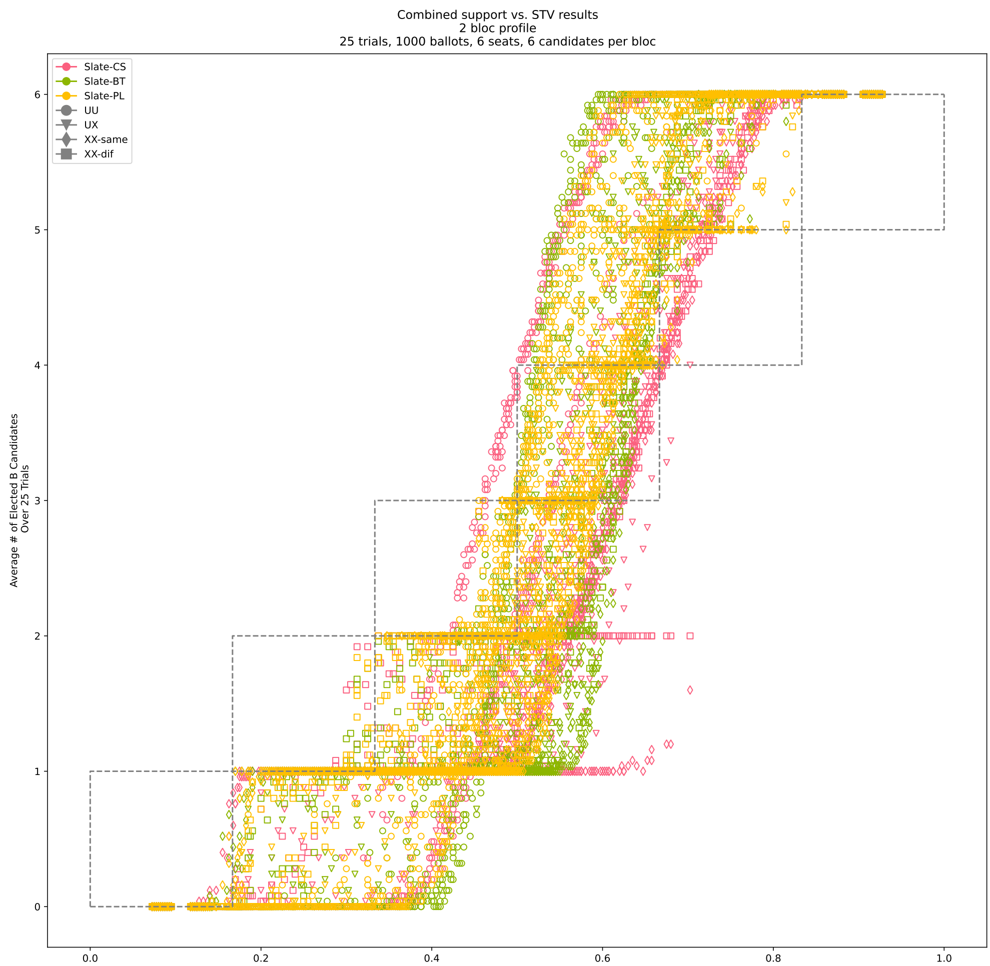

Fig. 4. This time $(r, s, k) = (6, 6, 6)$. We again independently vary the B proportion of the electorate, the generative model, the A and B cohesion, and the candidate strength settings. In this visualization, we have run 25 trials for each parameter tuple, recording the number of $\mathcal{B}$ candidates elected for each simulated profile. The x axis position is the combined support for $\mathcal{B}$ and the y-axis position records the average number of seats over the trials with each tuple of parameters. The dashed lines show the proportionality target rounded up and down to the nearest whole number of seats.

All of these observations invite further thought and investigation, starting with unpacking these omnibus diagrams. In the present paper they serve to illustrate the richness of this approach for understanding interactive effects for STV under realistically complex conditions.

4 ASYMPTOTIC PROPERTIES

In this section, we give proof of concept that the framework presented here is robust enough to admit provable statements about STV, a system of election for which theorems have so far been elusive.[13]

4.1 Single bloc asymptotics

In this section, we focus on the case of one bloc of voters and two slates of candidates. Note that even with a single bloc the fact that we have two slates means any lack of cohesion immediately leads to the richer types of crossover ballots that motivated our generative models.

For the Slate-PL and Name-PL models, we can prove theoretical results that offer a kind of asymptotic generalization of the well-known Proportionality for Solid Coalitions (PSC). We give asymptotics as the number of voters goes to infinity, since our models are probabilistic.

We start by giving bounds on the outcomes for a bloc voting under Slate-PL model. The results reveal that the choice of precise method for tallying votes has a profound impact on the expected outcomes. With that in mind, we define two different methods for deciding which candidates are elected in each round of an STV vote tallying process.

- **Simultaneous election:** if multiple candidates exceed the threshold for election in a certain round, they are all elected and their excess votes transfer down to the remaining candidates before the next round.
- **One-by-one election:** if multiple candidates exceed the threshold for election in a certain round, the one with the most votes is elected and their excess votes are transferred. The tallying process then proceeds to the next round.

Based on the way that election results are reported by the city of Cambridge, is appears that Cambridge follows the simultaneous election method.[14]

Proposition 4.1 (Slate-PL with simultaneous election). *Consider an election for k open seats, a single bloc of N voters, and two slates of candidates $\mathcal{A}$ and $\mathcal{B}$. Suppose that the voters vote according to a Slate-PL model with cohesion parameter $0.5 < \pi_A \leq 1$, and all voters rank the candidates within each slate in the same order. Suppose also that the number of candidates in each slate is more than k, and the votes are tallied using simultaneous election.*

(a) *For all $\varepsilon > 0$, the number of candidates elected from slate $\mathcal{B}$ is bounded below by $\lfloor (1-\pi_A)(k+1) - \varepsilon \rfloor$ and above by $\lfloor k/2 \rfloor$ asymptotically almost surely as $N \to \infty$.*

(b) *Suppose $\pi_A < 1$. As $k \to \infty$, the fraction of elected candidates which are from slate $\mathcal{B}$ (asymptotically almost surely as $N \to \infty$) tends to $1/2$.*

Note that the lower bound in (a) is precisely the number of thresholds exceeded by the first-place votes for slate $\mathcal{B}$.

Proof. We first derive the lower bound in (a). At any stage during the vote tallying process, let ω_A (resp. ω_B) denote the fraction of the original N ballots which have both not been discarded yet, and whose top vote is from $\mathcal{A}$ (resp. $\mathcal{B}$). Since we are concerned with results as $N \to \infty$, we may assume that these fractions are, up to an arbitrarily small error, deterministic quantities at each stage of the vote tallying process.

Note that the top candidate of any ballot is determined by its slate, since the ranking of all candidates within a slate is the same across ballots. This means that only one candidate from $\mathcal{A}$ and one candidate from $\mathcal{B}$ receive first-place votes at a time. Let $t = 1/(k+1)$ denote the Droop quota

[13] The assumption of solid coalitions, in particular, assumes away any role for transfer between blocs.

[14] See for instance https://www.cambridgema.gov/Election2023/Official/Council%20Round.htm

as a fraction of total votes. If $\omega_A > t$ and $\omega_B > t$, then one candidate is elected from each slate. After discarding votes and fractional transfers, the fractions ω_A and ω_B are updated as follows

$$\omega'_A = (\omega_A + \omega_B - 2t)\pi_A$$
$$\omega'_B = (\omega_A + \omega_B - 2t)(1 - \pi_A)$$

Thus, the ratio ω_A/ω_B returns to $\pi_A/(1-\pi_A)$ immediately after two candidates (one from each slate) are elected in a singe round. If $(1-\pi_A) \le t = 1/(k+1)$, then the bound holds trivially, so suppose $(1-\pi_A) > t$. Two candidates are elected in round i if $(1-\pi_A)(1-2t(i-1)) > t$. Setting $i = (1-\pi_A)(k+1) + \varepsilon$, we obtain

$$\begin{aligned}(1-\pi_A)(1-2t(i-1)) &= (1-\pi_A) - 2t(i-1)(1-\pi_A)\\ &\ge (1-\pi_A) - t(i-1) \qquad (\text{since } \pi_A \ge 0.5)\\ &> t\end{aligned}$$

so at least $\lfloor(1-\pi_A)(k+1)\rfloor$ candidates from slate $\mathcal{B}$ are elected. This proves the lower bound.

For the upper bound, note that if we start with $\omega_A > \omega_B$, then this inequality is maintained except for after a round where only an $\mathcal{A}$ candidate is elected. Following such a round, a single $\mathcal{B}$ candidate can be elected, after which the inequality $\omega_A > \omega_B$ is restored. Thus the only rounds in which a single $\mathcal{B}$ candidate can be elected are directly after rounds where a single $\mathcal{A}$ candidate is elected. It follows that at least as many $\mathcal{A}$ candidates as $\mathcal{B}$ candidates are elected.

To prove (b), note that the first election of single $\mathcal{A}$ candidate (rather than the simultaneous election of and $\mathcal{A}$ and $\mathcal{B}$ candidate) takes place when $(1-\pi_A)(\omega_A+\omega_B) < t$. Thus the number of candidates which can still be elected is at most

$$\begin{aligned}(\omega_A+\omega_B)/t &= (1-\pi_A)\left(1 + \frac{\pi_A}{1-\pi_A}\right)(\omega_A+\omega_B)/t\\ &< 1 + \frac{\pi_A}{1-\pi_A}.\end{aligned}$$

As $k \to \infty$, the ratio of this quantity to k goes to zero, which gives the required result. □

See Figure 5 for an empirical demonstration of Proposition 4.1. To obtain the exact asymptotics (as $N \to \infty$) plotted in the figure, we allow a fractional number of ballots of each kind, and assume that the number of ballots of each kind is exactly equal to the expectation under the model. We also assume that vote transfers are fractional and deterministic.

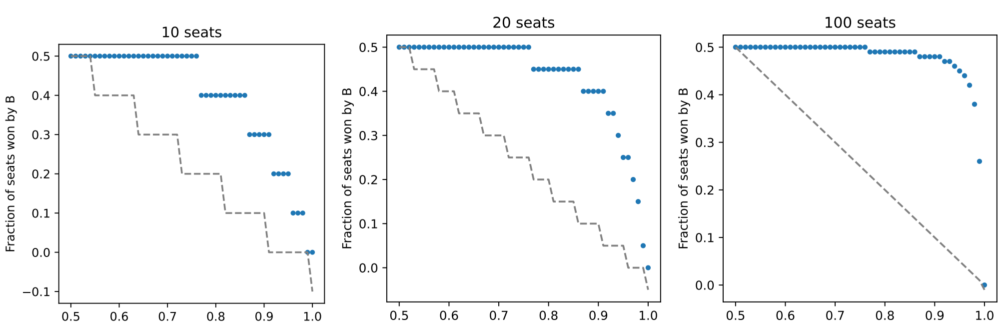

Fig. 5. A visualization of the lower bound and limiting behavior for a single bloc of voters described in Proposition 4.1. The dotted line indicates the lower bound in part (a) of the proposition, and the blue points are exact asymptotics as $N \to \infty$ for various values of π_A.

It is somewhat surprising that, as $k \to \infty$, $\mathcal{A}$ and $\mathcal{B}$ are equally represented even though all voters are in bloc A. Proposition 4.1 assumes simultaneous election transfers—this, together with the fact that there are fixed rankings over $\mathcal{A}, \mathcal{B}$, creates a situation where in nearly every round all first-place votes land on the top remaining $\mathcal{A}$ and $\mathcal{B}$ candidates, and both are elected.

We now consider the one-by-one vote tallying method. A practical difference between the simultaneous and one-by-one elections is that one-by-one election may exhibit a kind of leap-frogging, where a candidate who is over the threshold in round 1 may nonetheless be elected after a candidate who was below the threshold in round 1. This does not happen in simultaneous elections.

Proposition 4.2 (Slate-PL with one-by-one election). *Consider an election for k open seats, a single bloc of N voters, and two slates of candidates $\mathcal{A}$ and $\mathcal{B}$. Suppose that the voters vote according to a Slate-PL model with cohesion parameter $0.5 < \pi_A \leq 1$, and all voters rank the candidates within each slate in the same order. Suppose also that the number of candidates in each slate is more than k, and the votes are tallied using one-by-one election.*

Then, as $k \to \infty$, the fraction of candidates elected from $\mathcal{A}$ is lower bounded by

$$1 - 1/\lceil \log_{\pi_A}(1/2)) \rceil$$

and upper bounded by

$$1 - 1/(1 + \lceil \log_{\pi_A}(1/2)) \rceil)$$

asymptotically almost surely as $N \to \infty$.

Proof. Let ω_A^z (ω_B^z) denote the fraction of ballots which have not been discarded, whose top vote is from $\mathcal{A}$ ($\mathcal{B}$, respectively) at the start of round z.

Since $\pi_A > 0.5$, $\omega_A^1 = \pi_A > \omega_B^1 = 1 - \pi_A$ and an $\mathcal{A}$ candidate is elected in the first round. We have the following update:

$$\omega_A^2 = (\omega_A^1 - t)\pi_A = \pi_A^2 - \pi_A t$$
$$\omega_B^2 = \omega_B^1 + (\omega_A^1 - t)(1 - \pi_A)$$

If an $\mathcal{A}$ candidate is elected again in the second round, then

$$\omega_A^3 = (\omega_A^2 - t)\pi_A = ((\pi_A^2 - \pi_A t) - t)\pi_A = \pi_A^3 - \pi_A^2 t - \pi_A t$$
$$\omega_B^3 = \omega_B^2 + (\omega_A^2 - t)(1 - \pi_A)$$

Starting from the first round of the election, suppose $s - 1$ $\mathcal{A}$ candidates were elected thusfar. Now an $\mathcal{A}$ candidate is elected in seat s if $\pi_A^s \geq 1 - \pi_A^s - (s-1)t$. Letting $k \to \infty$ (so $t \to 0$), an $\mathcal{A}$ candidate is elected for the s-th seat if $\pi_A^s \geq 1 - \pi_A^s$, or $\pi_A^s \geq 1/2$, otherwise a $\mathcal{B}$ candidate is elected. It follows that $\lfloor s^* \rfloor$ candidates from $\mathcal{A}$ are elected consecutively at the start of the election, with $s^* = \log(1/2)/\log(\pi_A)$ satisfying $\pi_A^{s^*} = 1 - \pi_A^{s^*}$, followed by the election of the first $\mathcal{B}$ candidate. At this stage the fraction of $\mathcal{A}$ candidates elected is $1 - 1/(\lfloor s^* \rfloor + 1)$.

Let a *sequence* of rounds consist of the time between the elections of $\mathcal{B}$ candidates. The first sequence starts at the beginning of the election process and ends with the election of the first $\mathcal{B}$ candidate in seat $\lfloor s^* \rfloor + 1$. The second sequence starts from seat $\lfloor s^* \rfloor + 2$ and ends when a $\mathcal{B}$ candidate is next elected, etc.

Notice that at the start of the first sequence $\mathcal{A}$'s fraction of the overall first-place votes were π_A. At the start of any subsequent sequence the most recent update from electing a $\mathcal{B}$ candidate is

$$\omega_A' = \omega_A + (\omega_B - t)\pi_A$$
$$\omega_B' = (\omega_B - t)(1 - \pi_A)$$

from which it follows that $\mathcal{A}$'s share of the first-place votes is at least π_A. As a result, each sequence will elect at least as many $\mathcal{A}$ candidates as the first, from which we conclude that $\mathcal{A}$'s fraction of seats is at least $1 - 1/(\lfloor s^* \rfloor + 1)$.

For the upper bound, observe that if initially $\omega_A^0 = 1$, then after electing the first $\mathcal{A}$ candidate and updating, $\omega_A^1 = \pi_A$ as in our starting condition. Other words, when a sequence starts with $\omega_A = \pi_A$ it elects $\lfloor s^* \rfloor$ $\mathcal{A}$ candidates before the first $\mathcal{B}$ candidate, and if it starts with $\omega_A = 1$ it elects $\lfloor s^* \rfloor + 1$. At the start of every sequence, after electing a $\mathcal{B}$ candidate, $\pi_A \leq \omega_A \leq 1$, from which we conclude that at most $\lfloor s^* \rfloor + 1$ $\mathcal{A}$ candidates are elected for every $\mathcal{B}$ candidate. The bound follows. □

Figure 6 contains a visualization of Proposition 4.2 using the same method to compute exact asymptotics as in Figure 5.

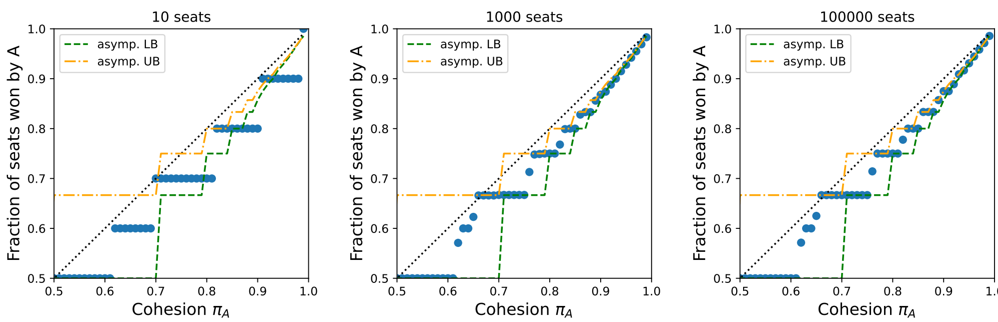

Fig. 6. Visualizations demonstrating the limit behavior described in Proposition 4.2 for $k = 10, 1000, 100000$. The blue points are exact asymptotics as $N \to \infty$ for various values of π_A. The theoretical bounds shown by dashed lines hold in the limit as $k \to \infty$ for $\pi_A > 0.5$. The dotted line is $y = \pi_A$, which is also A's combined share since there is no bloc B.

Finally, when one bloc votes by Name-PL, asymptotic results are easy to describe for extreme candidate strength scenarios, assuming there are more candidates in each slate than seats open, and equal numbers of candidates in each slate.

Proposition 4.3 (Name-PL with one-by-one election). *For ballots generated by a Plackett-Luce model, the STV winners are (a.a.s.) the top candidates by support value (up to a choice about how to break ties between equally supported candidates). Thus we obtain the following results a.a.s. as* $N \to \infty$.

(a) *If* $a_1 \gg a_2 \gg \ldots$ *and* $b_1 \gg b_2 \gg \ldots$, *then equal numbers of candidates are elected from both slates if there are an even number of seats open. If there are an odd number of seats open and* $\pi_A > 0.5$, *then one more* $\mathcal{A}$ *candidate is elected than* $\mathcal{B}$ *candidates.*

(b) *If the support is uniform and* $\pi_A > 0.5$, *then only* $\mathcal{A}$ *candidates are elected.*

Proof. To prove the first statement, consider a Plackett-Luce model with probability vector $(c_1, \ldots, c_k)$. For any partial ranking of candidates $\sigma' = C_1 < C_2 < \ldots < C_\ell$, let $F(\sigma', i)$ be the proportion of ballots which begin with $C_1 < C_2 < \ldots < C_\ell < C_i$. Asymptotically almost surely, if $i, j \notin \sigma'$, we have $c_i < c_j \implies F(\sigma', i) < F(\sigma', j)$. It follows that, initially, the candidates with the most first place votes are (a.a.s.) those with the highest support values. Moreover, after vote transfers, the candidates with the most first places will be (a.a.s.) those unelected and un-eliminated candidates with the highest support values. This proves the statement, and (a) and (b) follow as straightforward observations. □

4.2 Two-bloc asymptotics with fixed candidate order

We conclude our consideration of electoral outcomes with an observation that the asymptotics of two-bloc elections for the one-by-one variant of STV interpolate between solid coalitions and unpolarized voting in an intuitive way. (See Figure 7.)

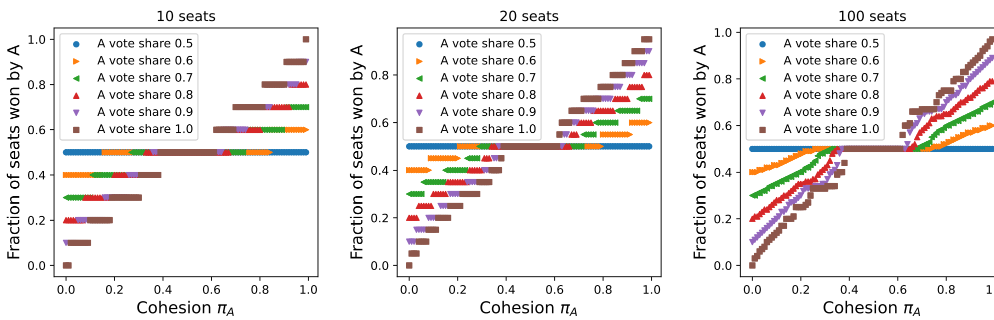

Fig. 7. Exact asymptotics (as the number of voters gets large) showing the share of seats won by the A bloc as their vote share and cohesion varies. The elections have $m = 10, 20, 100$ seats, with an inexhaustible supply of candidates. We use the Slate-PL model, suppose both blocs use the same fixed ordering over $\mathcal{A}$ and $\mathcal{B}$ and apply the one-by-one election variant of STV defined in §4.

One interesting (and real) artifact visible in these plots is that the outcome with seat share of 50% is a plateau that occurs for a range of cohesion values. To get an idea of the reason for this, note that since this plot assumes both blocs use a fixed candidate order $A_1, A_2, \ldots$ and $B_1, B_2, \ldots$, the first candidate elected with $\pi_A, N_A > .5$ will always be A_1. For large numbers of seats, where the election threshold is close to zero, there is a phase transition when $\pi^2 = (1 - \pi) + \pi(1 - \pi)$, occurring at $\pi = 1/\sqrt{2} \approx .707$, that determines whether the first transfer result in the election of A_2. For smaller π, enough support will transfer to B_1 that they are next to be elected. Similar polynomial thresholds determine how many A candidates are elected between successive B candidates. For π approaching $1/2$, the order of election will alternate $ABABAB\ldots$, giving $1/2$ seat share to each side.

5 CONCLUSION AND FUTURE WORK

In §3.2.2 we make first steps toward fitting models and parameters to realistic elections, with immediate payoff in a starkly improved correspondence to Scottish ranked elections than solid coalitions could offer. A more comprehensive fitting effort along these lines—simultaneously learning optimal blocs and slates from observed elections—is a natural future project. This would also point the way to new methods of measuring the degree of polarization, which can feed back usefully into voting rights law.

Our goal in this paper is to lay the groundwork to systematically study the tendency of systems to deliver more or less proportional outcomes for voters. Crucially, the framework we propose allows but does not require party labels, so that we can also consider emergent blocs with similar voting behavior. Finally, the new generative models outlined here can be theoretically explored, opening up rich directions for mathematical study, but can also give decision-makers a powerful toolkit for practical electoral reform.

REFERENCES

Elliot Anshelevich, Onkar Bhardwaj, Edith Elkind, John Postl, and Piotr Skowron. 2018. Approximating optimal social choice under metric preferences. *Artificial Intelligence* 264 (2018), 27–51.

Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. 2017. Justified representation in approval-based committee voting. *Social Choice and Welfare* 48, 2 (2017), 461–485.

William Benter. 2008. Computer based horse race handicapping and wagering systems: a report. In *Efficiency of racetrack betting markets*. World Scientific, 183–198.

Ralph Allan Bradley and Milton E. Terry. 1952. Rank Analysis of Incomplete Block Designs: I. The Method of Paired Comparisons. *Biometrika* 39, 3/4 (1952), 324–345. http://www.jstor.org/stable/2334029

Eric T Bradlow and Peter S Fader. 2001. A Bayesian Lifetime Model for the "Hot 100" Billboard Songs. *J. Amer. Statist. Assoc.* 96, 454 (2001), 368–381. https://doi.org/10.1198/016214501753168091

Markus Brill and Jannik Peters. 2023. Robust and Verifiable Proportionality Axioms for Multiwinner Voting. In *Proceedings of the 24th ACM Conference on Economics and Computation*. 301–301.

Barry C. Burden. 1997. Deterministic and Probabilistic Voting Models. *American Journal of Political Science* 41, 4 (1997), 1150–1169. http://www.jstor.org/stable/2960485

Joshua Clinton, Simon Jackman, and Douglas Rivers. 2004. The statistical analysis of roll call data. *American Political Science Review* (2004), 355–370.

Edith Elkind, Piotr Faliszewski, Jean-François Laslier, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. 2017. What do multiwinner voting rules do? An experiment over the two-dimensional euclidean domain. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 31.

James M Enelow and Melvin J Hinich. 1984. *The spatial theory of voting: An introduction*. CUP Archive.

Nikhil Garg, Wes Gurnee, David Rothschild, and David Shmoys. 2022. Combatting gerrymandering with social choice: The design of multi-member districts. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. 560–561.

Vasilis Gkatzelis, Daniel Halpern, and Nisarg Shah. 2020. Resolving the optimal metric distortion conjecture. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 1427–1438.

Isobel Claire Gormley and Thomas Brendan Murphy. 2007. A Latent Space Model for Rank Data. In *Statistical Network Analysis: Models, Issues, and New Directions*, Edoardo Airoldi, David M. Blei, Stephen E. Fienberg, Anna Goldenberg, Eric P. Xing, and Alice X. Zheng (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 90–102.

Isobel Claire Gormley and Thomas Brendan Murphy. 2008. Exploring voting blocs within the Irish electorate: A mixture modeling approach. *J. Amer. Statist. Assoc.* 103, 483 (2008), 1014–1027.

Todd Graves, C Shane Reese, and Mark Fitzgerald. 2003. Hierarchical models for permutations: Analysis of auto racing results. *J. Amer. Statist. Assoc.* 98, 462 (2003), 282–291.

J Gerald Hebert, Paul Smith, Martina Vandenburg, and Michael DeSanctis. 2010. The Realist's Guide to Redistricting: Avoiding the Legal Pitfalls. American Bar Association.

Valen E Johnson, Robert O Deaner, and Carel P Van Schaik. 2002. Bayesian analysis of rank data with application to primate intelligence experiments. *J. Amer. Statist. Assoc.* 97, 457 (2002), 8–17.

D Marc Kilgour, Jean-Charles Grégoire, and Angèle M Foley. 2020. The prevalence and consequences of ballot truncation in ranked-choice elections. *Public Choice* 184 (2020), 197–218.

Martin Lackner and Piotr Skowron. 2022. Approval-based committee voting. In *Multi-Winner Voting with Approval Preferences*. Springer.

Robin L Plackett. 1975. The analysis of permutations. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 24, 2 (1975), 193–202.

Keith T Poole and Howard Rosenthal. 1985. A spatial model for legislative roll call analysis. *American Journal of Political Science* (1985), 357–384.

Geoffrey Pritchard and Mark C Wilson. 2009. Asymptotics of the minimum manipulating coalition size for positional voting rules under impartial culture behaviour. *Mathematical Social Sciences* 58, 1 (2009), 35–57.

Ariel D Procaccia and Jeffrey S Rosenschein. 2006. The distortion of cardinal preferences in voting. In *International Workshop on Cooperative Information Agents*. Springer, 317–331.

Luis Sánchez-Fernández, Edith Elkind, Martin Lackner, Norberto Fernández, Jesús Fisteus, Pablo Basanta Val, and Piotr Skowron. 2017. Proportional justified representation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 31.

P Skowron, M Lackner, M Brill, D Peters, and E Elkind. 2017. Proportional rankings. In *International Joint Conference on Artificial Intelligence (IJCAI 2017)*. Association for the Advancement of Artificial Intelligence.

Hal Stern. 1990. Models for distributions on permutations. *J. Amer. Statist. Assoc.* 85, 410 (1990), 558–564.

Stanisław Szufa, Piotr Faliszewski, Łukasz Janeczko, Martin Lackner, Arkadii Slinko, Krzysztof Sornat, and Nimrod Talmon. 2022. How to Sample Approval Elections? *arXiv preprint arXiv:2207.01140* (2022).

Stanisław Szufa, Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. 2020. Drawing a map of elections in the space of statistical cultures. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems*. 1341–1349.

Alan D. Taylor. 2002. The Manipulability of Voting Systems. *The American Mathematical Monthly* 109, 4 (2002), 321–337. http://www.jstor.org/stable/2695497

Louis L Thurstone. 1927. A law of comparative judgment. *Psychological review* 34, 4 (1927), 273.

T Nicolaus Tideman and Florenz Plassmann. 2010. The structure of the election-generating universe. (2010). Manuscript.

GJG Upton and D Brook. 1975. The determination of the optimum position on a ballot paper. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 24, 3 (1975), 279–287.

A BALLOT COMPLETION

The distance between two (complete) ballot types is meant to measure the smallest number of swaps between adjacent symbols to turn one ballot into the other. Recall that in a ballot type, candidates of each slate are indistinguishable from one another. For example, for an election with $(r, s) = (2, 3)$ (i.e., $\mathcal{A} = \{A_1, A_2\}$ and $\mathcal{B} = \{B_1, B_2, B_3\}$), $\text{dist}(AABBB, ABBAB) = 2$ and $\text{dist}(AABBB, BBBAA) = 6$. For incomplete ballots and weak orders over candidates, we define this pairwise distance to be the *expected* smallest number of swaps between adjacent symbols required to turn one ballot into the other assuming that each way of breaking ties is equally likely. For example, in the above election with $r = 2$ and $s = 3$, the partial ballot AB has three completions: ABABB, ABBAB and ABBBA. As a result,

$$\text{dist}\left(\begin{smallmatrix} A \\ A \\ B \\ B \\ B \end{smallmatrix}, \begin{smallmatrix} A \\ B \\ - \\ - \\ - \end{smallmatrix}\right) = \frac{1}{3}\left[\text{dist}\left(\begin{smallmatrix} A \\ A \\ B \\ B \\ B \end{smallmatrix}, \begin{smallmatrix} A \\ B \\ A \\ B \\ B \end{smallmatrix}\right) + \text{dist}\left(\begin{smallmatrix} A \\ A \\ B \\ B \\ B \end{smallmatrix}, \begin{smallmatrix} A \\ B \\ B \\ A \\ B \end{smallmatrix}\right) + \text{dist}\left(\begin{smallmatrix} A \\ A \\ B \\ B \\ B \end{smallmatrix}, \begin{smallmatrix} A \\ B \\ B \\ B \\ A \end{smallmatrix}\right)\right] = \frac{1+2+3}{3} = 2.$$

Let $\text{sc}^{A|B}(\sigma)$ be the slate-ordered vector of Borda scores that result from ballot σ, where the first r entries $\text{sc}^{A|B}(\sigma)_1 \geq \text{sc}^{A|B}(\sigma)_2 \geq \cdots \geq \text{sc}^{A|B}(\sigma)_r$ are the scores of A symbols in decreasing order, followed by the scores of B symbols sorted similarly in positions $r + 1$ to $r + s$. In the case of partial or weak orders, we take the average of the sorted score vectors over all ways of breaking ties. For example, ballot $AABBB$ has score vector $(5, 4 \mid 3, 2, 1)$ and ballot $ABABB$ has score vector $(5, 3 \mid 4, 2, 1)$. (The bar is just a decoration to remind us of the break between the slates.) Partial ballot AB has an A symbol and two B symbols tied in positions 3, 4 and 5, resulting in score vector $(5, 2 \mid 4, 2, 2)$, while the weak preference order that rank the candidates $\{A_2, B_1\}$ above $\{A_1, B_2, B_3\}$ results in the slate-ordered score vector $(4.5, 2 \mid 4.5, 2, 2)$. The number of pairwise swaps between adjacent symbols that is required to turn ballot σ_1 into ballot σ_2 is related in a simple way to the L^1 distance between their score vectors.

Lemma A.1. *For ballot types σ_1, σ_2,*

$$dist(\sigma_1, \sigma_2) = \frac{1}{2}\left\| sc^{A|B}(\sigma_1) - sc^{A|B}(\sigma_2)\right\|_1 .$$

B ONE-BLOC PROFILES

B.1 Attributes of profile, split by candidate pool and strength scenario

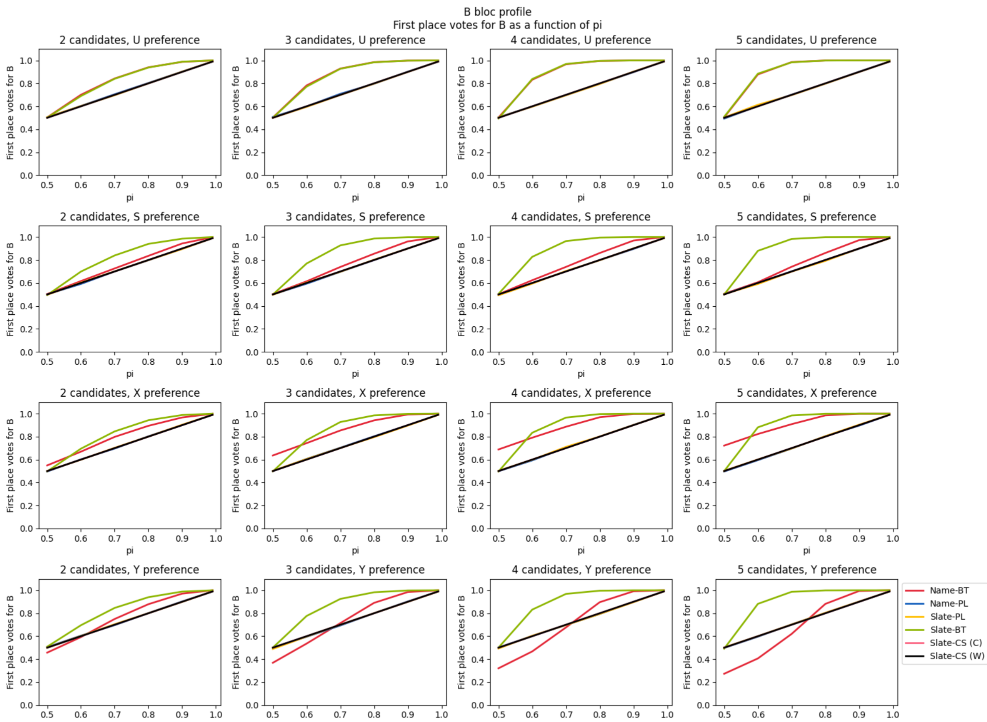

Fig. 8. The proportion of first-place votes for $\mathcal{B}$ candidates. Shown across different generative models, numbers of candidates, and strength scenarios. Notice that Slate-BT and Name-BT are the only two models for which first-place support may differ in expectation from the model parameter π.

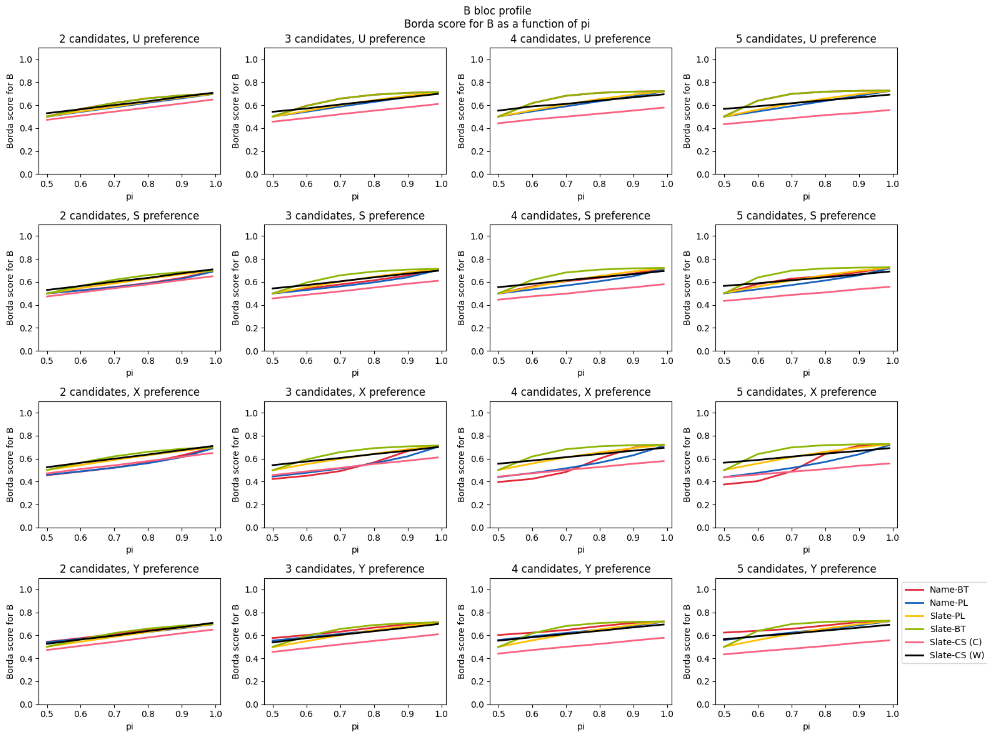

Fig. 9. The proportion of Borda points for $\mathcal{B}$ candidates. Shown across different generative models, numbers of candidates, and strength scenarios.

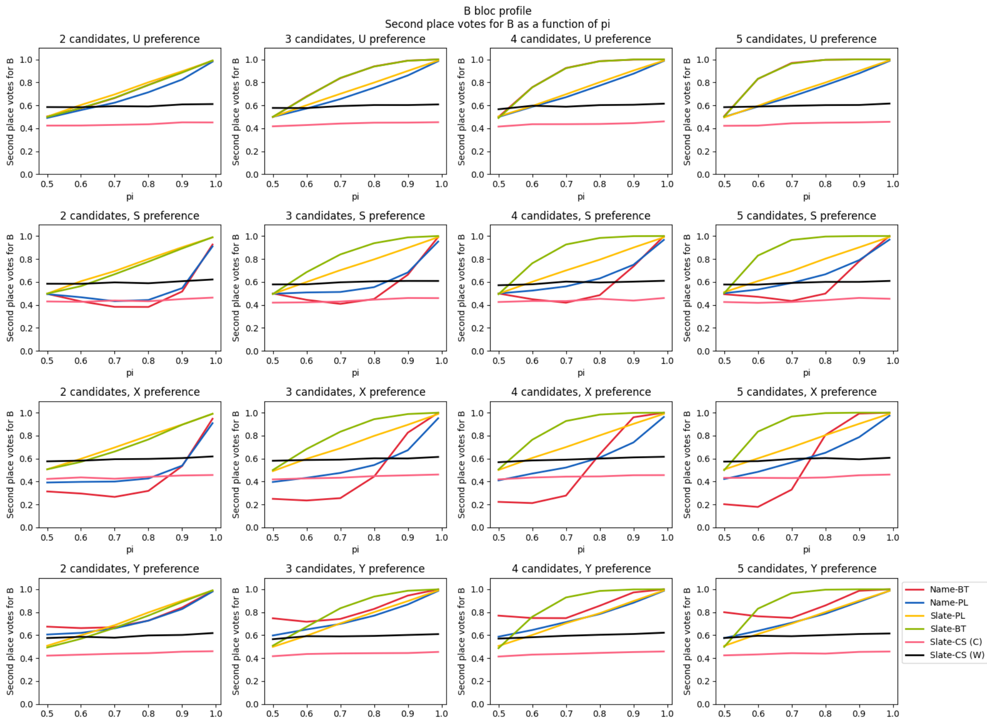

Fig. 10. The proportion of second-place votes for $\mathcal{B}$ candidates. Shown across different generative models, numbers of candidates, and strength scenarios. Notice that in the by-name models, the probability of ranking your own bloc's candidate second can actually be less than 50%, even in cases of high cohesion, if your slate has a strong candidate. (We regard this as evidence that the Slate models are more realistic, but others may hold different views.)

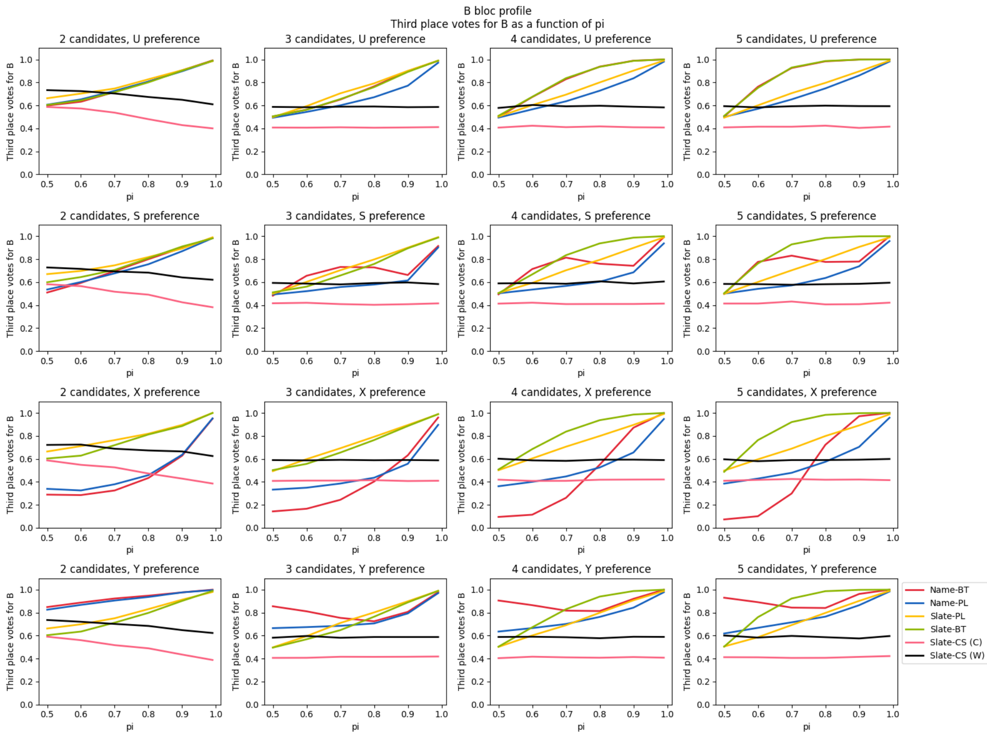

Fig. 11. The proportion of third-place votes for $\mathcal{B}$ candidates. Shown across different generative models, numbers of candidates, and strength scenarios.

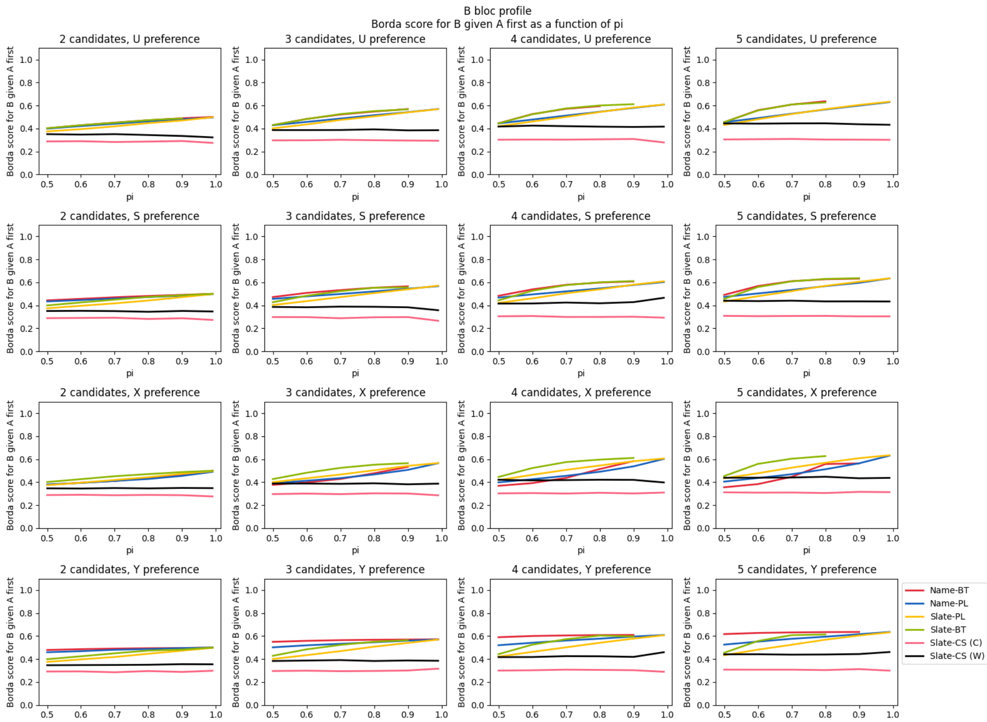

Fig. 12. The proportion of Borda points for $\mathcal{B}$ candidates given that a ballot started with an $\mathcal{A}$ candidate. Shown across different generative models, numbers of candidates, and strength scenarios.

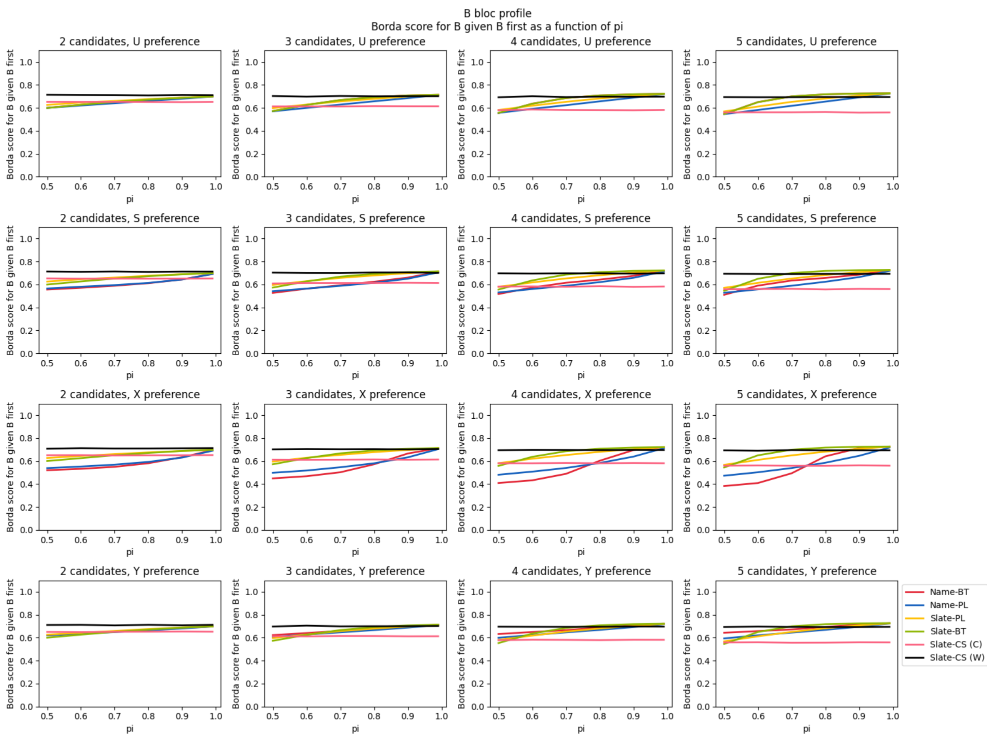

Fig. 13. The proportion of Borda points for $\mathcal{B}$ candidates given that a ballot started with a $\mathcal{B}$ candidate. Shown across different generative models, numbers of candidates, and strength scenarios.

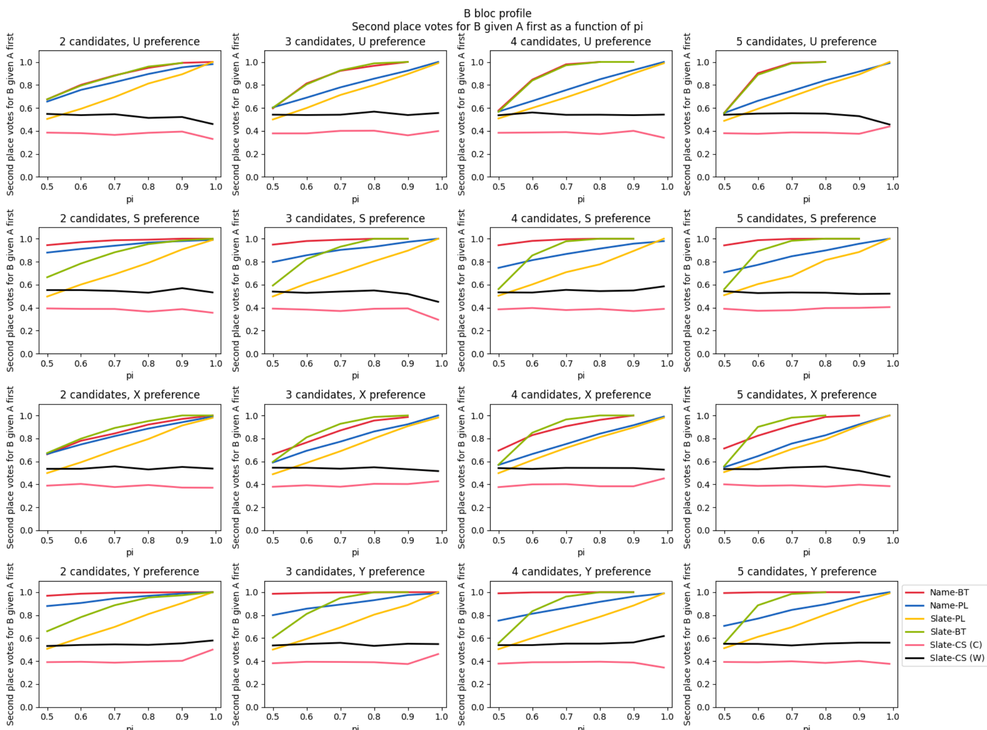

Fig. 14. The proportion of second-place votes for $\mathcal{B}$ candidates given that a ballot started with an $\mathcal{A}$ candidate. Shown across different generative models, numbers of candidates, and strength scenarios.

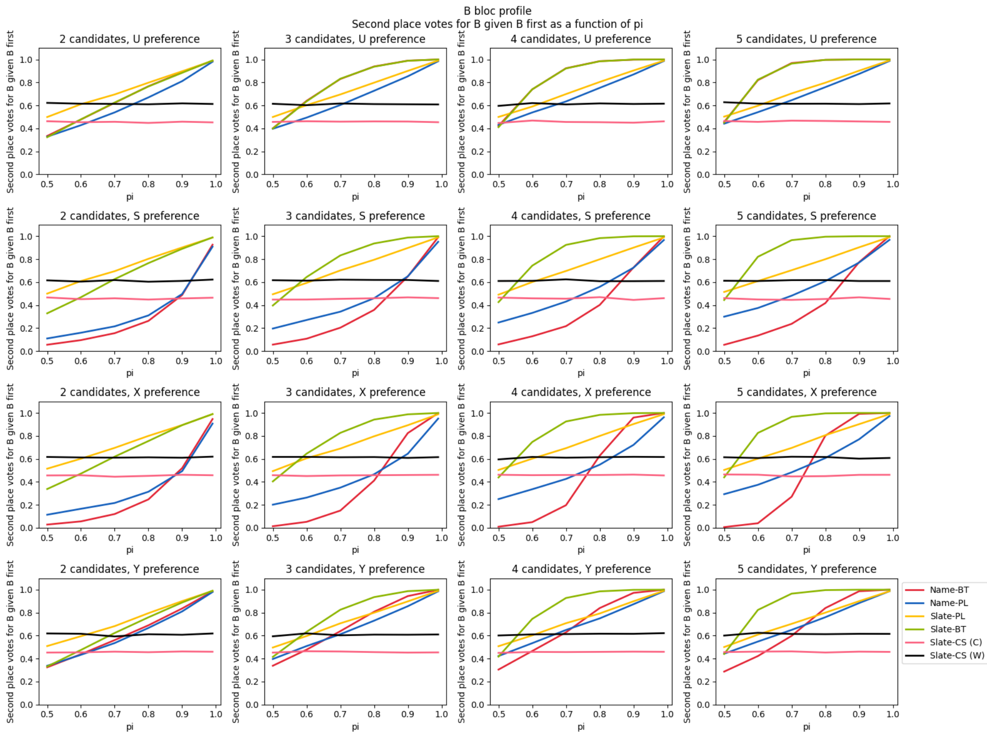

Fig. 15. The proportion of second-place votes for $\mathcal{B}$ candidates given that a ballot started with a $\mathcal{B}$ candidate. Shown across different generative models, numbers of candidates, and strength scenarios.

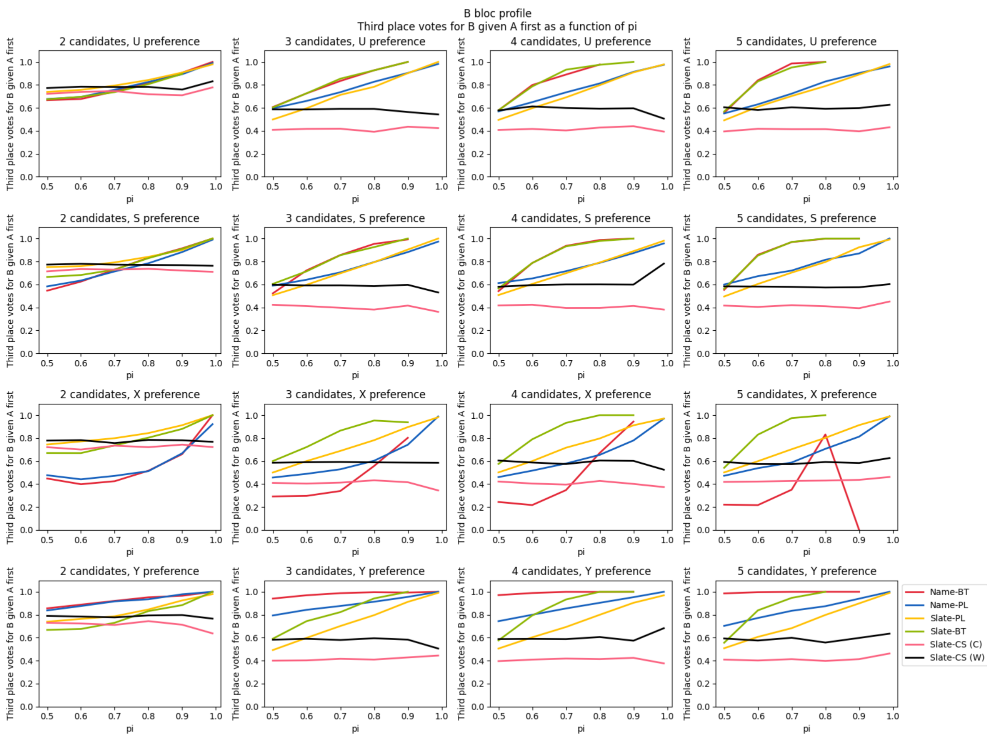

Fig. 16. The proportion of third-place votes for $\mathcal{B}$ candidates given that a ballot started with an $\mathcal{A}$ candidate. Shown across different generative models, numbers of candidates, and strength scenarios.

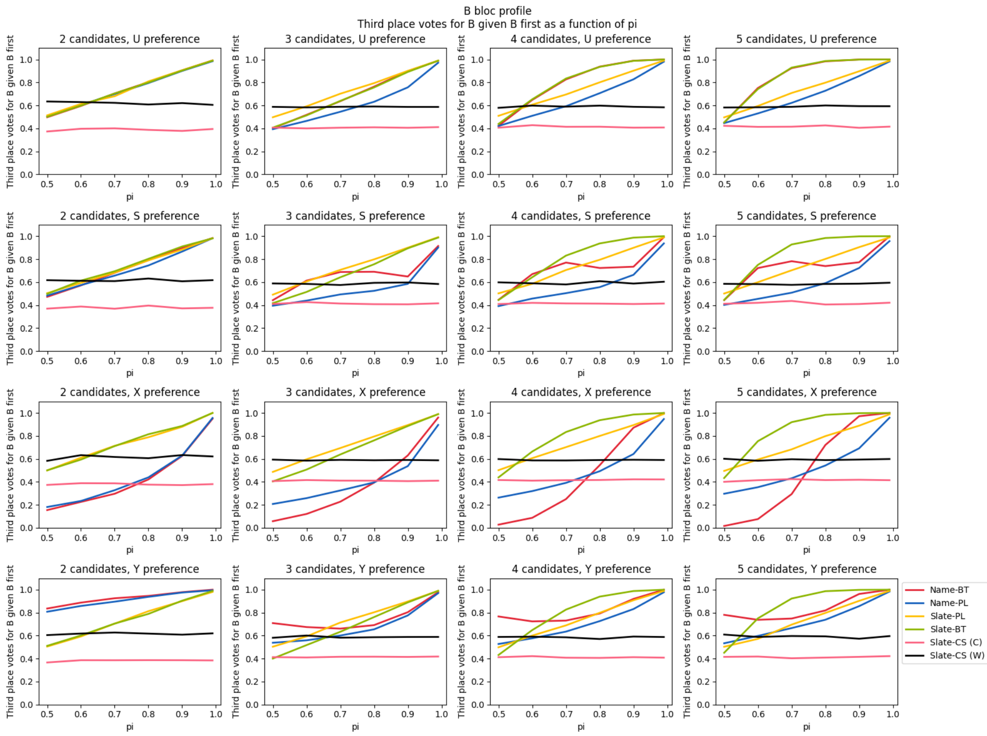

Fig. 17. The proportion of third-place votes for $\mathcal{B}$ candidates given that a ballot started with a $\mathcal{B}$ candidate. Shown across different generative models, numbers of candidates, and strength scenarios.

B.2 Attributes of profile, split by candidate pool and generative model

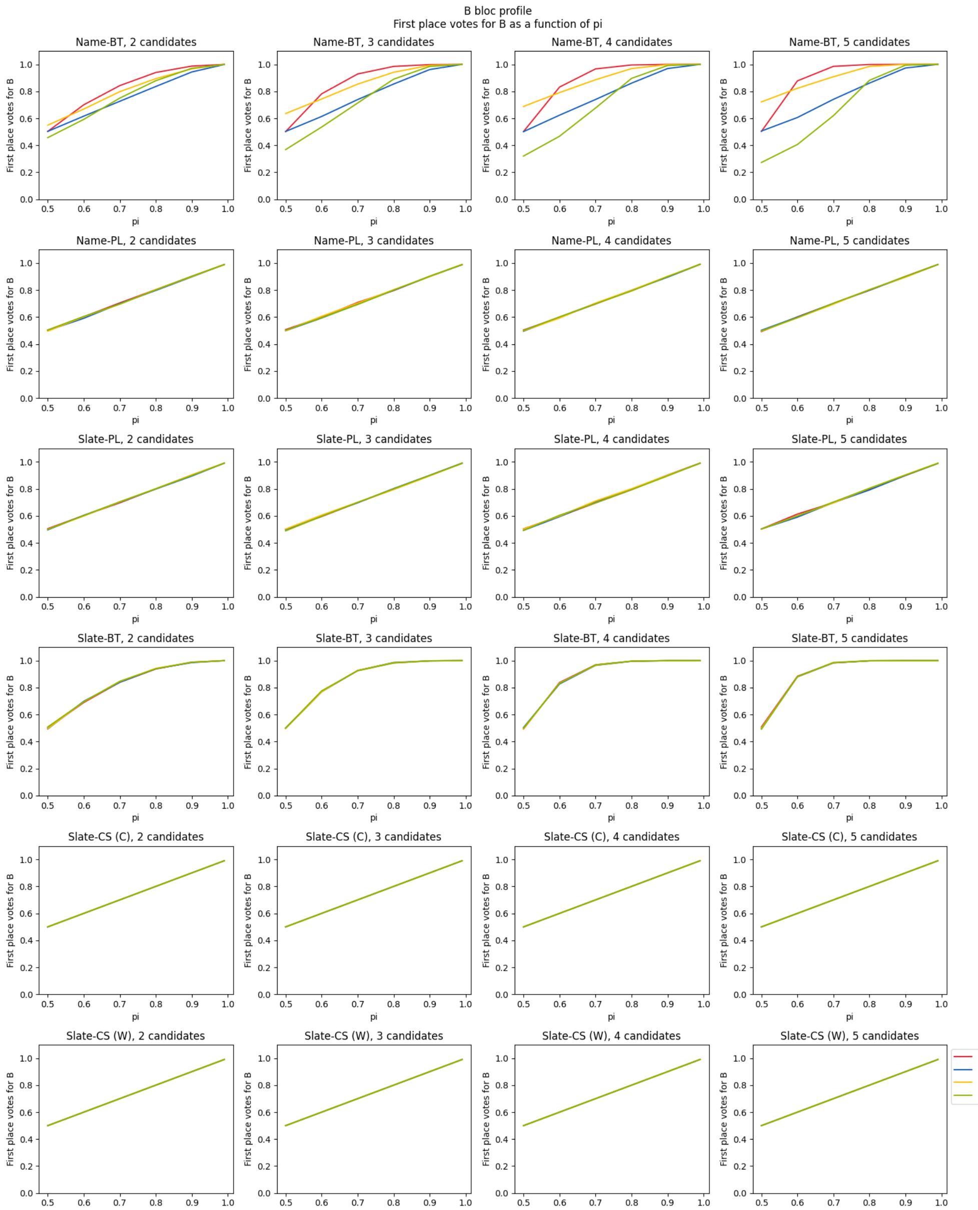

Fig. 18. The proportion of first-place votes for $\mathcal{B}$ candidates. Shown across different generative models, numbers of candidates, and strength scenarios.

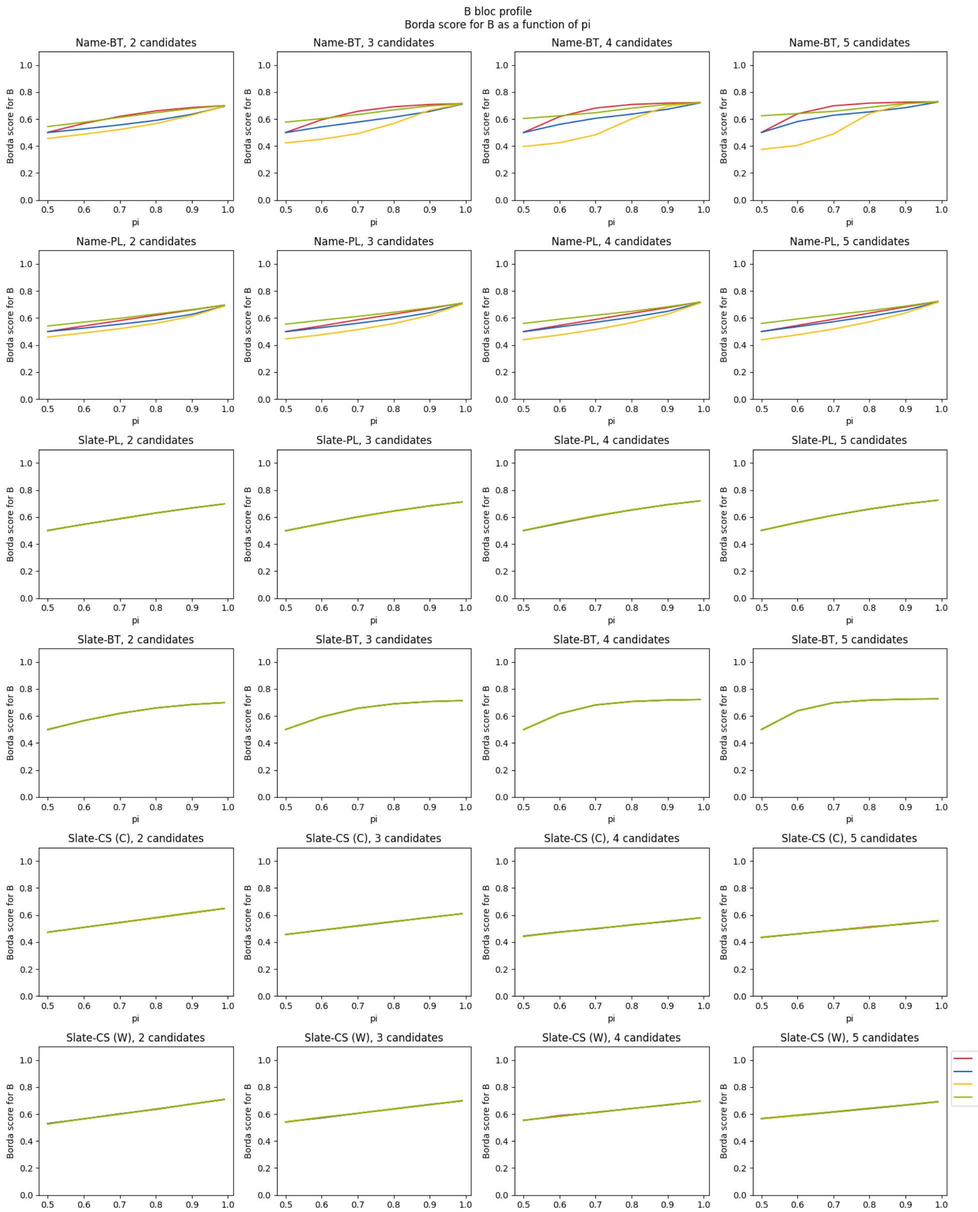

Fig. 19. The proportion of Borda points for $\mathcal{B}$ candidates. Shown across different generative models, numbers of candidates, and strength scenarios.

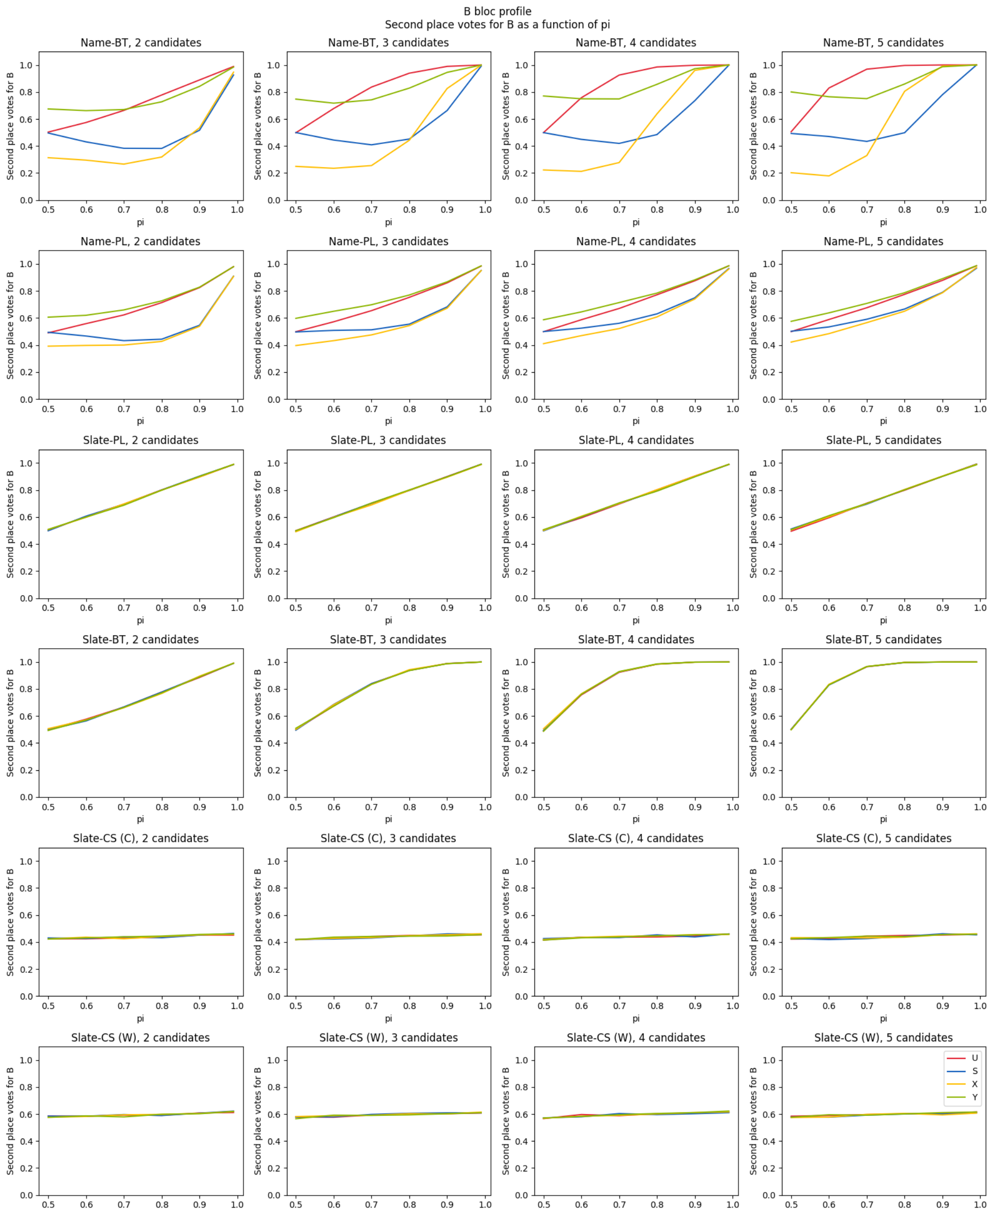

Fig. 20. The proportion of second-place votes for $\mathcal{B}$ candidates. Shown across different generative models, numbers of candidates, and strength scenarios. Notice that in the by-name models, the probability of ranking your own bloc's candidate second can actually be less than 50%, even in cases of high cohesion, if your slate has a strong candidate. (We regard this as evidence that the Slate models are more realistic, but others may hold different views.)

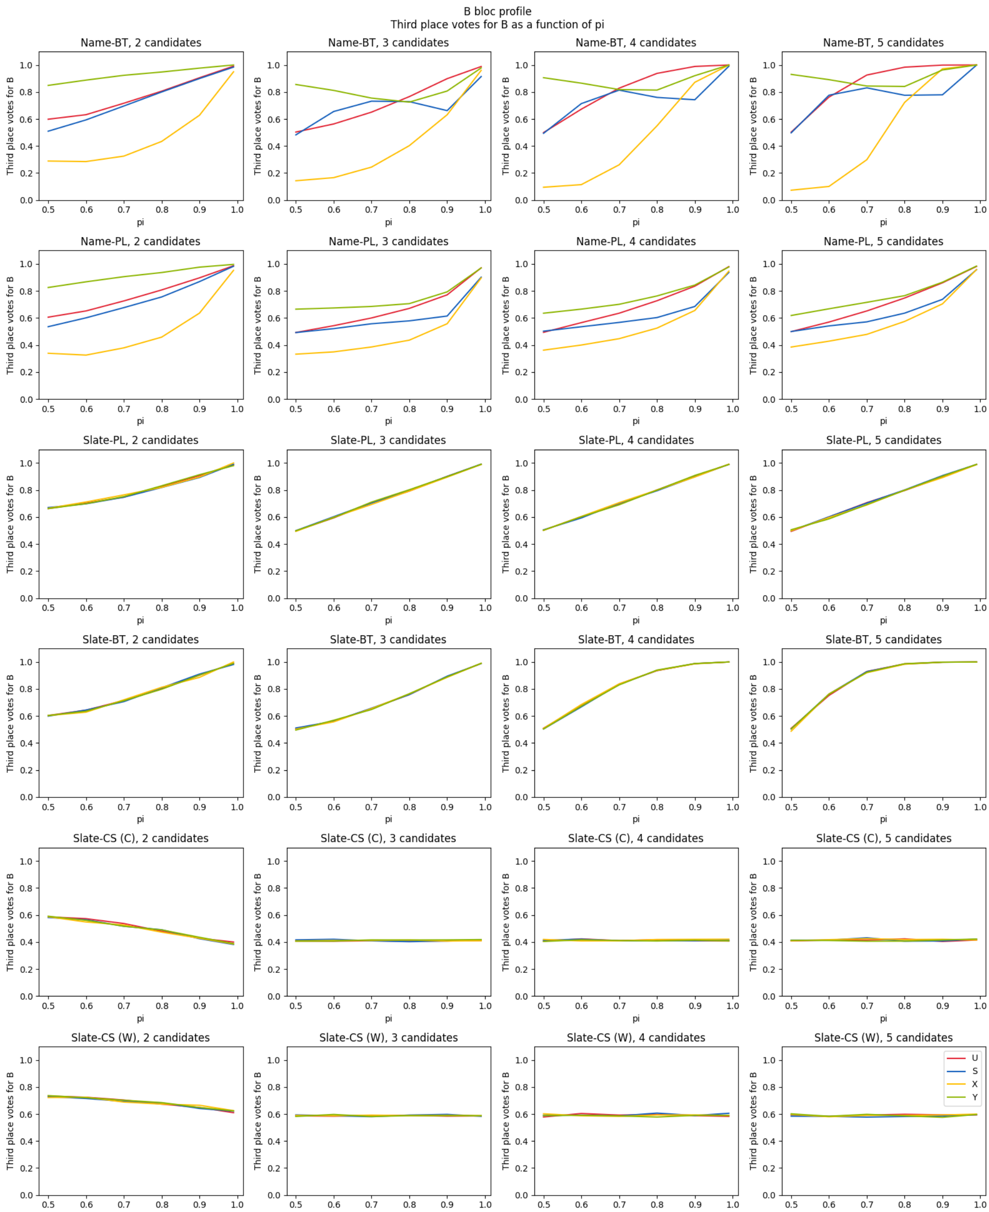

Fig. 21. The proportion of third-place votes for $\mathcal{B}$ candidates. Shown across different generative models, numbers of candidates, and strength scenarios.

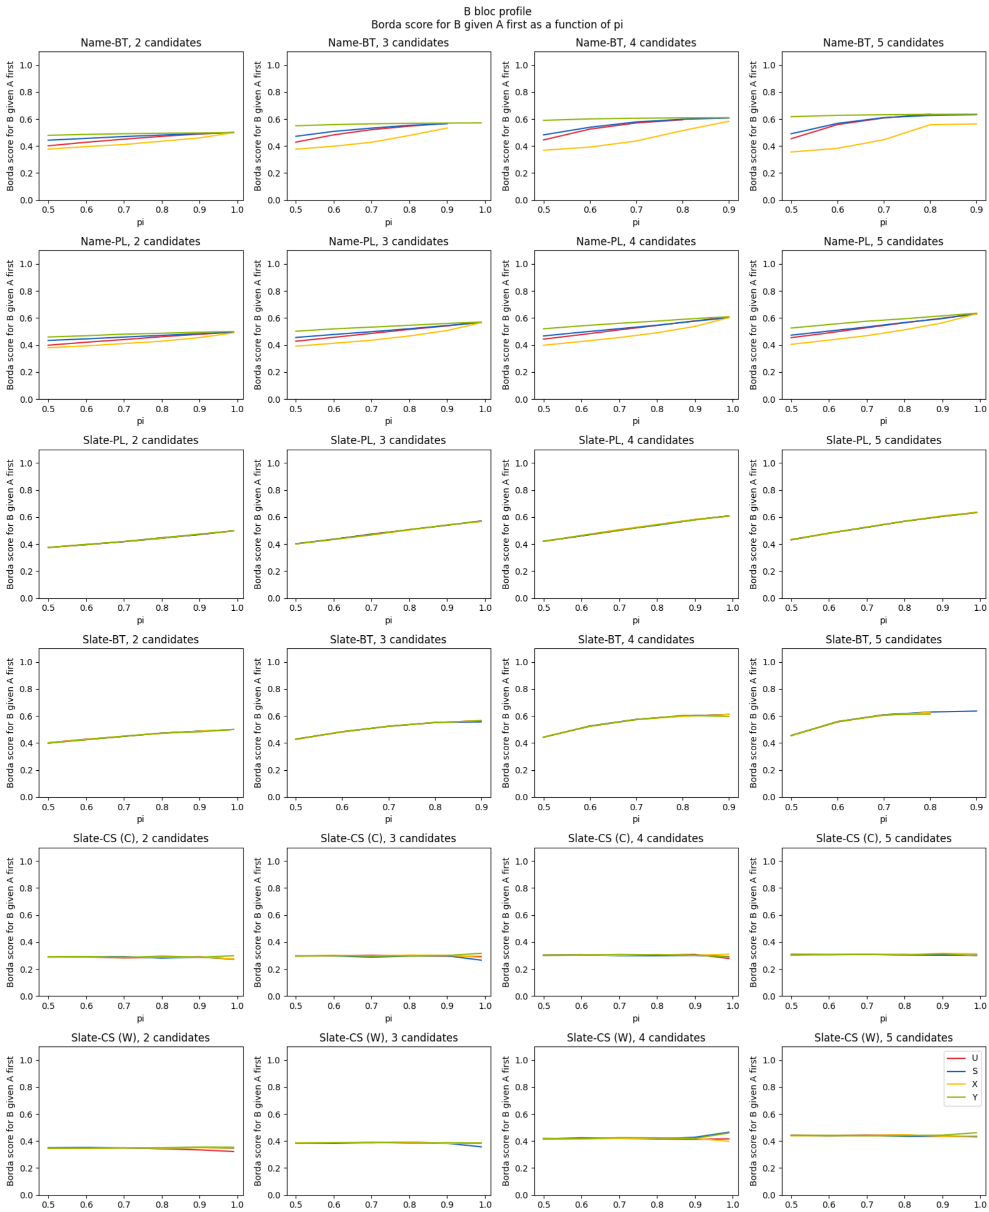

Fig. 22. The proportion of Borda points for $\mathcal{B}$ candidates. given that a ballot started with an $\mathcal{A}$ candidate. Shown across different generative models, numbers of candidates, and strength scenarios.

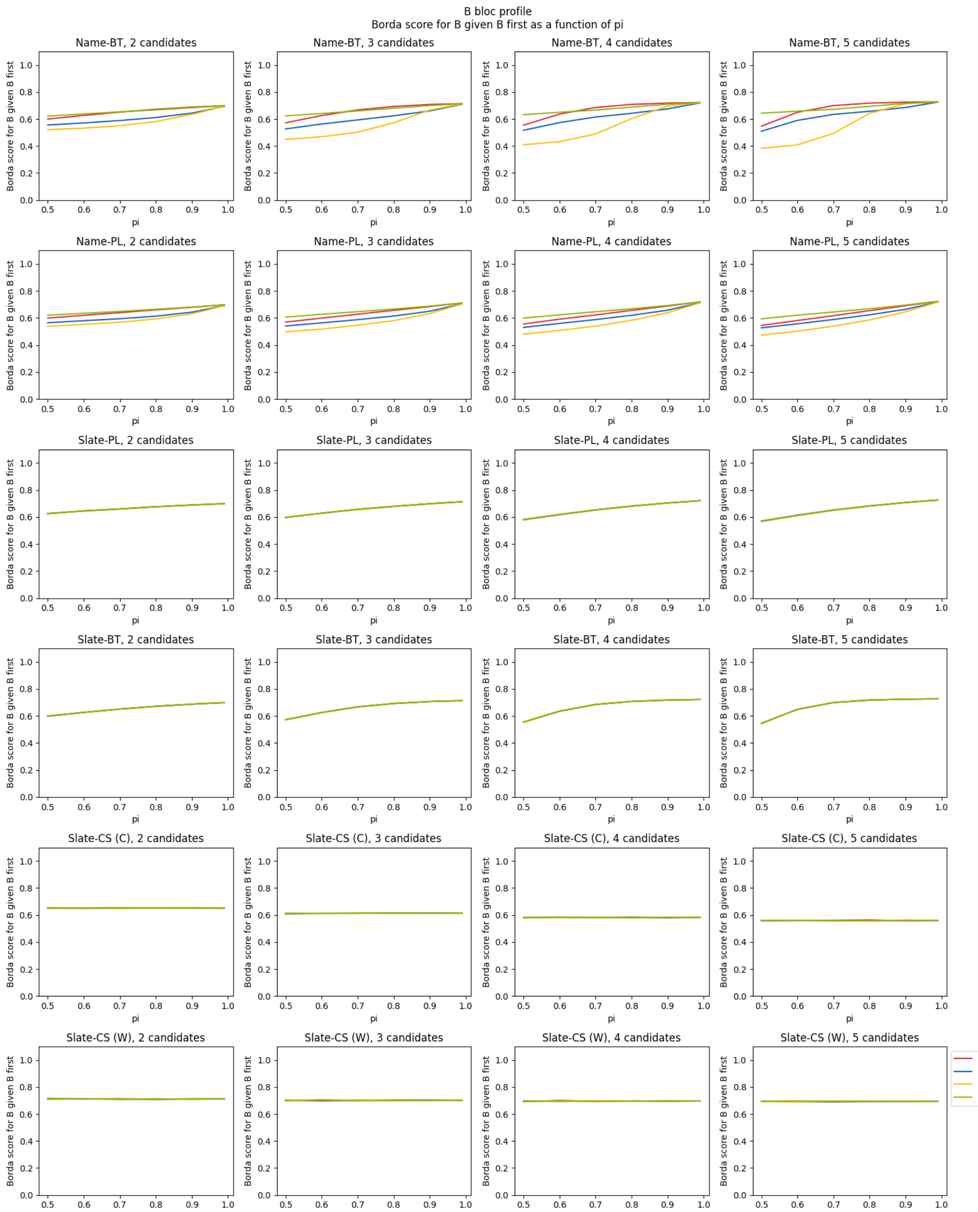

Fig. 23. The proportion of Borda points for $\mathcal{B}$ candidates. given that a ballot started with a $\mathcal{B}$ candidate. Shown across different generative models, numbers of candidates, and strength scenarios.

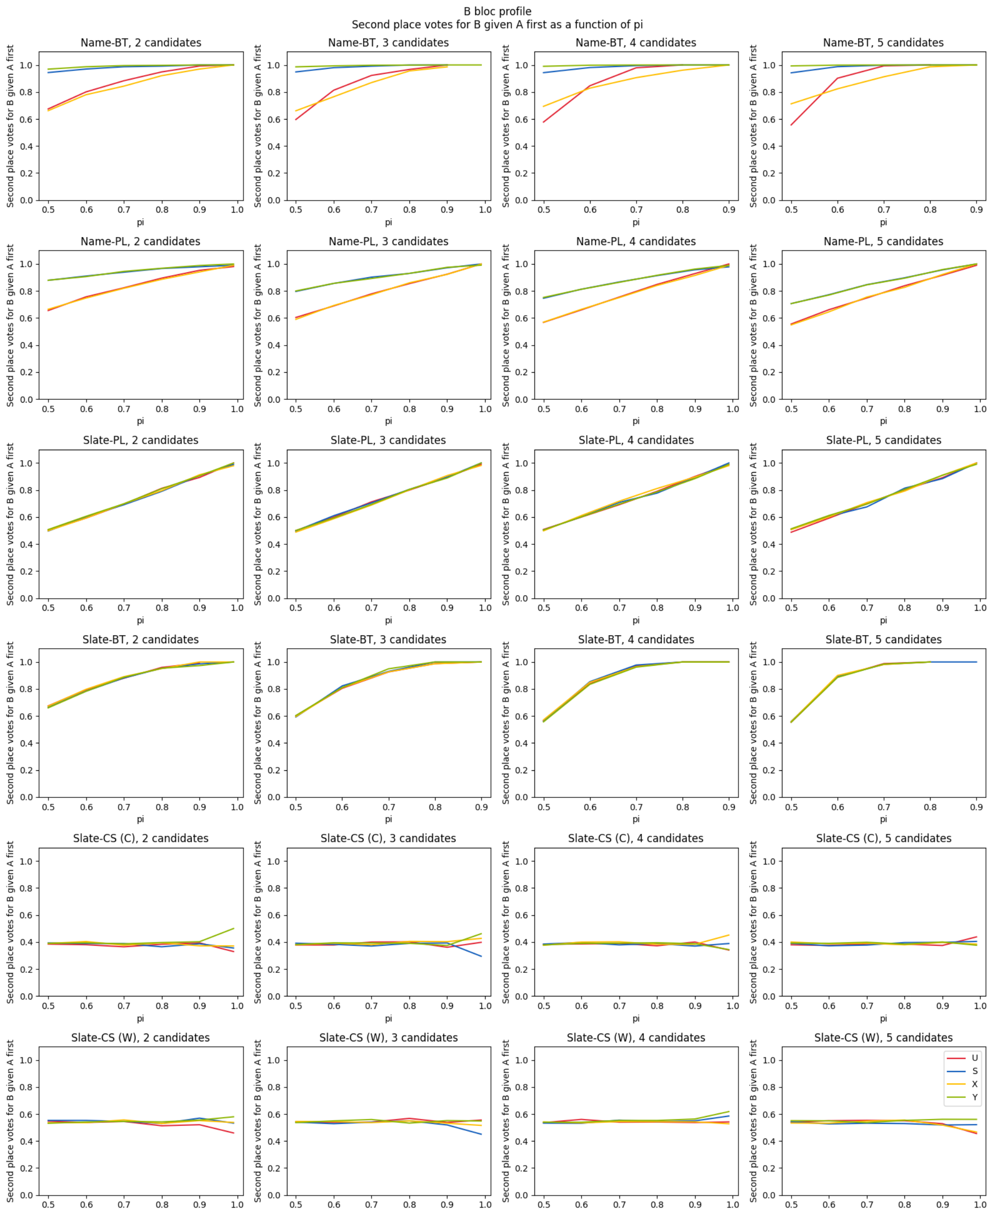

Fig. 24. The proportion of second-place votes for $\mathcal{B}$ candidates. given that a ballot started with a $\mathcal{A}$ candidate. Shown across different generative models, numbers of candidates, and strength scenarios.

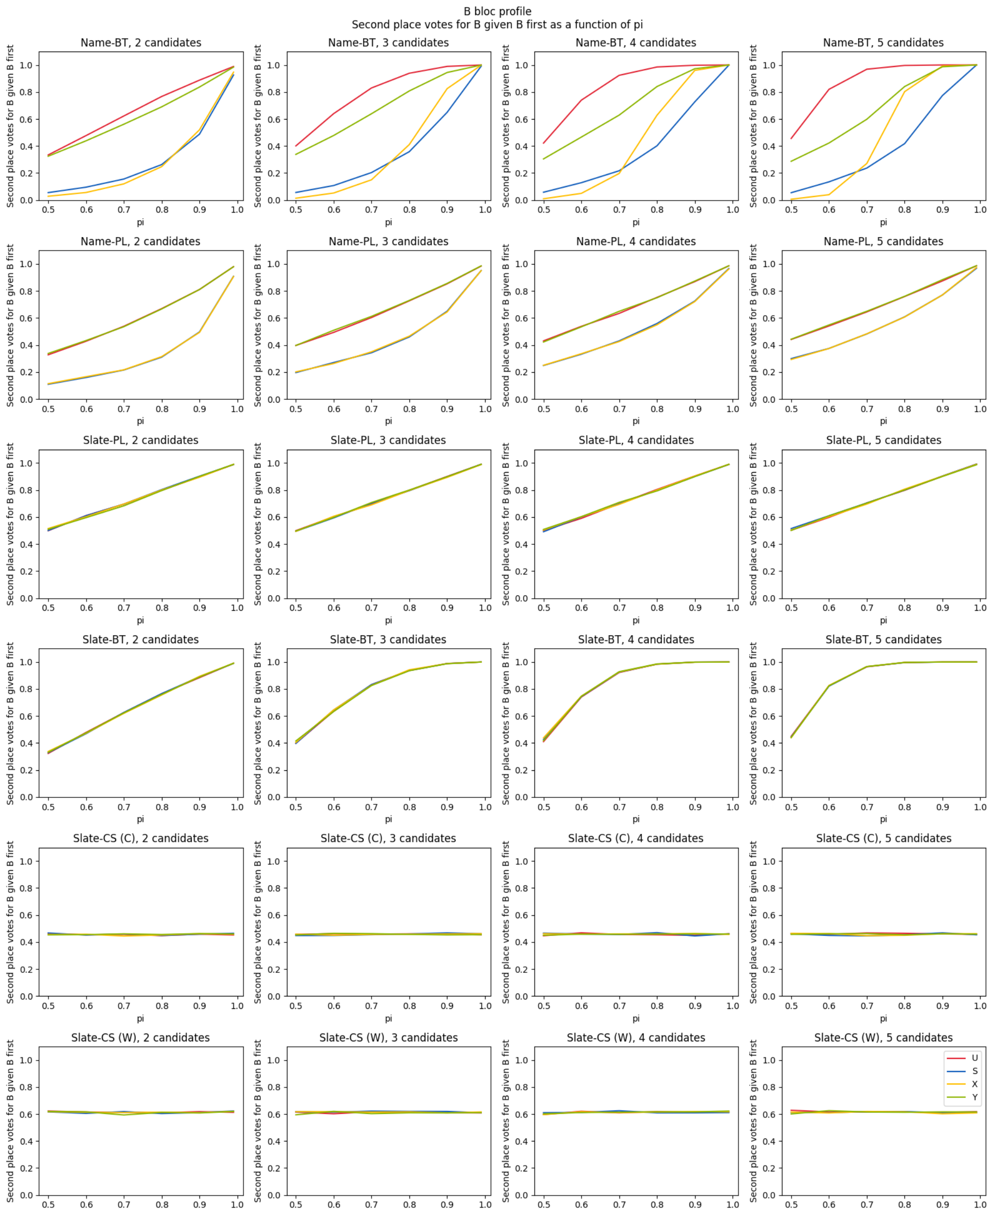

Fig. 25. The proportion of second-place votes for $\mathcal{B}$ candidates given that a ballot started with a $\mathcal{B}$ candidate. Shown across different generative models, numbers of candidates, and strength scenarios.

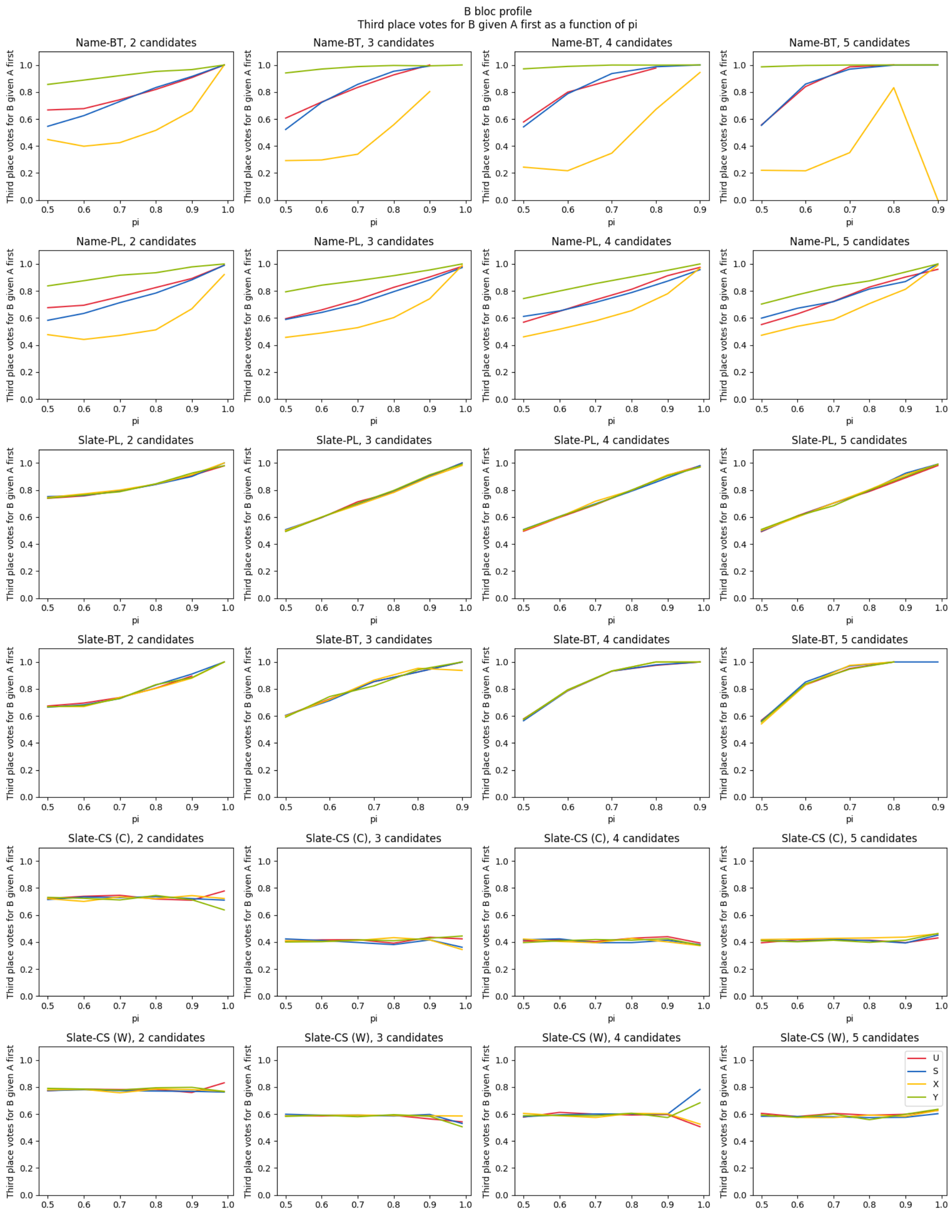

Fig. 26. The proportion of third-place votes for $\mathcal{B}$ candidates given that a ballot started with an $\mathcal{A}$ candidate. Shown across different generative models, numbers of candidates, and strength scenarios.

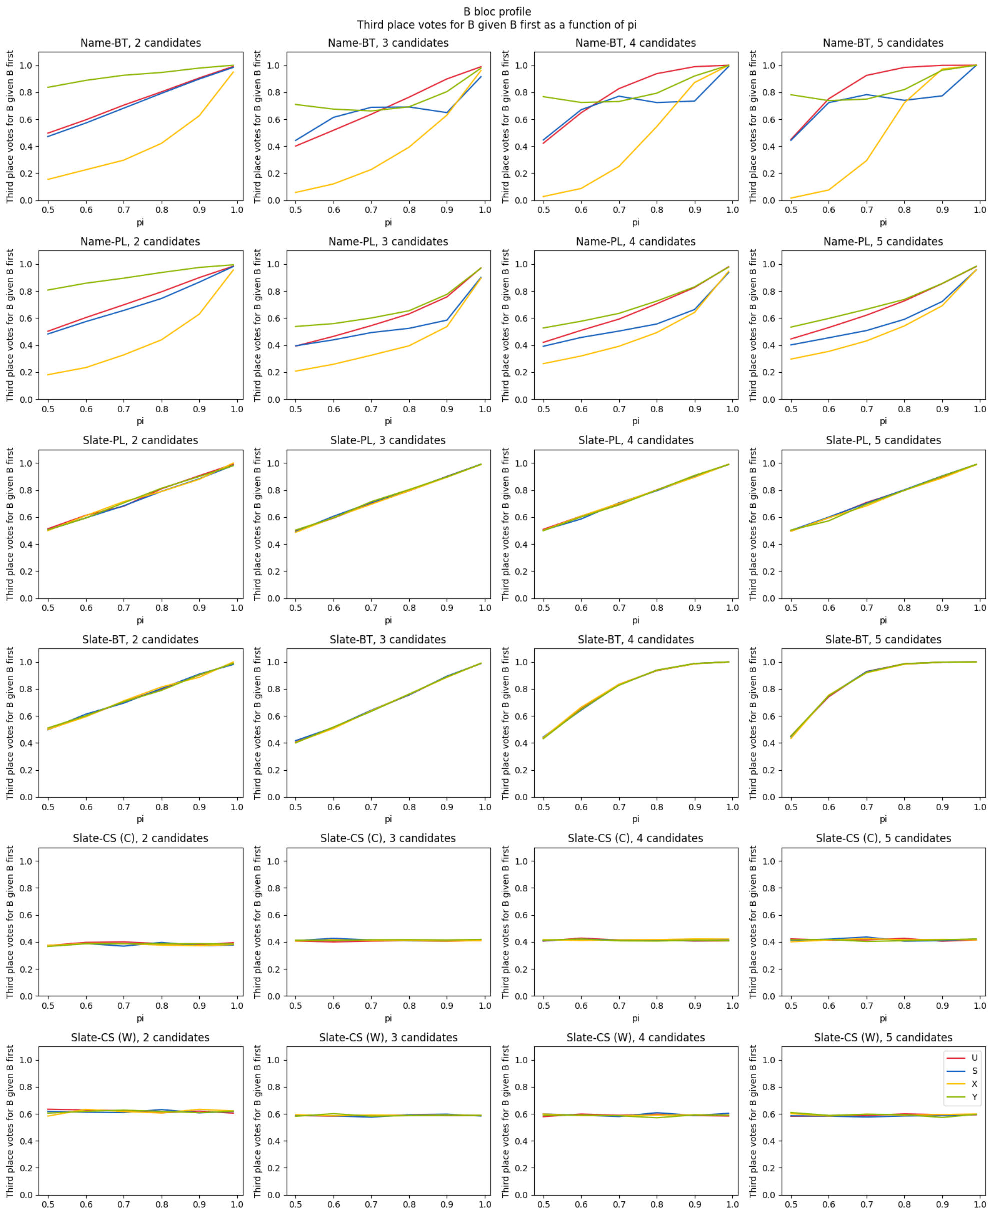

Fig. 27. The proportion of third-place votes for $\mathcal{B}$ candidates given that a ballot started with a $\mathcal{B}$ candidate. Shown across different generative models, numbers of candidates, and strength scenarios.

B.3 Attributes of profile, split by strength scenario and generative model

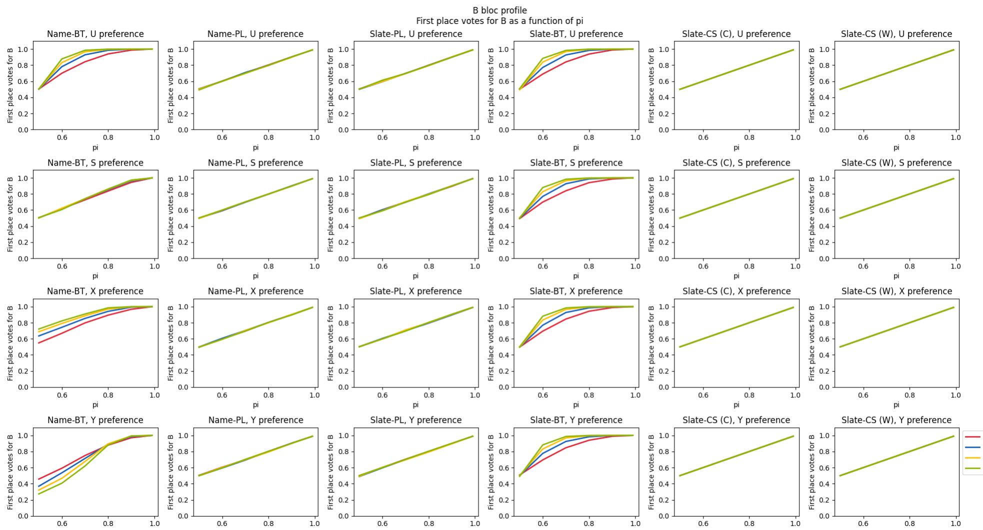

Fig. 28. The proportion of first-place votes for $\mathcal{B}$ candidates across different generative models, numbers of candidates, and strength scenarios.

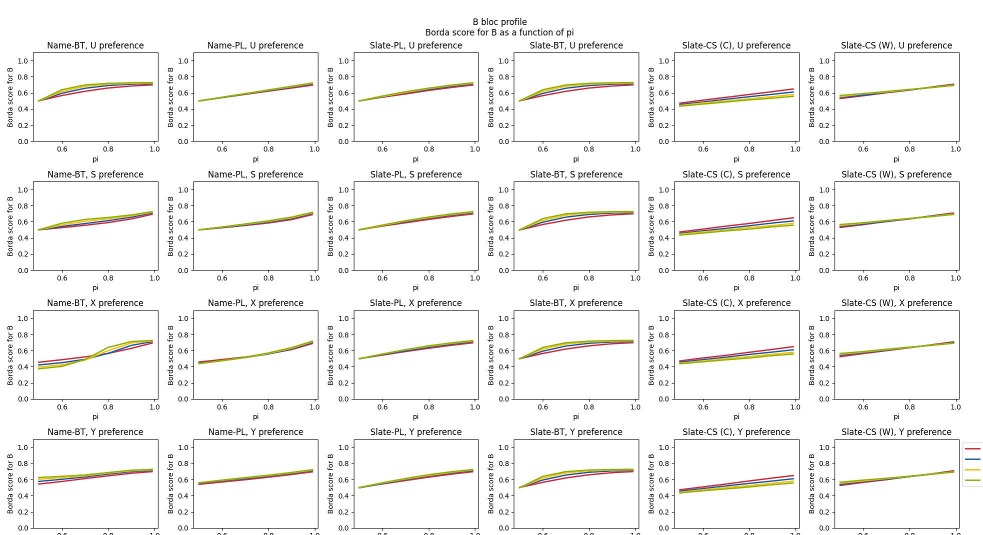

Fig. 29. The proportion of Borda points for $\mathcal{B}$ candidates across different generative models, numbers of candidates, and strength scenarios.

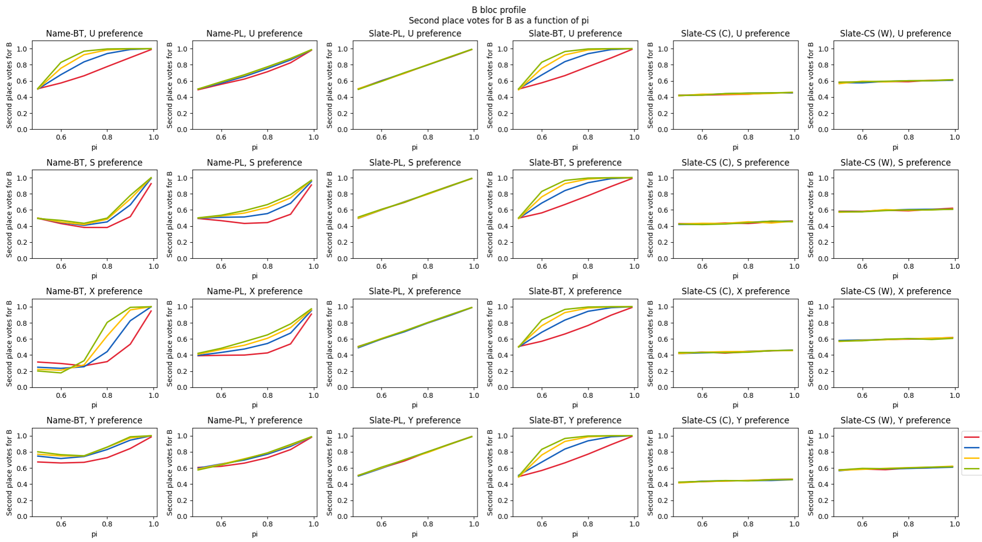

Fig. 30. The proportion of second-place votes for $\mathcal{B}$ candidates across different generative models, numbers of candidates, and strength scenarios. Notice that in the name models, the probability of ranking your own bloc second can actually be less than 50%, even in cases of high cohesion, given particular strength scenarios.

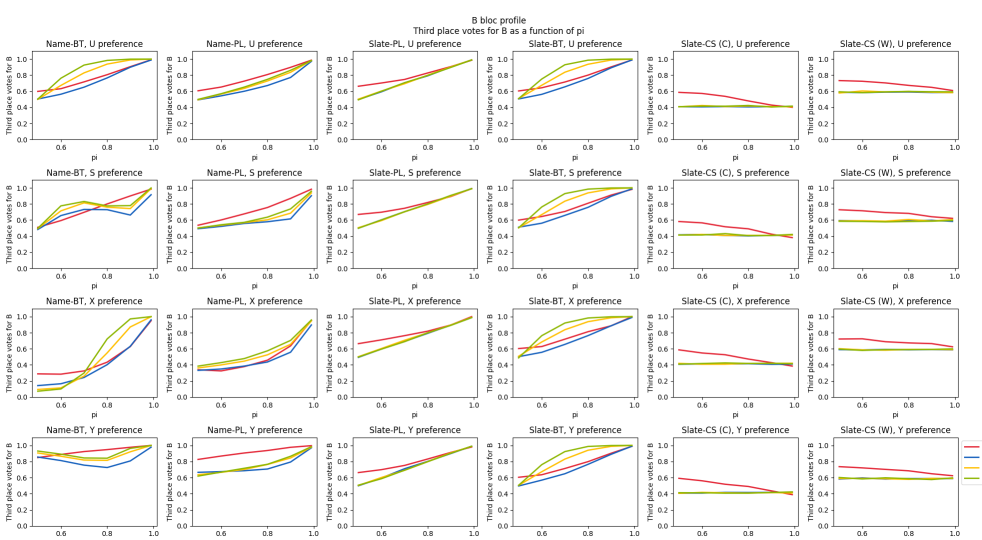

Fig. 31. The proportion of third-place votes for $\mathcal{B}$ candidates across different generative models, numbers of candidates, and strength scenarios.

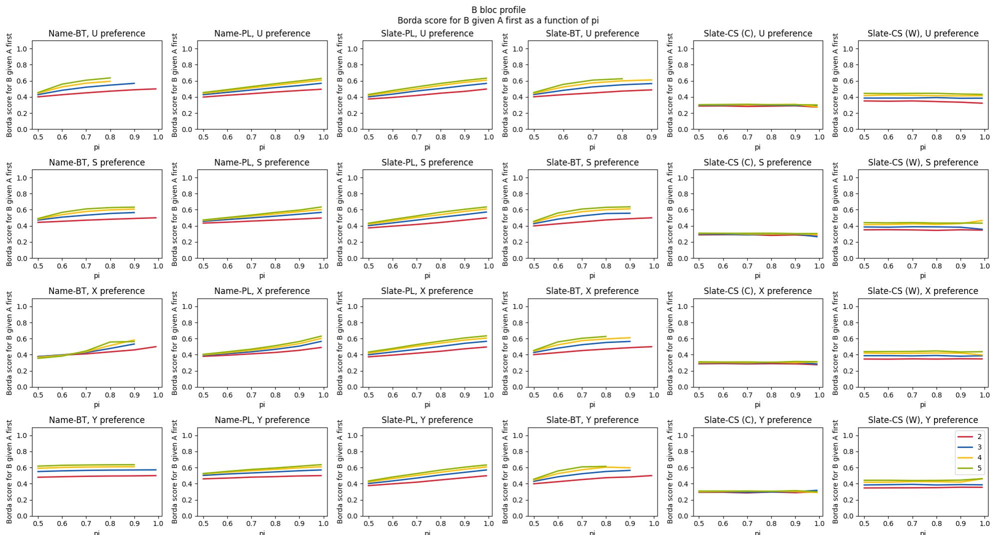

Fig. 32. The proportion of Borda points for $\mathcal{B}$ candidates given that a ballot started with an $\mathcal{A}$ candidate. Shown across different generative models, numbers of candidates, and strength scenarios.

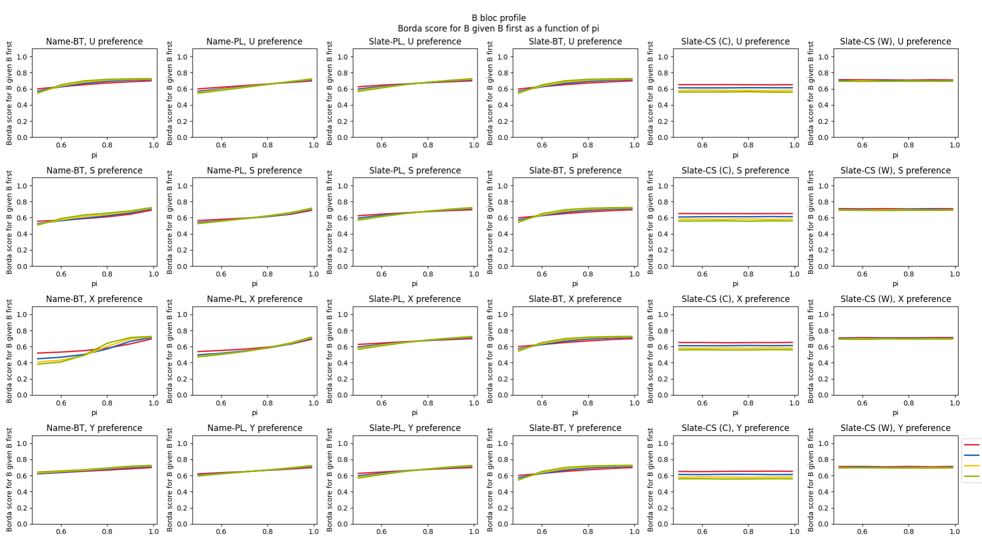

Fig. 33. The proportion of Borda points for $\mathcal{B}$ candidates given that a ballot started with a $\mathcal{B}$ candidate. Shown across different generative models, numbers of candidates, and strength scenarios.

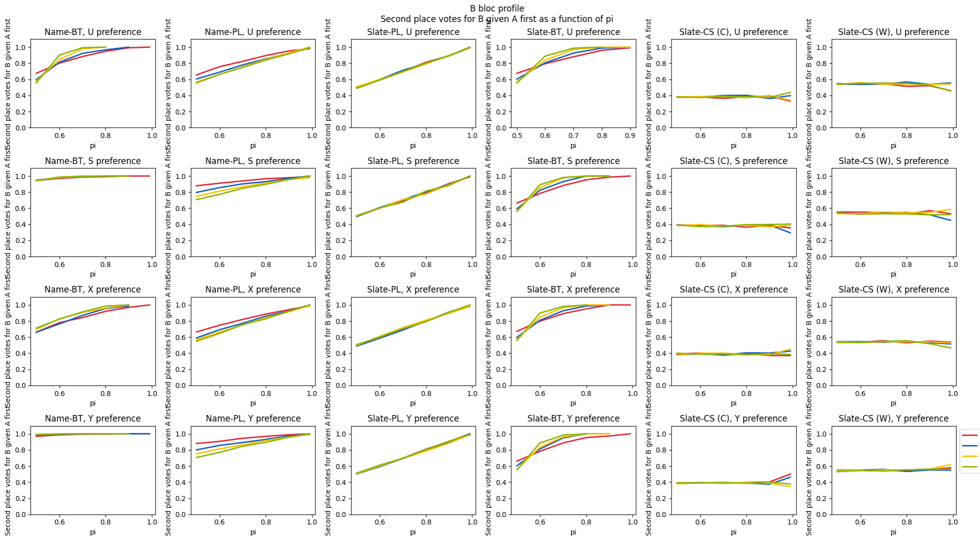

Fig. 34. The proportion of second-place votes for $\mathcal{B}$ candidates given that a ballot started with a $\mathcal{A}$ candidate. Shown across different generative models, numbers of candidates, and strength scenarios.

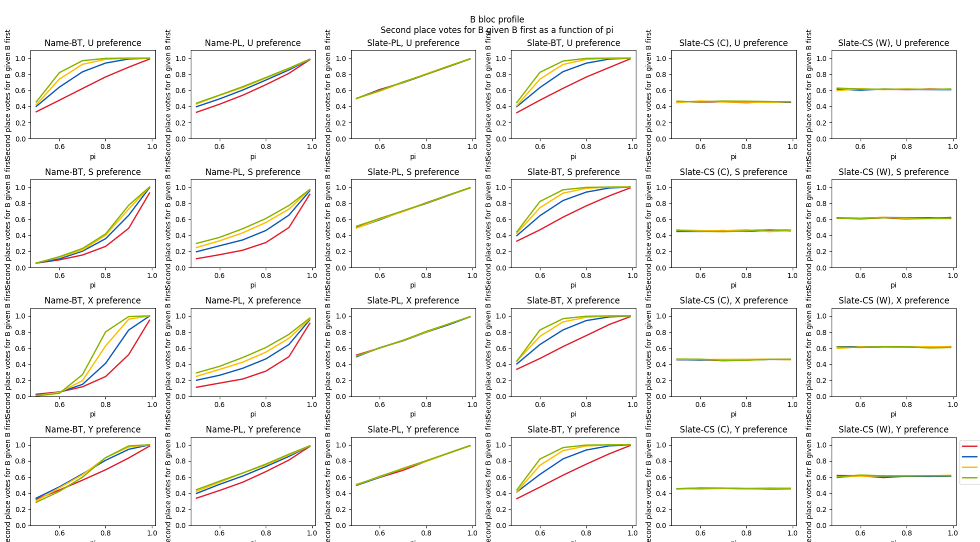

Fig. 35. The proportion of second-place votes for $\mathcal{B}$ candidates given that a ballot started with a $\mathcal{B}$ candidate. Shown across different generative models, numbers of candidates, and strength scenarios.

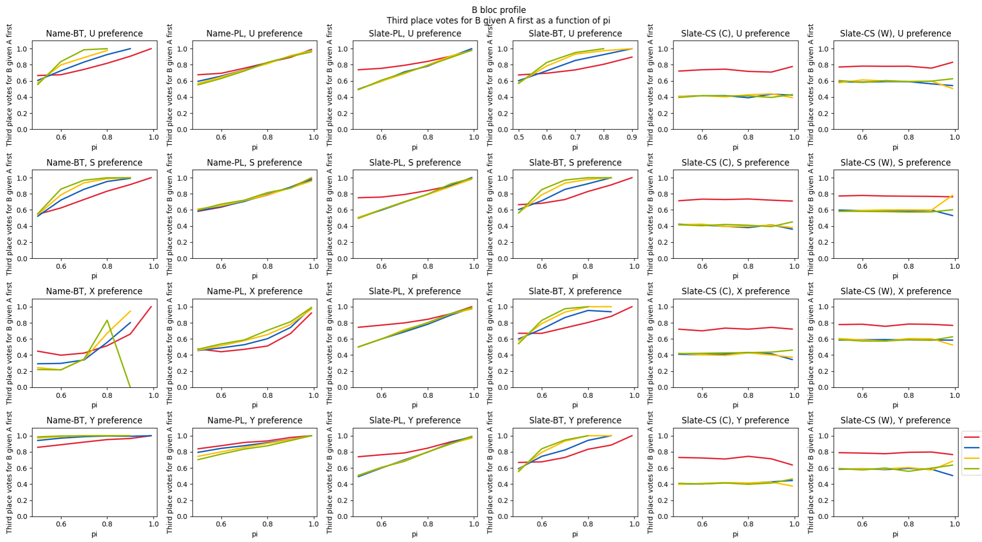

Fig. 36. The proportion of third-place votes for $\mathcal{B}$ candidates given that a ballot started with an $\mathcal{A}$ candidate. Shown across different generative models, numbers of candidates, and strength scenarios.

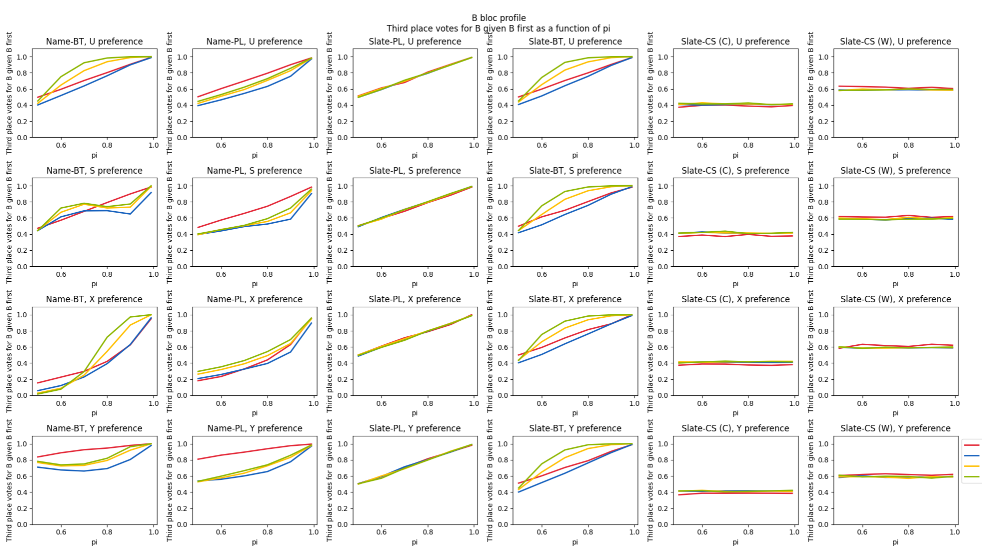

Fig. 37. The proportion of third-place votes for $\mathcal{B}$ candidates given that a ballot started with a $\mathcal{B}$ candidate. Shown across different generative models, numbers of candidates, and strength scenarios.

C MORE MDS PLOTS

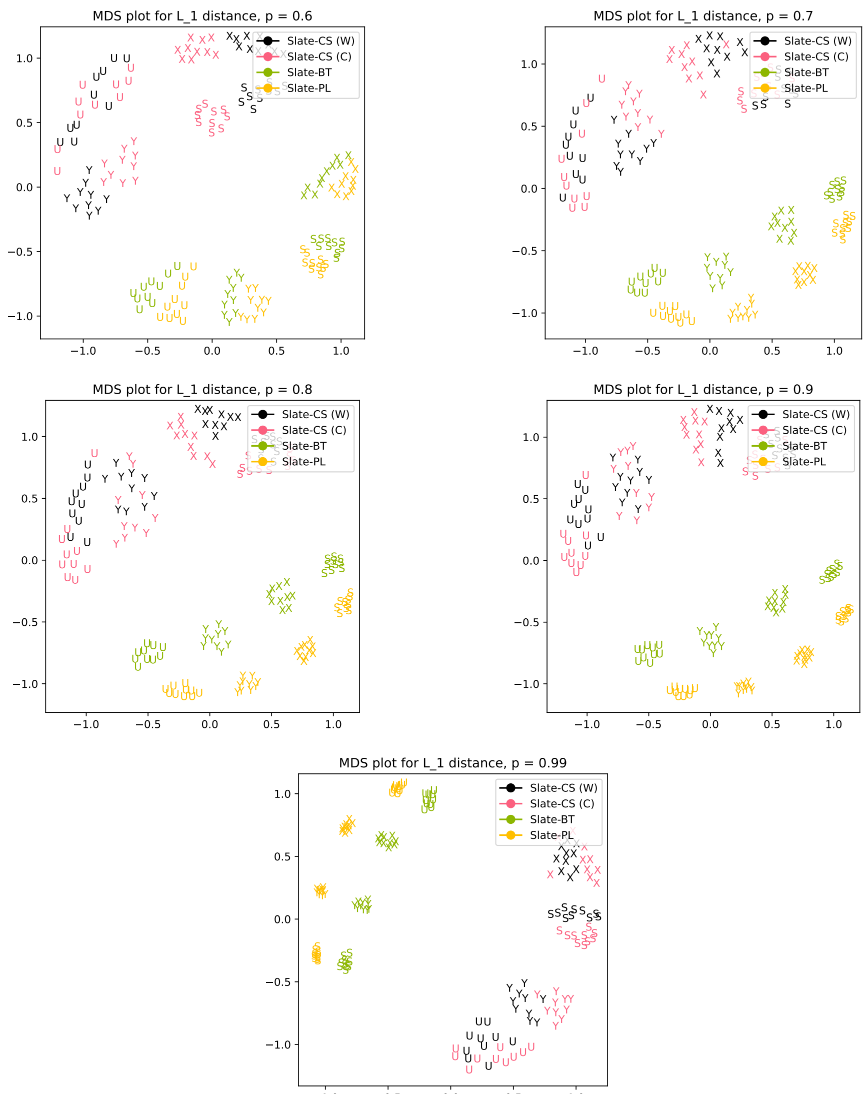

Fig. 38. Multi-dimensional scaling (MDS) plots for profiles with $r = s = 3$ (3 candidates per bloc), under a variety of generative models and candidate strength scenarios. The preference parameters π in each model are chosen to produce an expectation of p first-place votes for one's own slate (which means $\pi = p$ except for BT models, which require calibration). Each model is designated by a different color, and the candidate strength scenarios are denoted U, S, X, Y, as described above. The pairwise distances between profiles are computed with L^1 distance on the profiles. Each preference profile has 1000 ballots, and we have generated 10 profiles by each of the 16 model/strength pairs. As $p \to 1$, the main difference appearing in the models is that the BT and PL profiles become tightly clustered for each candidate strength scenario, while the CS profiles remain more variable.

D FITTING TO SCOTTISH ELECTIONS

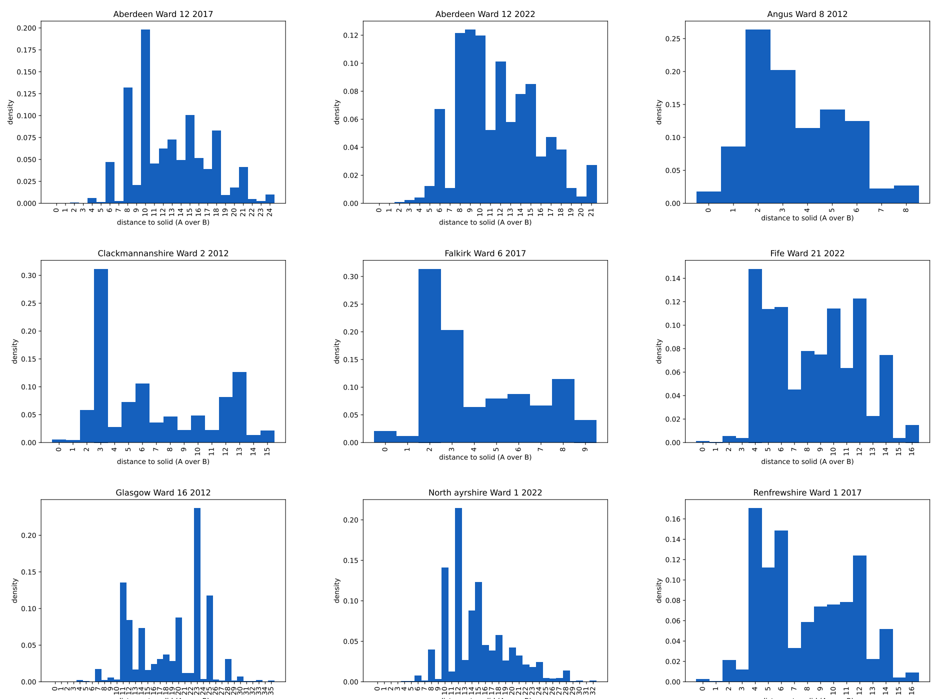

Fig. 39. Histograms showing the distribution of swap distances to solid A-over-B ballots in nine Scottish elections.

To conclude, we provide a full sweep of fitting outputs across the nine elections and various models in this paper.

We start with tables of fitting data for the two elections highlighted in the body of the text: Aberdeen Ward 12 and Falkirk Ward 6, 2017. In addition to optimizing the choice of π_B, we also simply use the first place vote share (FPV) and top-k Borda share.

Aberdeen Ward 12 2017	FPV π_B	d_{Wass}	Borda π_B	d_{Wass}	Opt. π_B	d_{Wass}
Name-PL	0.3332	0.8476	0.4079	1.4446	0.3338	0.8846
Name-BT	0.3332	1.9437	0.4079	1.6426	0.395	1.5827
Slate-PL	0.3332	3.2935	0.4079	1.9335	0.415	1.8544
Slate-BT	0.3332	10.4393	0.4079	7.6536	0.505	**0.5936**
CS ($B = W$)	0.3332	1.9777	0.4079	2.3372	0.0788	**0.5531**
CS ($B = C$)	0.3332	1.2303	0.4079	0.8829	0.445	**0.7082**

Table 2. Wasserstein distances (d_{Wass}) from swap distance distributions of generative models to the swap distance distribution of Aberdeen Ward 12 2017. The smallest three Wasserstein distances are bolded.

For the CS and Slate-type models, the candidate strength does not impact the ballot type and can be ignored. For the Name models, we estimate strength using first-place votes for each candidate.

Falkirk Ward 6 2017	FPV π_B	d_{Wass}	Borda π_B	d_{Wass}	Opt. π_B	d_{Wass}
Name-PL	0.3373	0.9901	0.4287	0.618	0.455	0.5819
Name-BT	0.3373	1.5758	0.4287	0.8319	0.4488	0.7939
Slate-PL	0.3373	1.3142	0.4287	0.4992	0.4388	0.4757
Slate-BT	0.3373	2.3792	0.4287	0.9821	0.4713	0.5025
CS ($B = W$)	0.3373	0.6173	0.4287	0.9211	0.215	**0.4459**
CS ($B = C$)	0.3373	0.7639	0.4287	**0.4553**	0.4638	**0.4237**

Table 3. Wasserstein distances (d_{Wass}) from swap distance distributions of generative models to the swap distance distribution of Falkirk Ward 6 2017. The smallest three Wasserstein distances are bolded.

We use Markov chain Monte Carlo (MCMC) methods to estimate the BT distribution in the two elections with more than 11 candidates since it is costly to compute the probability density function directly. In those cases we sample 10,000 ballots from the MCMC runs. All other simulations use the same number of ballots as in the observed election.

Plots for all elections and models follow.

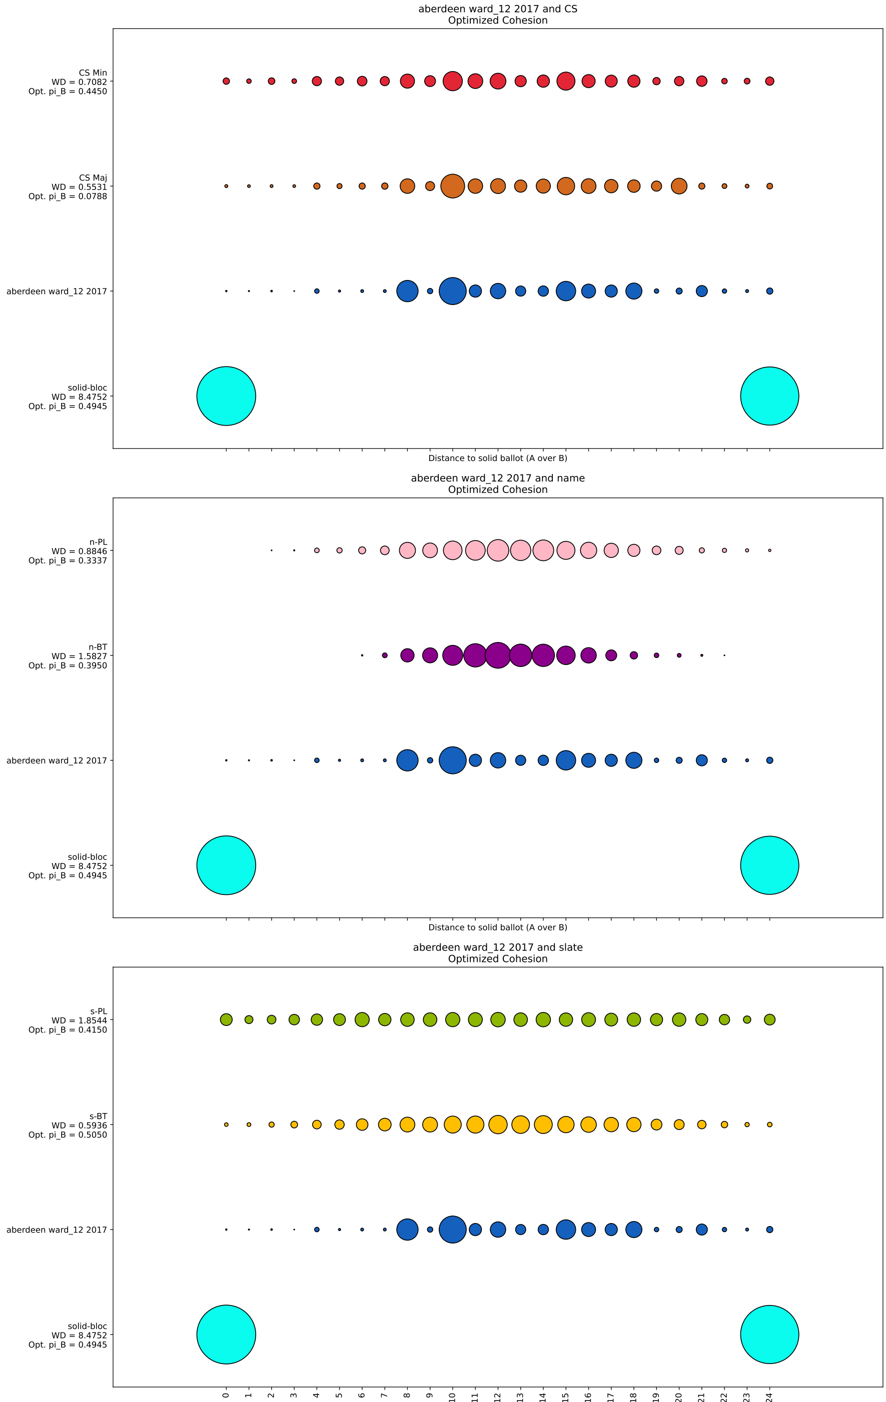

Fig. 40. Bubble plots showing the distribution of swap distances from our generative models, solid-bloc voting, and a real election to A-over-B ballots. Both the generative models and solid-bloc election are optimized via a grid search to choose a value for π_B that minimizes d_{Wass} to the real Aberdeen Ward 12 2017 election.

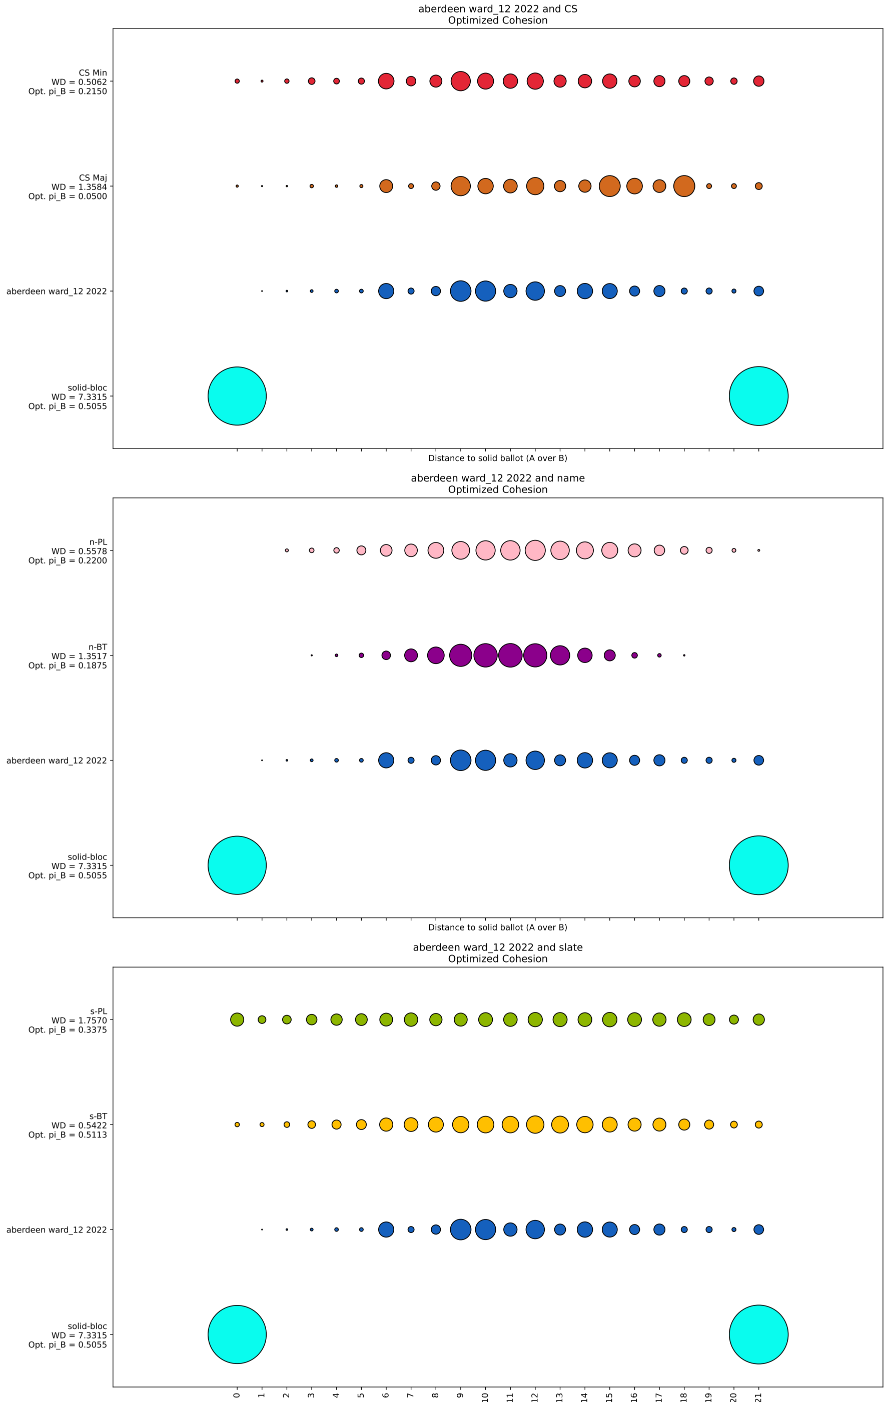

Fig. 41. Bubble plots showing the distribution of swap distances from our generative models, solid-bloc voting, and a real election to A-over-B ballots. Both the generative models and solid-bloc election are optimized via a grid search to choose a value for π_B that minimizes d_{Wass} to the real Aberdeen Ward 12 2022 election.

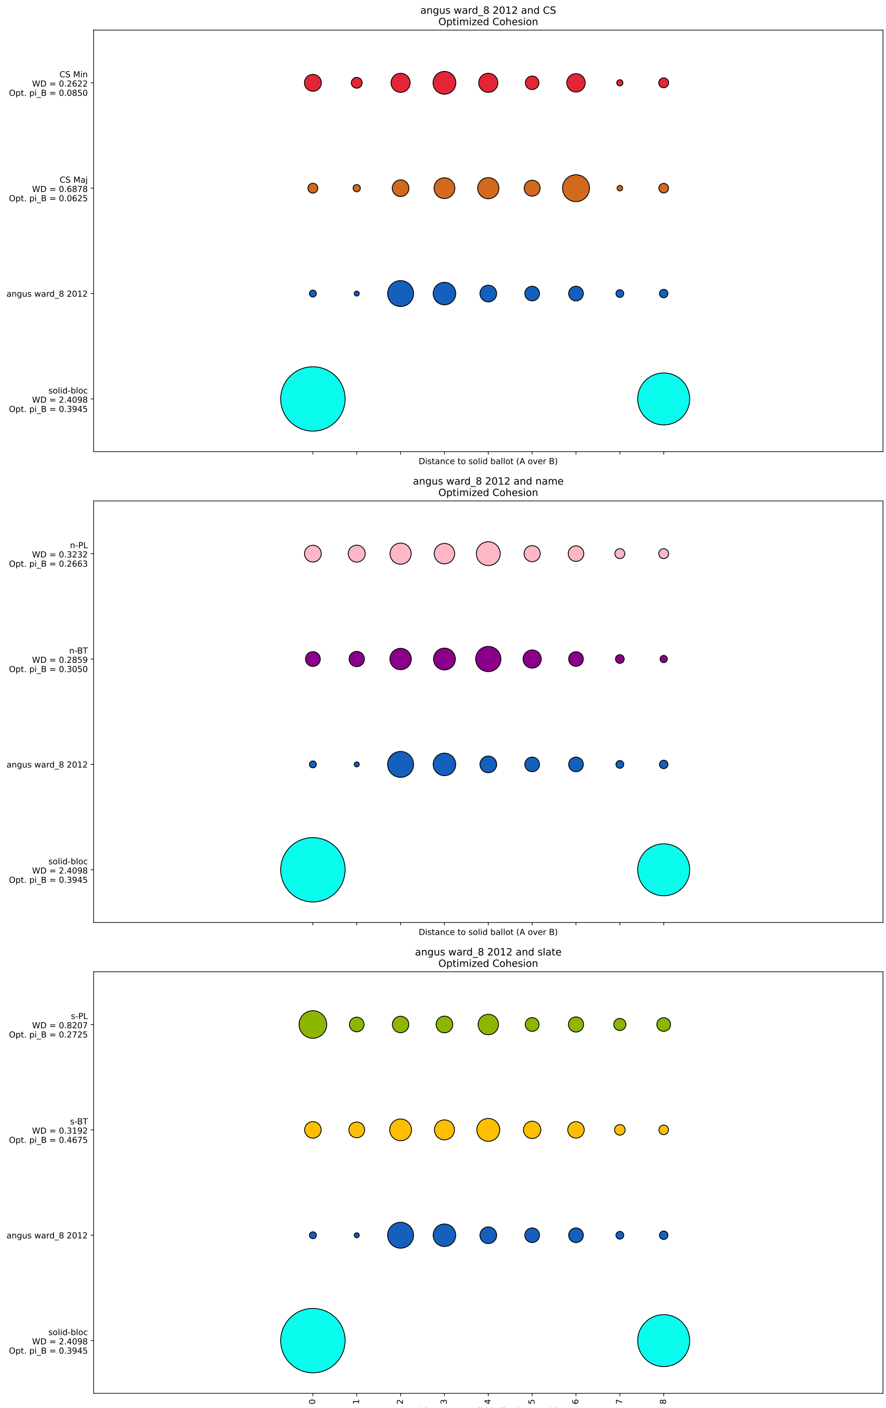

Fig. 42. Bubble plots showing the distribution of swap distances from our generative models, solid-bloc voting, and a real election to A-over-B ballots. Both the generative models and solid-bloc election are optimized via a grid search to choose a value for π_B that minimizes d_{Wass} to the real Angus Ward 8 2012 election.

Fig. 43. Bubble plots showing the distribution of swap distances from our generative models, solid-bloc voting, and a real election to A-over-B ballots. Both the generative models and solid-bloc election are optimized via a grid search to choose a value for π_B that minimizes d_{Wass} to the real Clackmannanshire Ward 2 2012 election.

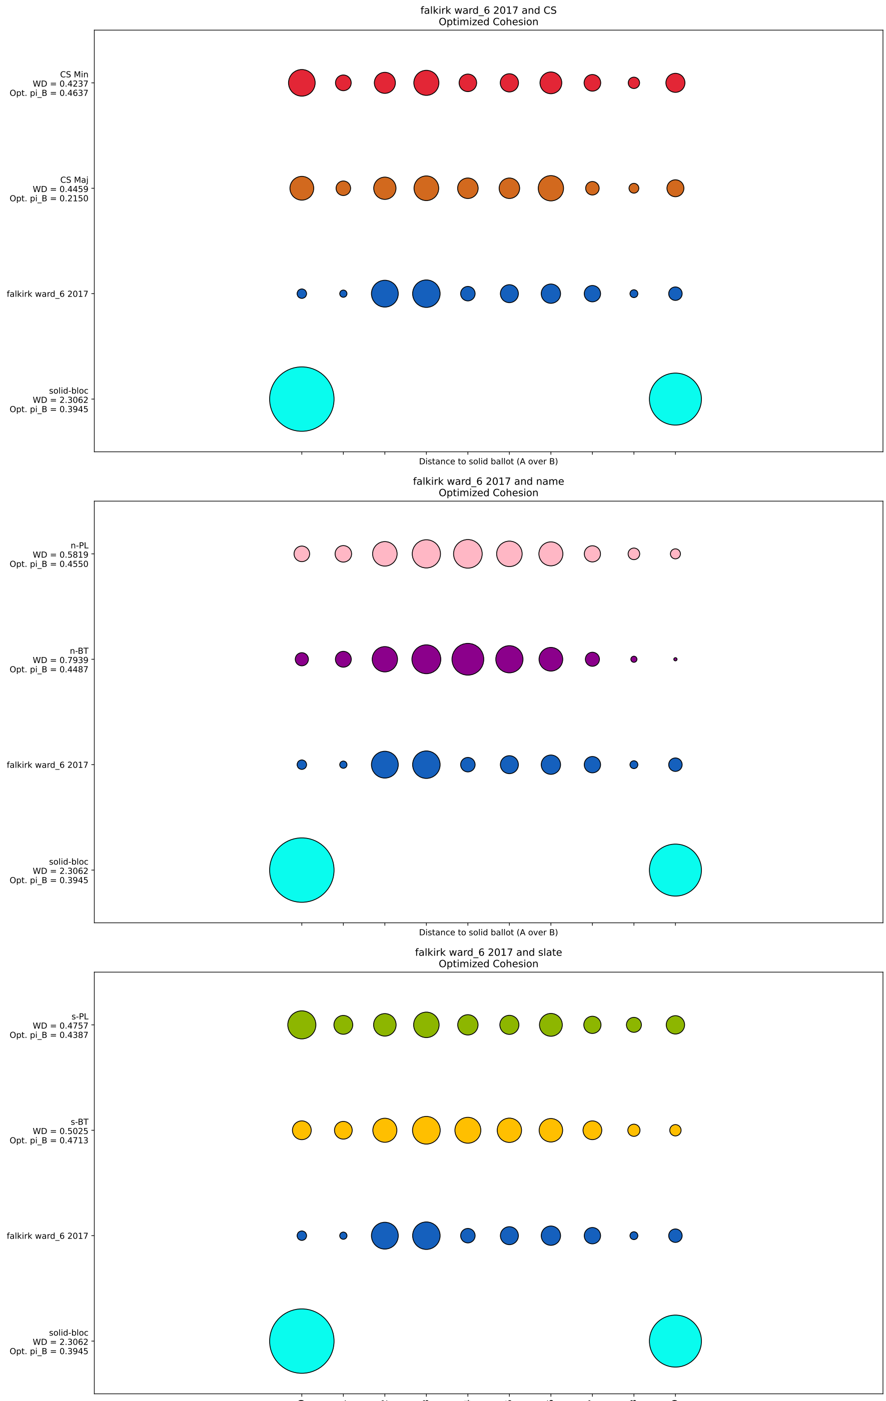

Fig. 44. Bubble plots showing the distribution of swap distances from our generative models, solid-bloc voting, and a real election to A-over-B ballots. Both the generative models and solid-bloc election are optimized via a grid search to choose a value for π_B that minimizes d_{Wass} to the real Falkirk Ward 6 2017 election.

Fig. 45. Bubble plots showing the distribution of swap distances from our generative models, solid-bloc voting, and a real election to A-over-B ballots. Both the generative models and solid-bloc election are optimized via a grid search to choose a value for π_B that minimizes d_{Wass} to the real Fife Ward 21 2022 election.

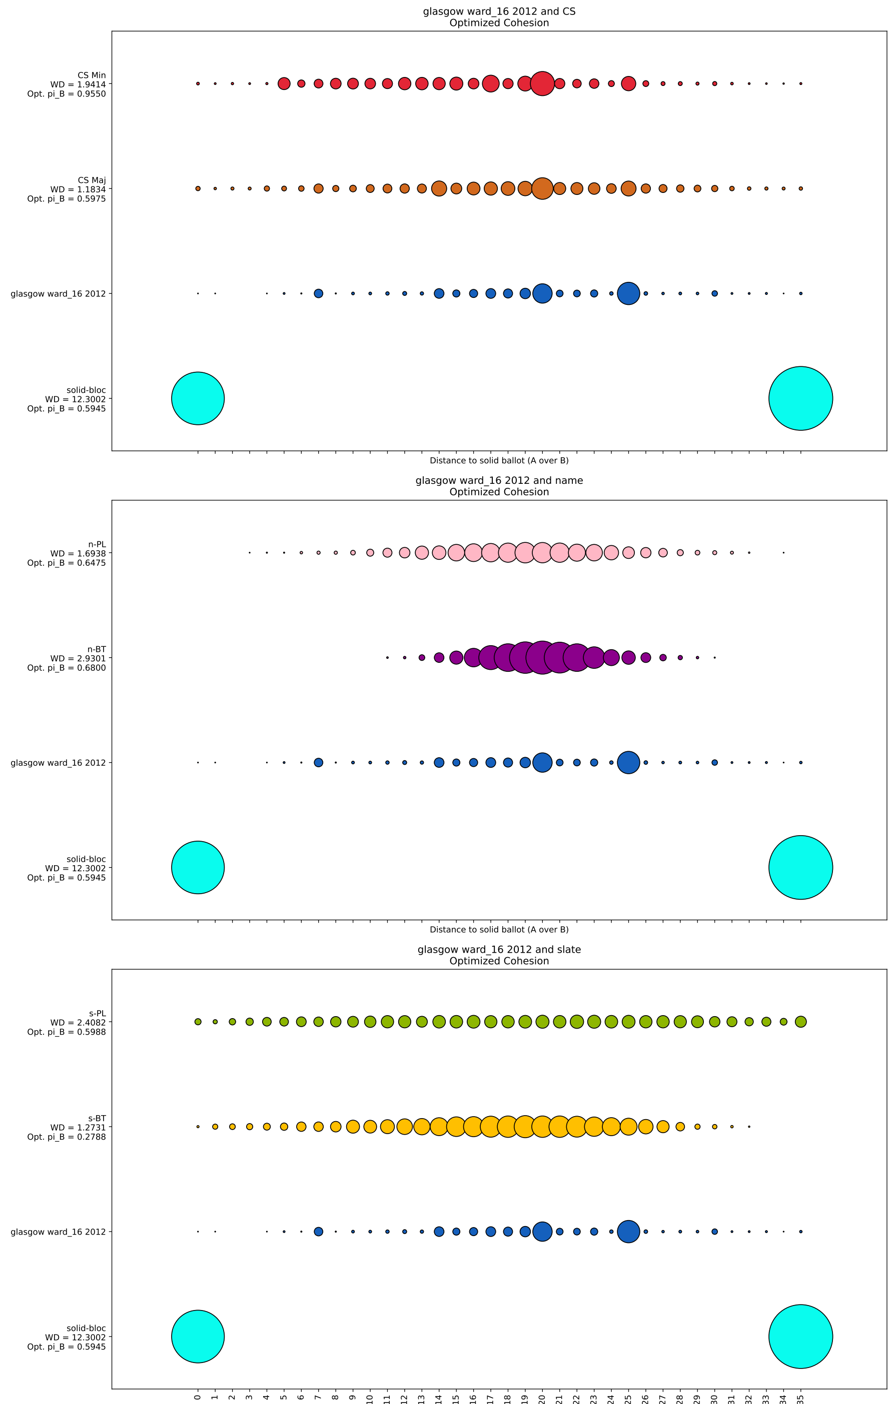

Fig. 46. Bubble plots showing the distribution of swap distances from our generative models, solid-bloc voting, and a real election to A-over-B ballots. Both the generative models and solid-bloc election are optimized via a grid search to choose a value for π_B that minimizes d_{Wass} to the real Glasgow Ward 16 2012 election.

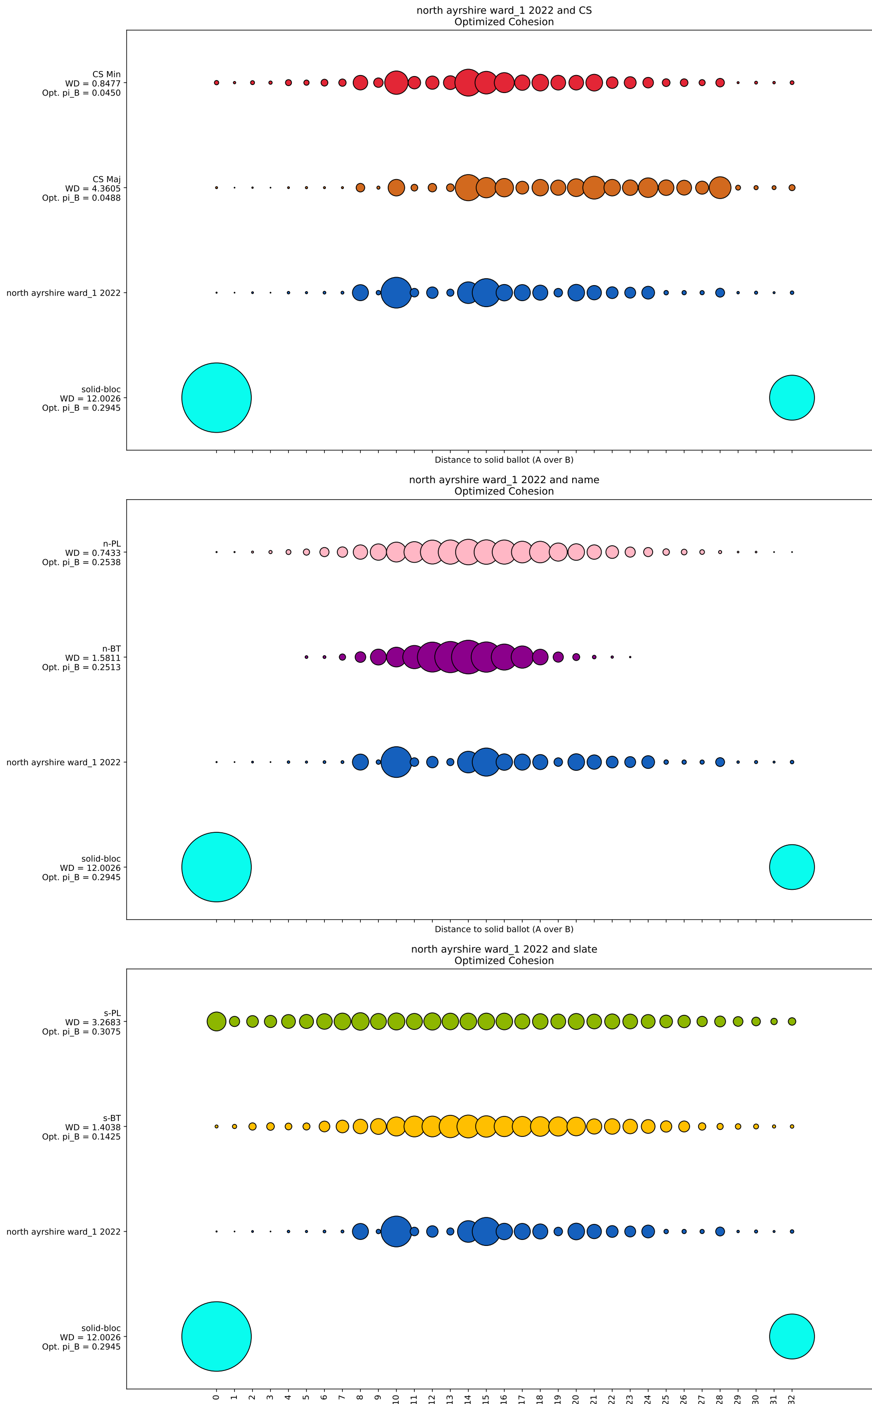

Fig. 47. Bubble plots showing the distribution of swap distances from our generative models, solid-bloc voting, and a real election to A-over-B ballots. Both the generative models and solid-bloc election are optimized via a grid search to choose a value for π_B that minimizes d_{Wass} to the real North Ayrshire Ward 1 2022 election.

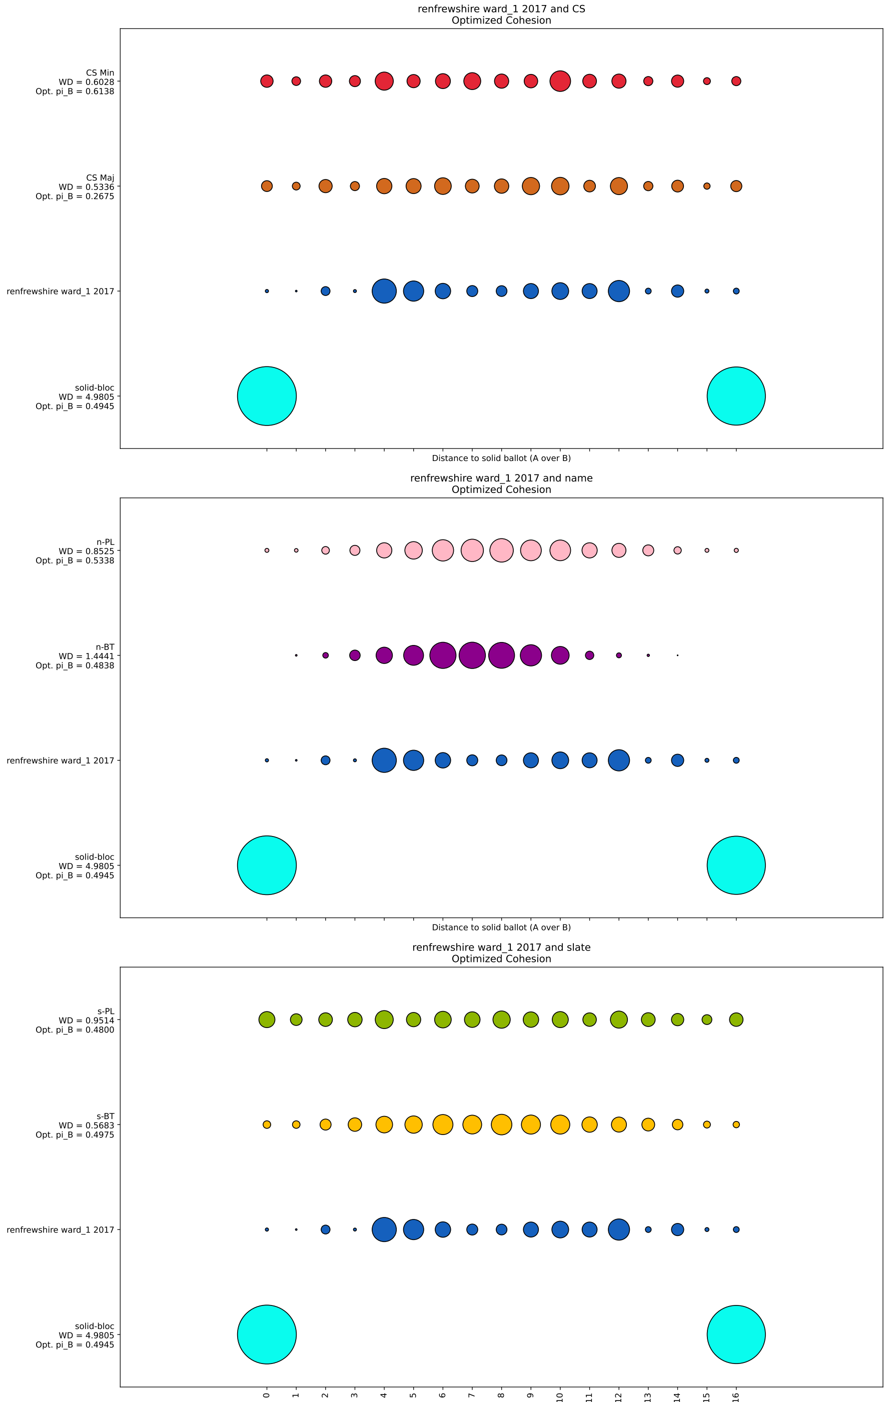

Fig. 48. Bubble plots showing the distribution of swap distances from our generative models, solid-bloc voting, and a real election to A-over-B ballots. Both the generative models and solid-bloc election are optimized via a grid search to choose a value for π_B that minimizes d_{Wass} to the real Renfrewshire Ward 1 2017 election.