

Ranked Choice Voting and Minority Representation

Gerdus Benade¹, Ruth Buck², Moon Duchin³, Dara Gold³, Thomas Weighill⁴

In this paper, we outline and demonstrate a data-driven methodology that voting rights advocates can use to compare the likely effectiveness of a single transferable vote system (STV) to single-member districts (SMD) for securing minority representation in local government. We incorporate both election data and demographics, and can apply variable assumptions on candidate availability and voter turnout. The core of our STV analysis uses four models of voter ranking behavior that take racial polarization into account; to assess districts, we use random district-generation algorithms developed at the MGGG Redistricting Lab.

We demonstrate this method on four case studies: judicial elections in Terrebonne Parish, Louisiana; the county commission of Jones County, North Carolina; and the city councils of Cincinnati, Ohio and Pasadena, Texas. We find that STV provides proportional or slightly better representation for the relevant minority group in each case, while districts vary widely in their effectiveness depending on local circumstances.

Contents

1	Introduction	2
2	Main contributions and findings	4
3	Theoretical results for STV with solid coalitions	6
4	Four models of ranking	8
5	Comparing alternative electoral systems	15
6	Case studies	17
7	Future work	30
8	Conclusion	30
A	How does STV work?	33
B	Notes on the data	34
C	Proof of Theorem 1	34
D	Details of PL and BT models	36

¹Boston University, ²Pennsylvania State University, ³Tufts University, ⁴University of North Carolina at Greensboro

1 Introduction

Electoral systems are methods that convert voter intent into representation, taking ballot data as input and identifying a winning subset of candidates as output. Though it has long been understood that no system exists that meets all fairness criteria simultaneously, many reasonable choices remain to be studied.

At-large plurality elections—in which every voter selects multiple candidates, up to to the number of open seats, and the candidates with the greatest number of votes are elected—have long been regarded as obstructing minority representation. Intuitively, this make sense: if a majority (however slim) of voters bloc-votes on a slate of candidates, then they can elect them all, making the preferences of even a sizeable minority irrelevant. The standard alternative to at-large elections is to use *single-member districts* (SMD) for plurality elections, with the single person who received the most votes carrying each territorial district. Indeed, the particular use of the Voting Rights Act of 1965 to challenge multi-member systems in favor of SMD has a long history. Several influential papers from the 1970s and early 1980s argued that at-large elections were much worse than districts for Black representation in particular [11, 14, 20]. By 1986, the Supreme Court had adopted a set of threshold conditions known as the “Gingles factors” which are explicitly keyed to the investigation of single-member districts as a remedy for systemic racial inequity. Since that time, the VRA has been an extremely successful tool for federal, state, and local elections, with representation for people of color increasing dramatically in the years to follow.¹

The first (of three) Gingles factors requires plaintiffs to show that the minority group in question is sufficiently numerous and geographically concentrated so that it is possible to draw a single-member district in which that group holds a majority.² It is certainly possible for a would-be plaintiff to fail this bright-line test and be denied recourse to the VRA. For example, a group that has one third of the population but is distributed extremely evenly over the jurisdiction might never form a majority in any zone of the relevant size.³

At-large plurality and SMD elections are not the only options, however. Ranked choice voting is a family of systems in which voters rank candidates in order of preference. In this paper we will focus attention on *single transferable vote* (STV) systems, a smaller family within ranked choice. In STV elections, there is a threshold level of support needed to be elected, depending on the number of seats to be filled. As candidates are either elected (by passing the threshold) or eliminated from contention, the votes supporting those candidates are transferred to the next options on their respective ballots. Specific mechanics vary; in this report we will focus on the vote-tallying mechanism used by Cambridge, MA for its City Council elections (see Appendix A for details). Since we are concerned with representation on multi-member local bodies such as judicial courts, county commissions or city councils, we will focus on that case rather than on the use of a similar mechanism to elect a single winner (which is sometimes known as *instant runoff voting*, or IRV). We note that STV can be employed for at-large elections when the size of the elected body is modest, or

¹Though originally conceived largely to protect the Black vote, the VRA is now used on behalf of a variety of racial, ethnic, and language minority groups. More recently, studies of the shift from at-large voting to districts in California elections brought on by the California Voting Rights Act have claimed to identify positive effects on representation for the Hispanic population [1, 6].

²Through subsequent court cases, the use of citizen voting age population (CVAP) has become standard for assessing a population majority, particularly for Latino and Asian plaintiffs.

³See [21, 10] for more discussion of the problem of districted representation for dispersed groups.

that it can be carried out in multi-member districts to fill a large representative body.⁴

Clearly, at-large STV avoids some of the drawbacks of districts: it is immune to gerrymandering (the agenda-driven drawing of district boundaries) and can represent both dispersed and concentrated population groups. Indeed, it is widely believed that STV tends to produce representation for various groups in rough proportion to their numbers; in some parts of the world (particularly Australia and Ireland), STV is even called by the name “PR,” or proportional representation. The social choice literature uses the term *Proportionality for Solid Coalitions* to describe this phenomenon—a bloc-voting group which unanimously prefers some collection of candidates over their competitors will secure at least their proportional share of representation—and sometimes treats this as a fairness axiom. We state this for STV in Section 3 as Theorem 1 and prove it in Appendix C.

Reform advocates will often want to predict how STV might fare in a particular jurisdiction, as opposed to SMD or some other electoral system. Needless to say, the comparison depends on many factors that are hard to model. Voting behavior, including turnout levels, can change as a result of electoral reform (see for instance [2]). But even holding voters and their preferences fixed from one system to another, it is impossible to infer rankings from a history of non-ranked elections. Contests where voters choose a single candidate tell us nothing about their second choices: a candidate could in theory be a universal second choice without ever appearing in a single cast vote. And at-large plurality ballots similarly reveal nothing about how voters would have ranked their selected candidates. So election history from jurisdictions with ranked voting is the most useful for predicting STV outcomes. Unfortunately that data is extremely limited; for instance, Cambridge, MA is the only U.S. city currently electing its city council using ranked choice, and many IRV jurisdictions only allow voters to rank up to three candidates.

Despite these obstacles, effective reform efforts must try to estimate the effect of STV on representation in a theoretically sound way. Waiting for more jurisdictions to adopt ranked voting is impractical, since the very adoption may depend on its perceived effect on representation. The MGGG Redistricting Lab has already been engaged in assessing likely STV outcomes in Santa Clara, CA [16], Chicago, IL [18], Lowell, MA [17] and Yakima County, WA [19]. In the current project, we significantly expand the ideas in this earlier work, developing four models for inferring voter rankings in the presence of various degrees of racially polarized voting (RPV).⁵ Our models are designed to incorporate any techniques for assessing RPV that are imported from voting rights scholarship and litigation.⁶

In this report, we begin by outlining some theoretical results from the social choice literature, then introduce the four statistical ranking models that allow us to move beyond the restrictive setting of the classical results. In Section 5, we outline our methodology for employing these models to compare alternative electoral systems for a jurisdiction with single-winner electoral history and known demographics. To demonstrate how this STV modeling works in practice, and to extract some important initial lessons about comparing STV to SMD outcomes, we present an in-depth

⁴Many authors recommend that the *magnitude*, or number of representatives elected, be between three and five to remain manageable for voters. For example, see [18] for a discussion of electing the 50-member Chicago City Council from ten 5-member districts. Cambridge uses at-large STV to elect a 9-member council, which imposes a voter burden that is arguably very high.

⁵Voting is called *racially polarized* when a minority racial group is found to cohesively support one set of candidates, while the majority group votes as a bloc for other candidates.

⁶The industry standard technique for assessing RPV is a suite of methods called *ecological inference*, or EI. Our application here is agnostic to whether the user selects a flavor of EI or uses some other method to estimate polarization.

study of four jurisdictions: Terrebonne Parish, Louisiana; Cincinnati, Ohio; Jones County, North Carolina; and Pasadena, Texas. As we will explain in Section 4, these were chosen because of some interesting feature of the location and representative body.

Note on terminology: We are interested in representation for a variety of groups in the United States who have historically been denied equal and fair political representation. Our four case studies focus on Black and Hispanic/Latino populations. In order to streamline the presentation of models and methods, we will refer to members of the population whose representation is being studied as “POC” (people of color), and the rest as “White.” We will also use the shorthand of referring to POC candidates, though to be more precise this always refers to candidates of choice for the minority group, whether or not these candidates identify as members of the group. Finally, we adopt the standard convention that “Black” refers to non-Hispanic Black population.

2 Main contributions and findings

The main contribution of this project is to introduce, refine, and apply the methodology outlined in Section 5 to compare SMD to STV systems for local jurisdictions. This methodology makes primary use of a suite of four ranking models presented Section 4, which are implemented in a user-friendly shiny app at vrdi.shinyapps.io/rcv-app. The reader can also find a Python codebase for larger-scale investigations at github.com/mggg/minority-RCV.

This paper uses four case studies. Three localities were chosen because of interesting Voting Rights Act legal challenges in the last decade (Terrebonne Parish, Jones County, and Pasadena). For the fourth, we wanted both a larger population and a large-magnitude electoral body; we selected Cincinnati because of its long history with RCV, as detailed further below. Our case studies are far from an exhaustive picture of possible settings for RCV, but together with previous work in Santa Clara, Chicago, Lowell, and Yakima, we are starting to put together a strong portfolio that is varied in scale, demographics, and region.

Jurisdiction	POC group	CVAP share	Typical STV outcomes	Favorable districts
Terrebonne Parish, LA	Black	.18	1–2 out of 5	0 out of 5
Cincinnati, OH	Black	.39	3–5 out of 9	7 out of 9
Jones County, NC	Black	.33	1–2 out of 5 2–3 out of 7	1 out of 5 1 out of 7
Pasadena, TX	Latino	.53	3–5 out of 8 2–4 <i>out of 8</i>	7 out of 8 3 <i>out of 8</i>

Table 1. Summary of findings in four case studies (see Figure 1 for a visualization). Typical STV outcomes capture the levels usually observed across a range of parameter settings and scenarios. Favorable districts records the plan with the largest number of districts that are majority-minority by citizen voting age population (CVAP) that could be found with an algorithmic search. Two different district magnitudes are considered in Jones County, for reasons described in the case study below. The italicized row highlights outcomes under a scenario of extremely low Latino turnout in Pasadena.




Figure 1. This visualization for Table 1 illustrates that the best-possible districted outcome is highly variable across our case studies, while the performance of STV is consistently in line with proportional representation of the POC population. Blue: the CVAP share for the identified POC group. Red: the range of typical STV outcomes. Amber: the most favorable districting plan found by a search, under favorable turnout assumptions. Gray: performance of favorable districting plan under low POC turnout.

Summary of findings

By investigating the models in a range of hypothetical scenarios (Section 4.6) across four case studies (Section 6), we highlight three main findings about the relative performance of STV and SMD systems.

1. **STV systems tend to elect POC candidates of choice in proportion to POC population.** In all four case studies, STV is predicted to award POC-preferred candidates a fraction of the seats roughly equal to the POC population share. This is due to the structural seats-to-votes proportionality of the STV system, in addition to observed levels of crossover voting by POC and White voters that roughly cancel out, across multiple reasonable models of voter ranking behavior.
2. **The range of representational outcomes in a districted system is highly sensitive to the size and distribution of the minority group.** In two cases, even the most proactive districting plans are dramatically sub-proportional (with a complete shutout in one case). In the other two, carefully drawn districts can in principle elect a super-majority of POC-preferred candidates. By contrast, STV is predicted to secure a roughly proportional result in every case, and is by construction independent of the level of geographic concentration of the POC population.
3. **The efficacy of single-member districts is especially questionable in the case of low turnout.** The case of Pasadena demonstrates that majority-minority districts are a risky proposition when turnout is low for a minority group. Districts near 50% can prove ineffective if the minority group has low relative turnout, which means that drawing effective districts requires a good understanding of the turnout in different parts of the jurisdiction.

We emphasize that our intent is to supply a toolkit for voting rights advocates, rather than a set of standalone conclusions. When a new locality is being assessed, it should be run through the models outlined here. The human geography will have a strong impact on the possible performance of districts, and in that regard every locality is unique. The STV modeling does not depend

on local geography, but does depend on local electoral conditions, including candidate availability from the minority group and the degree of agreement about which candidates are strongest. By making the models transparent and making the software tools available, we intend to make it easy to explore projections across a range of scenarios.

3 Theoretical results for STV with solid coalitions

Single transferable vote is often referred to informally as a “proportional representation system,” but it is important to distinguish it from common systems outside of the U.S., such as *party list* voting, that are necessarily proportional by design. STV, rather, is said to have a tendency to produce proportional outcomes, which is part of what the present paper aims to illuminate.

Unlike party-based proportional systems, STV requires no political parties or other official groupings of candidates for cohesive groups to gain representation. Any shared identity or interest of a sizeable group of voters can be translated to electoral success as long the group largely agrees on a corresponding subset of candidates.

The theoretical underpinning of this statement is a result known as Proportionality for Solid Coalitions, which we state here. A group of voters is called a *solid coalition* for a group of candidates \mathcal{C} if every member of the group always ranks candidates in \mathcal{C} above candidates not in \mathcal{C} . We will call \mathcal{C} the group’s *preferred candidates*. In this section, we will assume that all voters rank all candidates. A scenario in which the voting population decomposes into two groups voting solidly for disjoint sets of candidates is called *total polarization*.

Suppose there are m seats (sometimes called the *magnitude* of a multi-member district) and we have N voters. STV systems are conducted in rounds of retabulation to see which candidates have exceeded a threshold number of first-choice votes; we will employ the so-called *Droop quota*

$$t = \left\lfloor \frac{N}{m+1} \right\rfloor + 1$$

and defer further STV details to Appendix A.⁷ Any candidate who has more than t first-choice votes during the vote tallying process is designated as a winner.

Theorem 1 (Proportionality for Solid Coalitions). *Suppose t is the Droop quota for a particular election with N voters and m seats to be filled. Consider a solid coalition of at least kt voters for some whole number k , and suppose the coalition has at least k preferred candidates. Then at least k of the coalition’s preferred candidates will be elected.*

A full proof can be found in Appendix C. For realistic (large) values of N , the quota t is roughly $N/(m+1)$, so it is less than N/m . It follows that a solid coalition is guaranteed to elect at least a roughly proportional share of preferred candidates, assuming that there are enough preferred candidates to make this possible. This is just a lower bound; the coalition may have more electoral success, depending on support from other voters.

This is a promising feature of a system designed to make multi-winner elections more representative of the voting population. However, in practice few groups of voters vote 100% cohesively.

⁷Recall that $\lfloor x \rfloor$ denotes the largest whole number $\leq x$, i.e., rounding down to the nearest integer.

This necessitates the use of statistical models such as those introduced below, which can account for crossover voting patterns and for uncertainty.

The implications of strategic voting through so-called *vote-management systems*, also called *spreading the preferences*, is an important topic of study in the RCV literature. With appropriate orchestration, solid coalitions can elect the maximum number of candidates possible (recall that Theorem 1 only gives a lower bound on representation, not an exact prediction). See, e.g., [5, 4] for studies of strategic voting in Scottish elections, and [13] for a look at three other countries.

We will focus on a limited question: if the voting population is divided into two solid coalitions who share no preferred candidates, how many seats does each coalition secure? The following carefully designed example shows that the answer may depend on the details of the rankings.

Example 1. *Consider a contest in which $N = 400$ voters are made up of two solid coalitions of equal size, and they must elect for $m = 3$ positions, so that the Droop quota is $t = (400/4) + 1 = 101$. Suppose that coalition A supports candidates A_1 and A_2 , while Coalition B supports candidates B_1 and B_2 , with specific rankings as follows.*

$\times 200$	$\times 100$	$\times 100$
A_1	B_1	B_2
A_2	B_2	B_1
B_1	A_1	A_1
B_2	A_2	A_2

This notation means that 200 voters prefer $A_1 > A_2 > B_1 > B_2$, and likewise for the other two columns. Candidate A_1 exceeds the quota and is elected, with 99 of their votes then transferred to A_2 . In the next round, A_2 is eliminated for having the fewest first choice votes, leaving winning candidates $\{A_1, B_1, B_2\}$. This shows that the specific rankings of preferred candidates, and not just the size and solidity of the coalition, impacts the outcome.

This is a stylized scenario, however. For most group sizes, each coalition will elect the number of representatives guaranteed by Theorem 1 and no more. The only exceptions occur when the size of one coalition is very close to an integer multiple of the threshold. This is made precise in the following corollary, proved in Appendix C.

Corollary 1 (Exact Proportionality for Two Coalitions). *Suppose t is the Droop quota for an election to fill m open seats. Suppose that the voters are divided into two solid coalitions, A and B, who share no preferred candidates. Suppose that N_A , the number of voters in coalition A, satisfies*

$$(k + 1)t - m - 1 \geq N_A \geq kt$$

for k a whole number. Then the winners will include exactly k candidates preferred by coalition A and $m - k$ by coalition B, as long as this many preferred candidates are available to each group.

To get a sense for the limited relevance of vote management systems in two-party settings, consider a hypothetical Cambridge, MA election with two opposing solid coalitions each supporting nine candidates. These elections fill $m = 9$ seats and typically receive roughly 22,000 votes, so the threshold for election is $t = 2201$. Corollary 1 tells us that for 21,910 of the 22,001 possible coalition sizes (99.6%), vote management systems will not impact the composition of the elected candidates in terms of preferred groups.

In summary: under total polarization, we can typically exactly calculate the number of preferred candidates elected by each group as $\lfloor N_A/t \rfloor$, the size of the group divided by the quota, rounded down.

4 Four models of ranking

We now describe the four models of voter rankings proposed in this report and applied in the case studies below. A graphical representation of these models can be found in Figure 2.

4.1 Overview

From now on we will use C_1, C_2, \dots to symbolize the candidates preferred by the POC group as a whole, and W_1, W_2, \dots the other candidates. We distinguish between a *ballot type*, which is just a string encoding the group membership of each of the candidates selected by a voter (such as CCWCWW), and a *detailed ballot*, which describes which specific candidates occupy each position (such as $C_1C_3W_2C_2W_1W_3$). Our simulations must produce detailed ballots for each voter. The Plackett–Luce (PL) and Bradley–Terry (BT) models described below, based on classic ranking schemes in the social choice literature, will combine polarization parameters and candidate strength parameters to generate detailed ballots probabilistically. Two other models described below, an alternating crossover (AC) and a Cambridge sampler (CS) model, will first choose a ballot type from a prescribed list (with the choice based on estimated polarization) and then fill in candidate names (with the choice based on a candidate strength scenario) to build the detailed ballot.

All four models take polarization parameters π_C, π_W as input which record the likelihood of POC and White voters to vote for the set of candidates preferred by the group as a whole. That is, $\pi_C = \pi_W = 1$ would be the total polarization scenario from the last section, while $\pi_C = .8$ would indicate 80% group cohesion for people of color in supporting the same set of candidates. Using these parameters allows us to replace the unrealistic scenario of total polarization with more moderate group cohesion inferred from standard RPV methods. We display the polarization parameters in a table as follows.

	Candidates $\{C_i\}$	Candidates $\{W_i\}$
POC voters	π_C	$1 - \pi_C$
White voters	$1 - \pi_W$	π_W

Typically, but certainly not always, there is a clear group of POC-preferred candidates and the majority group prefers other candidates, so that $\pi_C > .5$ and $\pi_W > .5$.⁸ We describe how to estimate these support values from historical elections in Section 5. Note that the rows sum to one (since each group’s total support for all the candidates should sum to one), but that the columns need not (since C candidates may enjoy more support overall than W candidates or vice versa).

⁸For instance, the recent MGGG study of Yakima County, WA found that $\pi_C \approx .7$ and $\pi_W \approx .8$ for Spanish-surnamed candidates $\{C_i\}$ [19].




Figure 2. Graphical representations of the four models of ballot generation used in our case studies.

The PL and BT models assign detailed ballots to each voter by randomly dividing up the support encoded by the polarization values among candidates. We can think of this as the degree of consensus within each group about the strong candidates in the candidate pool. This division of support is governed by numerical parameters $\alpha_{CC}, \alpha_{CW}, \alpha_{WC}, \alpha_{WW}$. Here, the intermediate setting $\alpha = 1$ means that all divisions of support are equally likely, so splitting .8 support between candidates C_1 and C_2 as .1 + .7 is on equal terms with .4 + .4. Low values $\alpha < 1$ favor the existence of a consensus preferred candidate, such as a .7 + .1 split; and high values $\alpha > 1$ tend increasingly toward the equal weighting .4 + .4, so that any permutation of candidates is as likely as any other. For instance, $\alpha_{WC} = 2$ would mean that White voters, when they vote for POC-preferred candidates, have no strong tendency to favor C_1 over C_2 or vice versa. The other α values are defined similarly. For further details on the use of these parameters, see Appendix D.

No matter the parameter settings (outside of total polarization), the PL and BT models described below are capable of producing any detailed ballot, though some outcomes are vanishingly unlikely. For the CS and AC models, both the ballot types and the candidate strength scenarios are limited, as described below.

4.2 Plackett-Luce (PL): Individual draws

Plackett-Luce models are standard in the statistical ranking literature. Voter behavior is based on a *support vector* of n numbers summing to 1 that encodes relative likelihood to vote for each of n candidates. A voter fills in a ballot from top to bottom by drawing from a distribution weighted over the available candidates according to the values in the support vector. Appendix D contains further details.

For our PL model, one support vector will be fixed for all POC voters and a second for all White voters, used to build a detailed ballot for each voter. To build the support vector for each group, we begin with the overall polarization data π_C, π_W inferred from our RPV techniques, then randomly divide up the total support among C candidates and among W candidates by drawing weights from a symmetric Dirichlet distribution. This allows us to interpolate between even and uneven division of support among the candidates according to a candidate strength parameter α .

Let's look at an example. Suppose $\pi_C = .8$, meaning that POC voters tend to vote for C candidates 80% of the time, and for W candidates 20%. And suppose White voters tend to vote for C candidates 30% of the time and for W candidates 70% of the time ($\pi_W = .7$). One possible way to divide up these numbers is as follows.

	C_1	C_2	W_1	W_2
POC voters	0.64	0.16	0.11	0.09
White voters	0.1	0.2	0.3	0.4

Note that the POC voters divide their support for the W candidates roughly evenly, while the support for C candidates is heavily skewed towards Candidate 1. This is consistent with what we would expect if the candidate strength parameters were $\alpha_{CW} = .5$ and $\alpha_{CC} = 2$.

4.3 Bradley-Terry (BT): Paired comparison

Under a Bradley-Terry model, support vectors are used to calculate *head-to-head* preferences between candidates. For example, if Candidate X has a support value of 0.4 and Candidate Y has support 0.2, then a detailed ballot with X ranked above Y should be twice as likely to be formed as the reverse. Each hypothetical ballot of length n is weighted in proportion to the product of its pairwise support ratios. See Appendix D for details.

4.4 Alternating crossover (AC)

The AC model was introduced by MGGG in previous RCV studies [17, 19]. The model posits that each group consists of two types of voters: bloc voters and crossover voters. *Bloc voters* vote for in-group candidates first and then out-group candidates. *Crossover voters* select an out-group candidate first and then alternate between in-group and out-group candidates. For example, in a race with three POC-preferred candidates and three White-preferred candidates, the options are shown in Table 2.

Voter group	Voter type	Ballot type
POC	bloc	CCCWWW
POC	crossover	WCWCWC
White	bloc	WWWCCC
White	crossover	CWCWCW

Table 2. The four AC ballot types for a race with 6 candidates, three POC-preferred and three White-preferred.

Note that the full ballot type is determined entirely by the group membership of the voter and the first-choice candidate type. A natural choice is to set the share of bloc voters to match the polarization parameter. For example, if White voters support W candidates at 70%, then 70% of White voters are modeled as bloc voters and 30% are crossover.

From the ballot type, we proceed to a detailed ballot. Following [17], we consider four different scenarios. (These were chosen as anecdotally plausible situations of interest and are just four of an infinite variety of possible scenarios.)

- **Scenario A.** All voters agree on the rank order of the POC candidates and the White candidates.
- **Scenario B.** All voters agree on the rank order of the White candidates, and all White voters agree on the ordering of the POC candidates, but POC voters randomly vary the order of the POC candidates.
- **Scenario C.** All voters randomly order the White and the POC candidates on their ballots.
- **Scenario D.** All the POC voters agree on the rank order of the POC candidates and the White candidates, but White voters randomly vary the order of both groups of candidates.

Each of these scenarios can be roughly mapped onto a choice of α settings. To facilitate comparisons in the tables below, we will use the following correspondence.

$$A \leftrightarrow (1/2, 1/2, 1/2, 1/2) \quad B \leftrightarrow (2, 1/2, 1/2, 1/2) \quad C \leftrightarrow (2, 2, 2, 2) \quad D \leftrightarrow (1/2, 1/2, 2, 2).$$

4.5 Cambridge sampler (CS)

A natural idea for trying to model the kind of ballots submitted by voters is to use existing data, and the richest supply of ranked ballots from American elections is from the 9-member city council of Cambridge, MA. The Cambridge Sampler (CS) model uses the group-level ballot types from ten years' worth of Cambridge city council elections to generate ballot types in a simulated election.




Figure 3. The top 30 most frequent ballot types in Cambridge elections in the five elections from 2009 to 2017 inclusive. Here, C denotes POC-identified candidates and W denotes White-identified candidates. The CS model samples from this distribution, conditioned on the voter's first choice.

The model generates ballots as follows. Using the support for each class of candidate from each racial group, voters are broken down into four groups: POC voters placing a C candidate first, POC voters placing a W candidate first, White voters placing a C candidate first and White voters placing a W candidate first. For each voter, based on their first vote, a ballot type is selected using the distribution of ballot types in the Cambridge elections. For example, if 5% of Cambridge ballots which ranked a C candidate first were of the form CWC , then a voter whose first choice is a C candidate has a 5% chance of filing a CWC ballot. See Figure 3 for the top 30 most frequent ballot types in Cambridge elections.

Once all the ballot types have been allocated, voters fill them in with candidates from the appropriate group based on a scenario chosen from Scenarios A–D in the previous section. If a ballot calls for more candidates from a particular group than are available (for example, a ballot type of $WWWWC$ when there are only 3 White candidates running), then the ballot is truncated when it can no longer be continued. One feature of the Cambridge sampler is that it generates incomplete

ballots, i.e., with not all candidates ranked, which is a common feature of real RCV elections (as can be seen in Figure 3). Indeed the single most frequently observed ballot type is a *bullet vote* for one White candidate with no other choices marked.

4.6 Model comparison

To investigate how the models behave in terms of representation, we compare model outputs in a number of hypothetical elections with various demographics and polarization parameters. As an additional point of reference, we will add a pared-down linear proportionality heuristic given by

$$\text{Seat share for POC-preferred} = (\pi_C)(\text{POC CVAP}) + (1 - \pi_W)(\text{White CVAP}).$$

This heuristic is built by first supposing that each group turns out to vote at the same rate as the other, so that the group’s citizen voting age population share is the same as their share of the electorate. Under that assumption, the right-hand side records the share of the vote for POC-preferred candidates, by computing the POC level of support times their share of the voters and adding White crossover support times their share of the voters. Seat share is set equal to vote share to set up the heuristic for what linear proportionality would predict.

We present a visualization in Figure 4 to compare the four models against the proportionality heuristic. To generate the figure, we consider a hypothetical election with six open seats, six POC-preferred candidates and six White-preferred candidates. Comparing all models and all parameter settings is prohibitive, so we focus on one candidate strength setting for comparisons. In the case of the PL and BT models, we consider only $\alpha = (1, 1, 1, 1)$ in this figure. For the AC and CS models, we take the average of the four scenarios A through D. This corresponds to the right-most column in our case studies (e.g., Table 3). The POC share of electorate and the polarization parameters are varied in the figure.⁹

We see that despite the major differences in construction, the models have similar average behavior. One notable difference (observed in the top row of each square) is that the CS model is less likely to predict a total lockout of POC-preferred candidates than the others, especially when the POC percentage in the electorate is low (10% and 20%). This is likely driven by one distinguishing feature of the CS model: its use of incomplete ballots. To lock out the minority group in this setup, it would be necessary for significant numbers of voters to rank *all* available W candidates.¹⁰

Summarizing the findings from this comparison of model averages in Figure 4, we can see that all four models are reasonably in line with the linear heuristic, with PL being the closest match. All four models thus comport with a central proportionality trend in STV elections. This suggests that the linear heuristic may be useful as a crude way to ballpark STV outcomes without full simulations.

⁹In practice, we often observe relatively high levels of polarization, such as $\pi_C \approx .7$ and $\pi_W \approx .8$ measured in Yakima [19].

¹⁰It is also crucial to note that the same incomplete ballots can cause extreme behavior for the Cambridge sampler in certain candidate strength scenarios, but this is not visible in this figure because it averages over all candidate strength scenarios. That tendency of the CS model will be apparent in the case studies below.




Figure 4. Comparing the four models (PL/BT/AC/CS) to the linear proportionality heuristic on a hypothetical six-seat contest with six POC-preferred candidates and six White-preferred candidates. The columns give four different choices for the POC share of the electorate. The x -axis of each colored square indicates the cohesion of POC voters for preferred candidates (π_C), and the y -axis is the similar cohesion parameter for White voters (π_W). The small units are color-coded by the average number of POC-preferred candidates elected over 100 simulations with $N = 1000$ voters. For PL and BT we set $\alpha = (1, 1, 1, 1)$, and AC and CS use the average of the four candidate strength scenarios.

5 Comparing alternative electoral systems

We outline our methodology for analyzing SMD and STV election systems in elections with a measured level of racial polarization.

Step 0: Choosing a POC group.

Depending on the situation, we may want to study representation for a specific POC population group, or for a coalition of groups like the Latino/Asian plaintiffs in Lowell [17]. The choice depends on a number of factors including the demographics of the jurisdiction (larger minority groups are more likely to secure representation) and cohesion among minority voters of different racial groups (this can often be gauged from the RPV analysis in Step 2). It is also sometimes best to run the analysis several ways before fixing a choice of group.

Below, we will use citizen voting age population (CVAP) from the most recent American Community Survey (ACS), available as five-year rolling averages published by the Census Bureau. This gives us an estimate of the population that is eligible to vote, not taking incarceration statistics into account.

Step 1: SMD analysis.

Once we have selected a minority group as the relevant POC population for a jurisdiction, we generate a sample of districting plans with the appropriate number of districts for the governing body in question (drawn from an actual or proposed number of districted seats), using a Markov chain algorithm called Recombination or ReCom [8, 9] to propose iterative alterations to form a chain of possible plans. We use an acceptance function to decide whether to accept each proposed change, tilted to target plans with higher numbers of majority-minority districts.¹¹ We then compute the demographic statistics across the sample in a box plot (such as in Figure 7). The districts are arranged across the x axis, from lowest to highest POC share in each plan. We also display the highest-scoring plan that was encountered in the search as a demonstration plan.

Even though having a majority of CVAP is neither a requirement nor a guarantee of electoral success, this gives us an idea of how many districts might be effective for the POC group under the most favorable conditions.

The reason to use an optimization technique to try to find an especially favorable districting plan, rather than comparing to a truly neutral ensemble of alternatives, is to model the choice that is likely to be faced by reform advocates: whether to use RCV or to seek a remedy with a high

¹¹To accomplish this, we preferentially accept plans which score high by the following formula.

$$\text{score of plan } P = F(P) = \sum_{i=1}^m D(i)$$

where m is the number of districts in plan P and D is a piecewise linear function given by

$$D(i) = \begin{cases} 1 & \text{CVAP}_i \geq .5 \\ (\text{CVAP}_i - .35)/.15 & .35 \leq \text{CVAP}_i \leq .5 \\ 0 & \text{CVAP}_i \leq .35 \end{cases}$$

where CVAP_i is the POC share of CVAP in district i .

number of majority-minority districts. For instance, a recent VRA lawsuit in Lowell, Massachusetts ended in a choice to voters between a citywide RCV system and a districted system in which two of eight districts were mandated to be at least 50% Hispanic + Asian by CVAP—the maximum number of such districts that the court believed possible.¹²

Step 2: RPV analysis.

In order to apply our STV models, we need to determine the polarization parameters to use as input. We will outline one proposed method for determining these values based on historical data, but we emphasize that users are free to choose whatever method they like to obtain these values. For example, these parameters could be set based on survey data or exit polls, or could be chosen in an exploratory fashion.

In each case study below, we select one recent single-winner election as our test case for inferring what percentage of minority and majority voters supported each candidate. We attempted to select an election that would be considered highly probative in a VRA challenge, using the following criteria.

- The contests must have each voter selecting a single candidate. To see why, consider a situation where voters vote for five candidates each, unranked. A universal first choice candidate among voters receives just as many votes as a universal fifth choice candidate, making them indistinguishable.
- The contests should have an available candidate of choice for the relevant minority group. This will often but by no means always be a member of the same group.
- The contests should be as closely connected to the multi-winner election being analyzed as possible so as to capture the specific voting patterns of the region being studied. In particular, the same kind of election as the one being analyzed (endogenous) is preferred to statewide races (exogenous) when possible. Recent elections are also preferable to older ones.

A fuller analysis would be based on a larger collection of elections, as in [16, 19].

The secret ballot means that we do not have any ground truth as to how many voters in each group voted for a specific candidate. We therefore employ a statistical inference method commonly used for RPV in research and litigation: a version of King’s Ecological Inference (EI) [15] implemented in the R package `ei`.

Step 3: STV analysis.

To apply each of the models (PL/BT/AC/CS), we use support levels from RPV for the polarization parameters π_C, π_W and report outputs from each of several different candidate strength settings. We run 100 simulated elections with $N = 1000$ ballots cast in each and compute the average number of candidates elected from the relevant racial group.¹³

¹²Voters were actually asked to indicate their approval for each system separately. Districts got around 60% support and RCV received just under 50% support. A special master is now drawing the city’s first-ever city council districts.

¹³While many elections have more than 1000 ballots cast, we choose this value for computational efficiency and since it allows for enough randomness to generate a variety of possible outcomes.

As one final variation, we consider one candidate pool with equal numbers of POC-preferred and White-preferred candidates, and another where there are fewer POC candidates available.

For an example of how this data is compiled into a table, see Table 3.

6 Case studies

6.1 Terrebonne Parish, LA

6.1.1 Background




Figure 5. Terrebonne Parish demographics. (Note that the southern part of the map is composed mainly of water blocks in the Gulf of Mexico.)

Terrebonne Parish is the coverage region for the 32nd District Court in Louisiana. In 2017, the NAACP Legal Defense Fund (often abbreviated LDF) challenged the election system used to elect five judges to this district court. LDF won the case,¹⁴ but the decision was overturned on appeal in 2020.¹⁵ Under the current election system, the 32nd district is divided into five divisions for the sole purposes of nomination and election of judges. To run for election, a candidate must designate which division they are running in, and to be elected they must attain a majority of the votes in that division, either in a primary or a subsequent two-way runoff election. All voters may vote in the primary and runoff elections, voting for up to one candidate in each division, regardless of their party affiliation or the affiliation of the candidates. In the 2017 ruling, the court found that Terrebonne Parish’s system of designated seats was even worse for Black representation than true at-large plurality (where each voter would vote for up to five candidates from the full pool of candidates). Even though Terrebonne Parish is about 18% Black by CVAP, a Black candidate had never been elected to the court, with one exception: Judge Juan Pickett, who ran unopposed for his division in 2014. In overturning the decision, the Fifth Circuit U.S. Court of Appeals pointed

¹⁴Terrebonne Parish NAACP v. Jindal, CIVIL ACTION NO. 14-069-JJB-SCR (M.D. La. Jul. 21, 2014)

¹⁵Terrebonne Parish Branch NAACP v. Jindal, 274 F. Supp. 3d 395 (M.D. La. 2017)

Houma City Court 2014		
	Black-preferred	other candidates
Black voters	0.85	0.15
non-Black voters	0.06	0.94

Houma City Marshal 2014		
	Black-preferred	other candidates
Black voters	0.82	0.18
non-Black voters	0.06	0.95

President 2012		
	Black-preferred	other candidates
Black voters	1.00	0.00
non-Black voters	0.13	0.87

Figure 6. Polarization estimates for three elections in Terrebonne Parish. For the STV analysis, we used the Houma City Court estimates shown in bold.

out that a system based on five (geographically defined) districts each electing one judge ran too much against the “state’s substantial interest in linking judicial positions to the judges’ parish-wide jurisdiction.” This naturally prompts the question of whether at-large STV would be an effective way to avoid districts while also structurally promoting Black representation on the Court. We therefore compare favorable districted options to STV to see if either is likely to improve Black representation.

For the STV analysis, we estimate the support for Black candidates using one of the elections used by Dr. Richard Engstrom in his expert testimony in the case: the 2014 Houma City Court race. One Black candidate (Cheryl Carter) ran in the race against two White candidates (Matthew Hagen and Randy Alfred). Estimates are shown in Figure 6¹⁶. For comparison, we also show estimates for two other elections, with numbers taken directly from Dr. Engstrom’s testimony.

The results of our STV simulations using these estimates are shown in Table 3. SMD analysis using BCVP is shown in Figure 7.

6.1.2 Findings

In principle, since Black residents make up 18% of CVAP, it might be barely possible to have two majority-Black districts out of five if the Black population were sufficiently concentrated.¹⁷ With the actual spatial arrangement of demographics, we find that even one district with the highest possible BCVP will still have only a narrow majority.¹⁸ This leaves some doubt as to whether any districting plan can provide effective opportunities to elect a Black-preferred candidate to the court. One promising fact is that the turnout rate among Black residents, at least for the 2014 contest considered in Table 6, was roughly equal to that of non-Black residents, according to the

¹⁶Estimates were obtained using 2×2 ecological inference (EI) which compares voting with demographics of each precinct. The demographic breakdown of each precinct was obtained from a voter registration file provided by the Parish. Louisiana has self-identified race as a field in the voter file.

¹⁷This would additionally need to leverage CVAP imbalance across districts.

¹⁸The plaintiffs in the 2017 case presented an demonstration plan with 53% Black CVAP using older ACS data.

voter registration file. In the STV analysis (Table 3), most scenarios elect 1 to 2 Black candidates. These estimates back up a claim that at-large STV would perform as well or better than a standard SMD system, while avoiding one of the Court’s main objections to a districted system as remedy.




Figure 7. On the left, district-level Black CVAP shares for a sample of 5-district plans for Terrebonne Parish. These plans were generated by an algorithm designed to search for majority-Black districts. The best-scoring plan encountered in the search is shown on the right.

5 seats 5 C / 5 W candidates	Scenario A	Scenario B	Scenario C	Scenario D	average
	(.5, .5, .5, .5)	(2, .5, .5, .5)	(2, 2, 2, 2)	(.5, .5, 2, 2)	(1, 1, 1, 1)
PL (Individual draws)	1.4	1.4	0.7	0.8	1.0
BT (Paired comparisons)	1.4	1.4	1.0	0.9	1.1
Alternating crossover	1.0	1.0	0.9	0.9	0.9
Cambridge sampler	2.0	2.0	0.0	0.4	1.1

5 seats 3 C / 5 W candidates	Scenario A	Scenario B	Scenario C	Scenario D	average
	(.5, .5, .5, .5)	(2, .5, .5, .5)	(2, 2, 2, 2)	(.5, .5, 2, 2)	(1, 1, 1, 1)
PL (Individual draws)	1.4	1.4	0.9	0.9	1.1
BT (Paired comparisons)	1.4	1.4	1.0	1.0	1.1
Alternating crossover	1.0	1.0	1.0	1.0	1.0
Cambridge sampler	2.0	2.0	0.0	1.0	1.2

Table 3. Expected number of Black-preferred candidates elected to $m = 5$ open seats in a ranked choice election in Terrebonne Parish. Expectation is computed by averaging over 100 simulations with $N = 1000$ voters under two different pools of available candidates. Each model either uses specified scenarios for candidate strength or uses the corresponding parameters $(\alpha_{CC}, \alpha_{CW}, \alpha_{WC}, \alpha_{WW})$ shown in the second row.

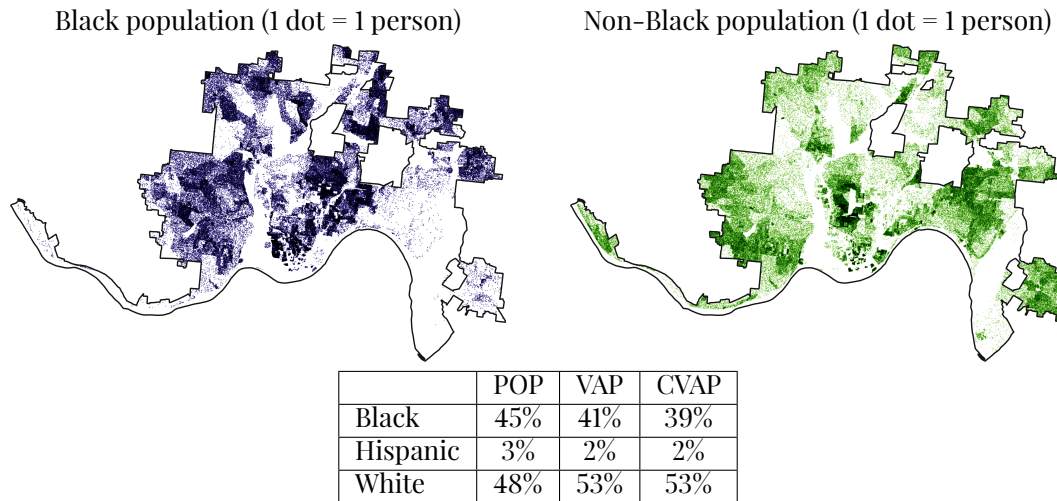


Figure 8. Cincinnati demographics.

6.2 Cincinnati, OH

6.2.1 Background

Ranked choice voting in Cincinnati has a long and complicated history. In 1925, Cincinnati moved to a ranked choice voting system for its city council after a campaign by the City Charter Committee largely aimed at dismantling the “Republican machine”. The City Charter Committee failed to endorse a single Black candidate in subsequent elections, though, and it would take until 1931 for Cincinnati’s first Black city council member to be elected (Frank A.B. Hall), on the Republican ticket. Cincinnati would abandon STV in 1957 in favor of an at-large plurality voting system. For more on this time period and for references for the above, see [3]. Today, Cincinnati still elects nine city council members via at-large plurality in a single election. It is interesting to model the effects of a possible return to STV in Cincinnati, especially in terms of Black representation.

We estimate current support for Black candidates using the 2017 Mayoral runoff with a 2×2 EI run based on CVAP. Yvette Simpson, a Black woman, lost the election to the incumbent John Cranely, a White man. Estimates are shown in Figure 9. We also show two other elections for comparison. In these other two elections, Black voters voted cohesively for their candidate of choice, while non-Black voters were split. Note that a Democrat and a Republican ran in these two elections (with the Democrat being the Black-preferred candidate in each case), while the 2017 Mayoral runoff was between two Democrats. For this reason, we think the 2017 Mayoral runoff is a better estimator of racial polarization since it is less affected by partisanship.

The results of our STV simulations using these estimates are shown in Table 4. Districting analysis is shown in Figure 10.

6.2.2 Findings

For the hypothetical nine-seat SMD system, Figure 10 shows a demonstration plan with seven majority-Black districts. A neutral algorithm which did not take demographics into account typ-

Mayor 2017		
	Black-preferred	other candidates
Black voters	0.76	0.24
non-Black voters	0.29	0.71

Cincinnati County Commissioner 2018		
	Black-preferred	other candidates
Black voters	0.98	0.02
non-Black voters	0.575	0.425

Governor 2018		
	Black-preferred	other candidates
Black voters	0.98	0.02
non-Black voters	0.63	0.37

Figure 9. Polarization estimates for three elections in Cincinnati. For the STV analysis, we used the Mayor 2017 estimates shown in bold.

ically generated three majority-Black districts, with a fourth not uncommon. Looking at Table 4, most STV estimates predict three to five seats out of nine won by Black candidates. Thus, if we are confident in the ability of majority-Black districts to elect Black representatives, then the ceiling on Black representation is much higher for a districted system than for STV. Of course, seven districts with only slightly more than 50% Black population can be risky given variable turnout and support for Black candidates. Despite its at-large voting system, Cincinnati elected four Black city council members in 2017. Among those elected were 6 Democrats (two also running as Charterites), one Independent and two Republicans (one a Charterite as well). This indicates that bloc voting, along both partisan and racial lines, is not so strong as to completely fence out minorities from representation. Even though it is not a likely site for VRA challenge at present, we can still consider the different prospects under different systems of election.

One parameter choice which deserves special attention is the Cambridge Sampler (CS) model with Scenario C. The expected number of Black-preferred candidates projected to win is 2.5 when nine Black-preferred candidates are running, and 5.0 when five are running. The reason for the poor performance in the former situation is that the Black vote is split among too many candidates, *and* voters rank only a few candidates on their ballots. This striking dependence on the size of the candidate pool is not present under the other models largely because those models assume that each voter submits a full ranking. This phenomenon – that performance under CS Scenario C depends heavily on the size of the candidate pool – also appears in the Jones County, Cincinnati and Pasadena case studies below. It indicates an important lesson: if there are many strong POC-preferred candidates running in an election, it is crucial that POC voters rank as many of them as possible in order to elect a proportional number.




Figure 10. On the left, district-level Black CVAP shares for a sample of 9-district plans for Cincinnati. These plans were generated by an algorithm designed to search for majority-Black districts. The best-scoring plan encountered in the search is shown on the right.

9 seats 9 C / 9 W candidates	Scenario A	Scenario B	Scenario C	Scenario D	average
		(.5, .5, .5, .5)	(2, .5, .5, .5)	(2, 2, 2, 2)	(.5, .5, 2, 2)
PL (Individual draws)	4.5	4.8	4.3	3.9	4.4
BT (Paired comparisons)	4.5	4.7	4.1	3.7	4.4
Alternating crossover	4.0	5.0	4.0	3.0	4.0
Cambridge sampler	3.9	5.0	2.5	2.1	3.4

9 seats 5 C / 9 W candidates	Scenario A	Scenario B	Scenario C	Scenario D	average
		(.5, .5, .5, .5)	(2, .5, .5, .5)	(2, 2, 2, 2)	(.5, .5, 2, 2)
PL (Individual draws)	3.9	4.3	4.3	3.7	4.1
BT (Paired comparisons)	3.9	4.3	4.3	3.6	4.1
Alternating crossover	4.0	5.0	4.6	3.5	4.3
Cambridge sampler	3.8	5.0	5.0	3.3	4.3

Table 4. Expected number of Black-preferred candidates elected to $m = 9$ open seats in a ranked choice election in Cincinnati. Expectation is computed by averaging over 100 simulations with $N = 1000$ voters under two different pools of available candidates. Each model either uses specified scenarios for candidate strength or uses the corresponding parameters $(\alpha_{CC}, \alpha_{CW}, \alpha_{WC}, \alpha_{WW})$ shown in the second row.

6.3 Jones County, NC

6.3.1 Background

In 2017, the Jones County Board entered into a consent decree to change the voting system used to elect County Commissioners.¹⁹ Prior to this decree, five Commissioners were elected every four years using an at-large partisan system consisting of a primary and general election. In each

¹⁹Hall et al v. Jones County Board of Commissioners et al

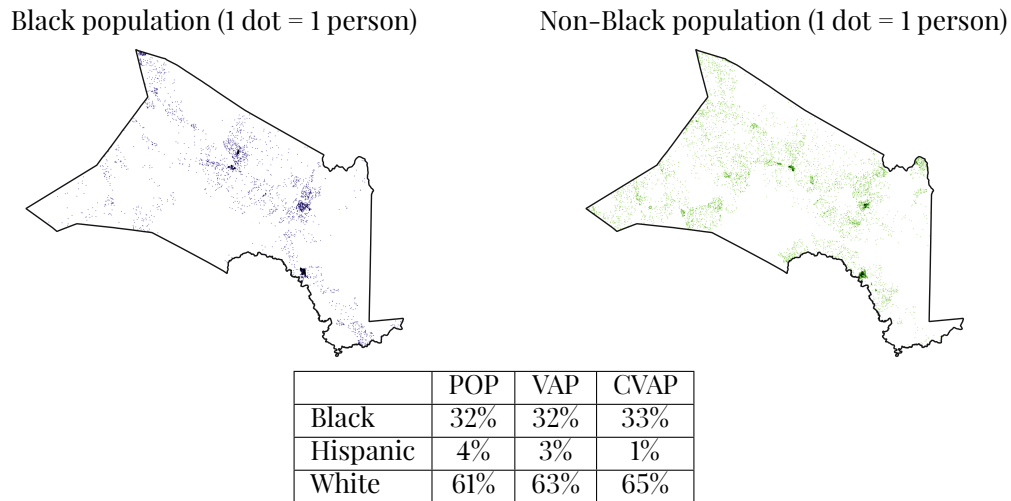


Figure 11. Jones County demographics.

primary, voters could vote for up to five of the candidates running for that party's nomination. The top five vote recipients from each party then advanced to the general, where voters voted for up to five candidates, with the top five vote recipients being elected to the Board. Despite the county being roughly 32% Black, no Black candidate had ever been elected to the commission. A districting plan including two majority-Black districts (by voting age population) out of seven was put into place in time for the 2018 election. As many expected, two Black candidates, Charlie Dunn and James Harper, were elected.

A single-winner local election in Jones County in which a Black candidate ran was not readily available, so we instead use a statewide election: the 2016 North Carolina Gubernatorial Election. This election was partisan, with the Democratic candidate Roy Cooper narrowly defeating the Republican Pat McCrory. In Jones County, there was a large degree of polarization between Black and non-Black voters in this race, as can be seen in Figure 12.²⁰ In general, the use of Party-ID elections to infer voting preferences can subsume racial polarization in broader partisan patterns. However, the Jones County Board of Commissioner elections also use Party-ID, with most candidates affiliated with the two major parties, which gives us some reason to believe that the gubernatorial election will exhibit relevant behavior.²¹ If anything, support for Cooper may overestimate White willingness to support a local Black-preferred candidate; therefore, the White support in local elections may be lower than the 13% in Table 12. Figure 12 also shows two other statewide elections for comparison, both Party-ID.

Our STV analysis is carried out for two different commission sizes: five seats (the size before 2017) and seven seats (the size after 2017). The results are in Table 5. District-based analysis, again for two commission sizes, is shown in Figure 13.

<u>Governor 2016</u>		
	Black-preferred	other candidates
Black voters	0.98	0.02
non-Black voters	0.13	0.87

<u>President 2016</u>		
	Black-preferred	other candidates
Black voters	0.98	0.02
non-Black voters	0.14	0.86

<u>President 2012</u>		
	Black-preferred	other candidates
Black voters	0.97	0.03
non-Black voters	0.21	0.79

Figure 12. Polarization estimates for three elections in Jones County. For the STV analysis, we used the Governor 2016 estimates shown in bold.

6.3.2 Findings

The district-based analysis in Figure 13 shows that one majority-minority district (by CVAP) is the most that can be expected for both the five- and seven-seat options.²² The Black population in Jones County is simply too spread out to easily make any more majority-Black districts (see Figure 11). The STV analysis in Table 5 shows that two Black candidates are likely to be elected out of five seats, and two to three out of seven. These estimates indicate that STV would likely secure a proportional number of seats for Black-preferred candidates, and outperform a district system.

²⁰Estimates were obtained using 2×2 ecological inference (EI) with VAP data. An important caveat of this data is that early and absentee votes in North Carolina are reported at the county level and thus were excluded from this analysis.

²¹In particular, both Black commissioners currently serving ran as Democrats.

²²If we use VAP instead of CVAP, then a second narrowly majority-Black district can be found. In fact, in the court case, the parties agreed to a seven-district plan with two majority-Black districts (53% and 57%) by VAP.

5 seats 5 C / 5 W candidates	Scenario A	Scenario B	Scenario C	Scenario D	average
	(.5, .5, .5, .5)	(2, .5, .5, .5)	(2, 2, 2, 2)	(.5, .5, 2, 2)	(1, 1, 1, 1)
PL (Individual draws)	2.2	2.2	2.0	1.9	2.1
BT (Paired comparisons)	2.0	2.1	1.4	1.3	1.6
Alternating crossover	2.0	2.0	2.0	2.0	2.0
Cambridge sampler	2.0	2.0	0.3	1.0	1.3
5 seats 3 C / 5 W candidates	Scenario A	Scenario B	Scenario C	Scenario D	average
	(.5, .5, .5, .5)	(2, .5, .5, .5)	(2, 2, 2, 2)	(.5, .5, 2, 2)	(1, 1, 1, 1)
PL (Individual draws)	2.1	2.2	2.0	1.9	2.0
BT (Paired comparisons)	2.1	2.1	1.7	1.6	1.9
Alternating crossover	2.0	2.0	2.0	2.0	2.0
Cambridge sampler	2.0	2.0	2.5	1.5	2.0
7 seats 7 C / 7 W candidates	Scenario A	Scenario B	Scenario C	Scenario D	average
	(.5, .5, .5, .5)	(2, .5, .5, .5)	(2, 2, 2, 2)	(.5, .5, 2, 2)	(1, 1, 1, 1)
PL (Individual draws)	3.1	3.3	2.7	2.5	2.84
BT (Paired comparisons)	3.0	2.9	2.1	2.0	2.4
Alternating crossover	2.9	3.0	2.0	2.0	2.5
Cambridge sampler	3.0	3.0	0.5	2.0	2.1
7 seats 4 C / 7 W candidates	Scenario A	Scenario B	Scenario C	Scenario D	average
	(.5, .5, .5, .5)	(2, .5, .5, .5)	(2, 2, 2, 2)	(.5, .5, 2, 2)	(1, 1, 1, 1)
PL (Individual draws)	2.9	3.1	2.9	2.6	2.95
BT (Paired comparisons)	2.8	2.9	2.2	2.1	2.6
Alternating crossover	2.9	3.0	3.0	2.0	2.7
Cambridge sampler	3.0	3.5	3.7	2.0	3.0

Table 5. Expected number of Black-preferred candidates elected to $m = 5$ or $m = 7$ open seats in a ranked choice election in Jones County. Expectation is computed by averaging over 100 simulations with $N = 1000$ voters under two different pools of available candidates. Each model either uses specified scenarios for candidate strength or uses the corresponding parameters (α_{CC} , α_{CW} , α_{WC} , α_{WW}) shown in the second row.




Figure 13. On the left, district-level Black CVAP shares for a sample of 5-district plans (top) and 7-district plans (bottom) for Jones County. These plans were generated by an algorithm designed to search for majority-Black districts. The best-scoring plan encountered in the search is shown on the right.

6.4 Pasadena, TX

6.4.1 Background

In 2014, in the wake of *Shelby County v Holder* and the removal by the Supreme Court of the pre-clearance requirement under Section 5 of the Voting Rights Act, Pasadena changed its city council election system from eight districts (8-0) to six districts and two at-large seats (6-2). Under the 6-2 system, the two at-large seats are elected using numbered places (Place G and Place H), and so are contested separately. In both systems, for both the at-large numbered places and the district seats, a majority of the vote is required to be elected. If no candidate receives a majority of the votes for a given seat, a runoff is held between the top two vote recipients to decide a winner. The city was sued following the change to a 6-2 plan by plaintiffs represented by the Mexican American Legal Defense and Educational Fund (MALDEF) for alleged dilution of Hispanic voting power.²³ Under recent ACS data, Pasadena is about 53% Hispanic by CVAP. MALDEF argued that the lower turnout

²³Patino v. City of Pasadena, 229 F. Supp. 3d 582 (S.D. Tex. 2017)




Figure 14. Pasadena demographics.

rate among Hispanics placed the group at a disadvantage when electing at-large candidates. The federal court ruled that Pasadena must return to the original eight-district system and moreover, due to evidence of intentional discrimination, that the city would be “bailed in” to preclearance under Section 3 of the Voting Rights Act. An appeal was launched by the city, but dropped after a settlement was reached in 2017.

One city council election was run under the 6-2 system, and one of the two at-large seats was contested by a Hispanic candidate, Oscar Del Toro, who lost with 39.1% of the vote. This election (the at-large race for Place H) was also presented at trial and analyzed by Dr. Richard Engstrom. The ecological inference results taken from Engstrom’s analysis for the plaintiffs in the case are shown in Figure 15. Two other elections are also shown for comparison. The Place G City Council race in 2015 was between two non-Hispanic candidates, which is why we chose to use the Place H estimates for our STV analyses.

The results of our STV simulations using these estimates are shown in Table 6. These simulations assume equal turnout between Hispanic and non-Hispanic residents. In the case of Pasadena, the relatively low turnout of Hispanic residents was highlighted during the court case as particularly important in determining the election of Hispanic candidates. For this reason, we also report simulations where the relative Hispanic turnout is lower than the non-Hispanic turnout. In particular, the estimates reported in Engstrom’s analysis imply that only 18% of the voters in the at-large Place H election were Hispanic. The results for the low Hispanic turnout scenario are shown in Table 6. Districting analysis (without adjusting for turnout) is contained in Figure 16.

6.4.2 Findings

The district-based analysis in Figure 16 shows that prospects for majority-Hispanic districts, even by CVAP, are very good in Pasadena, with as many as 7 being theoretically possible even with three above 70%. This optimistic result is tempered by the issue of low Hispanic turnout, which could prevent even a 75% Hispanic district from electing a Hispanic-preferred candidate, and presents

	Black-preferred	other candidates
Black voters	0.87	0.13
non-Black voters	0.28	0.72

	Black-preferred	other candidates
Black voters	0.71	0.29
non-Black voters	0.42	0.58

	Black-preferred	other candidates
Black voters	0.96	0.04
non-Black voters	0.29	0.71

Figure 15. Polarization estimates for three elections in Pasadena. For the STV analysis, we used the Place H City Council 2015 estimates shown in bold.

a serious risk for districts closer to 50%.²⁴ For some context on the effectiveness of majority-Hispanic districts, the 6-2 map used in 2015 contained three majority-Hispanic districts. A Hispanic candidate was elected from two of these, in the other the Hispanic candidate narrowly lost. A third Hispanic candidate, Cody Ray Wheeler, was elected from a 47% Hispanic CVAP district.²⁵

The STV estimates in Table 6 indicate that under equal turnout, roughly 4 Hispanic-preferred candidates are likely to be elected. Notably, under lower turnout this number drops only by about one seat. This is because the turnout effect is mitigated by non-Hispanic support for Hispanic-preferred candidates (estimated at 28%). Thus, the ceiling for representation via districts, wherein every majority-Hispanic district elects a Hispanic-preferred candidate, is much higher than the estimated representation under STV. This relies heavily on the ability of candidates to drive up Hispanic turnout within their own districts, however. On the other hand, STV allows Hispanic-preferred candidates to also leverage non-Hispanic support from high-turnout areas of the city. Note that this requires wider campaigning, which has proven challenging for Hispanic candidates in Pasadena.²⁶

²⁴The literature is unclear what the general effect of majority-minority districts are on minority turnout, though at least one study suggests a positive effect [2]. Hispanic turnout decreased after the introduction of the 6-2 map, which shows that the change affected participation among Hispanics (who widely opposed the change).

²⁵Wheeler testified in court that it would have been harder for him to be elected if he had a Spanish surname, and that he did not advertise his Latino ethnicity.

²⁶Del Toro and another Hispanic candidate, Ybarra, testified that they were not well received in South Pasadena and were the target of racist remarks.

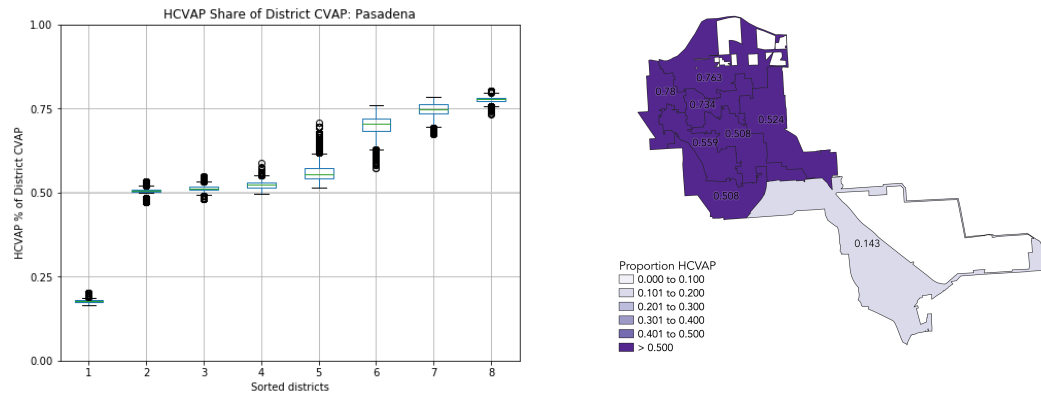


Figure 16. On the left, district-level Hispanic CVAP shares for a sample of 8-district plans for Pasadena. These plans were generated by an algorithm designed to search for majority-Hispanic districts. The best-scoring plan encountered in the search is shown on the right.

8 seats	Scenario A	Scenario B	Scenario C	Scenario D	average
8 C / 8 W candidates	(.5, .5, .5, .5)	(2, .5, .5, .5)	(2, 2, 2, 2)	(.5, .5, 2, 2)	(1, 1, 1, 1)
PL (Individual draws)	4.5	5.2	5.0	4.2	4.9
BT (Paired comparisons)	4.6	5.0	4.6	4.1	4.6
Alternating crossover	4.0	5.0	5.0	4.0	4.5
Cambridge sampler	3.1	5.0	5.8	3.0	4.2

8 seats	Scenario A	Scenario B	Scenario C	Scenario D	average
4 C / 8 W candidates	(.5, .5, .5, .5)	(2, .5, .5, .5)	(2, 2, 2, 2)	(.5, .5, 2, 2)	(1, 1, 1, 1)
PL (Individual draws)	3.5	3.9	3.9	3.5	3.8
BT (Paired comparisons)	3.5	4.0	4.0	3.4	3.8
Alternating crossover	4.0	4.0	4.0	4.0	4.0
Cambridge sampler	3.0	4.0	4.0	3.4	3.6

8 seats – low turnout	Scenario A	Scenario B	Scenario C	Scenario D	average
8 C / 8 W candidates	(.5, .5, .5, .5)	(2, .5, .5, .5)	(2, 2, 2, 2)	(.5, .5, 2, 2)	(1, 1, 1, 1)
PL (Individual draws)	3.7	3.7	3.0	2.9	3.3
BT (Paired comparisons)	3.5	3.6	2.7	2.7	3.1
Alternating crossover	3.0	3.0	2.0	1.1	2.3
Cambridge sampler	3.0	3.7	0.4	1.1	2.0

8 seats – low turnout	Scenario A	Scenario B	Scenario C	Scenario D	average
4 C / 8 W candidates	(.5, .5, .5, .5)	(2, .5, .5, .5)	(2, 2, 2, 2)	(.5, .5, 2, 2)	(1, 1, 1, 1)
PL (Individual draws)	3.0	3.2	3.2	3.0	3.1
BT (Paired comparisons)	3.0	3.1	2.9	2.9	3.0
Alternating crossover	3.0	3.0	3.0	2.0	2.8
Cambridge sampler	3.0	3.8	4.0	2.8	3.4

Table 6. Expected number of Latino-preferred candidates elected to $m = 8$ open seats in a ranked choice election in Pasadena. Expectation is computed by averaging over 100 simulations with $N = 1000$ voters under two different pools of available candidates. Each model either uses specified scenarios for candidate strength or uses the corresponding parameters $(\alpha_{CC}, \alpha_{CW}, \alpha_{WC}, \alpha_{WW})$ shown in the second row. Under the low turnout assumption, 18% of the electorate is Hispanic; this estimate is based on a recent at-large city council election.

7 Future work

We briefly outline some ways in which the work in this report could be expanded or continued. A key factor that has not been sufficiently explored here is that of *incomplete ballots*. With the exception of the CS model, our models generate full rankings by each voter of all the available candidates. As Figure 3 makes clear, however, incomplete ballots—and even votes for only one candidate, or *bullet votes*—are common in real elections. Future work should incorporate ballot truncation into all of the models, either by relying on the empirically observed ballot lengths in election data, or through the use of an additional parameter or mechanism.

While the table of outputs for each case study (e.g., Table 6) captures a significant variety of scenarios that could play out in an STV election, it can be hard to scan for patterns from a table alone. In future work, we plan to design a hyperparameter search, i.e., a method for searching the parameter space to find settings that produce extreme outcomes, both in terms of expected value and in terms of variance. With a method that finds corner cases, we can reveal the threats and opportunities for ranked choice voting more systematically and clearly.

Finally, as more jurisdictions adopt ranked choice voting, we look forward to more (and more varied) ranked ballot data becoming available in the near future. A key obstacle to settling on a single method for simulating STV elections is the dearth of ranked ballot images for multi-winner races in the United States. If model specification is over-reliant on data from a single source (like Cambridge city council elections) or with particular structure (like IRV where rankings only go three deep), it can lead to over-fitting. Local conditions or other common structural elements are unlikely to generalize. For this reason, we present a range of models and scenarios instead of picking a preferred one. Three of the four models are built without the use of any training data, relying instead on established statistical methods in the case of PL and BT. Richer training data will make it more feasible to fit a model and parameters that comport with observed ranking behavior.

8 Conclusion

In this report, we consider the potential for STV to better protect voting rights for historically disenfranchised communities than the more standard SMD approach. To do this, we must model the effect of STV on minority representation in a theoretically sound way. We have outlined a methodology for comparing alternative electoral systems which builds on the foundation of established legal practice in VRA cases. The early lessons from applying this methodology in four case studies indicate that while representation under district systems can vary widely, RCV consistently provides proportional (or slightly better) representation for minority groups.

References

- [1] Carolyn Abott and Asya Magazinnik. At-large elections and minority representation in local government. *American Journal of Political Science*, 2020.
- [2] Matt A Barreto, Gary M Segura, and Nathan D Woods. The mobilizing effect of majority-minority districts on latino turnout. *American Political Science Review*, pages 65–75, 2004.

-
- [3] Robert A Burnham. Reform, politics, and race in Cincinnati: Proportional representation and the City Charter Committee, 1924–1959. *Journal of Urban History*, 23(2):131–163, 1997.
- [4] Alistair Clark. Party organization and concurrent multi-level local campaigning: The 2007 Scottish elections under MMP and STV. *Party Politics*, 18(4):603–622, 2012.
- [5] Alistair Clark and Lynn Bennie. Strategic actors or muddling through? Political parties, local campaigning and the single transferable vote. *Representation*, 44(4):327–340, 2008.
- [6] Loren Collingwood and Sean Long. Can states promote minority representation? assessing the effects of the California Voting Rights Act. *Urban Affairs Review*, 2019.
- [7] Douglas E Critchlow, Michael A Fligner, and Joseph S Verducci. Probability models on rankings. *Journal of mathematical psychology*, 35(3):294–318, 1991.
- [8] Daryl DeFord and Moon Duchin. Redistricting reform in Virginia: Districting criteria in context. *Virginia Policy Review*, 2019.
- [9] Daryl DeFord, Moon Duchin, and Justin Solomon. Recombination: A family of Markov chains for redistricting. *arXiv preprint arXiv:1911.05725*, 2019.
- [10] Moon Duchin, Taissa Gladkova, Eugene Henninger-Voss, Ben Klingensmith, Heather Newman, and Hannah Wheelen. Locating the representational baseline: Republicans in Massachusetts. *Election Law Journal*, 18(4):388–401, 2019.
- [11] Richard L Engstrom and Michael D McDonald. The underrepresentation of blacks on city councils: Comparing the structural and socioeconomic explanations for South/non-South differences. *The Journal of Politics*, 44(4):1088–1099, 1982.
- [12] Isobel Claire Gormley and Thomas Brendan Murphy. Exploring voting blocs within the Irish electorate: A mixture modeling approach. *Journal of the American Statistical Association*, 103(483):1014–1027, 2008.
- [13] Shaun Bowler Bernard Grofman et al. *Elections in Australia, Ireland, and Malta under the Single Transferable Vote: Reflections on an embedded institution*. University of Michigan Press, 2000.
- [14] Albert Karnig and Susan Welch. *Black representation and urban policy*. University of Chicago Press, 1980.
- [15] Gary King. *A solution to the ecological inference problem: Reconstructing individual behavior from aggregate data*. Princeton University Press, 2013.
- [16] MGGG Redistricting Lab. Study of voting systems for Santa Clara, CA. *White paper*, 2018. mggg.org/santaclara.
- [17] MGGG Redistricting Lab. Findings on the City of Lowell’s election systems. *White paper*, 2019. mggg.org/Lowell-Detailed-Report.
- [18] MGGG Redistricting Lab. Study of reform proposals for Chicago city council. *White paper*, 2019. mggg.org/chicago.
- [19] MGGG Redistricting Lab. Analysis of county commission elections in Yakima County. *White paper*, 2020. mggg.org/yakima.

- [20] Delbert Taebel. Minority representation on city councils: the impact of structure on blacks and Hispanics. *Social Science Quarterly*, 59(1):142–152, 1978.
- [21] Delbert A Taebel, Richard L Engstrom, and Richard L Cole. Alternative electoral systems as remedies for minority vote dilution. *Hamline J. Pub. L. & Pol'y*, 11:19, 1990.

A How does STV work?

Voting in a ranked choice election is fairly simple, but the mechanisms by which the winners are calculated are numerous and can be quite complicated. The main areas in which these mechanisms differ are the quota used for election and the method by which votes are transferred down the ballot after elimination or election of a candidate. In order to keep things simple, we restrict our attention in this report to the system used by Cambridge MA to elect its city council (the only city council currently elected by STV in the U.S.). The Cambridge system uses a Droop quota, the most common quota used in STV elections. For transferring votes, Cambridge uses a form of the random sample method. We now explain the Cambridge system in full.²⁷

Once ranked ballots have been collected, winners are determined in a series of “counts”. A *quota* is set to the smallest whole number greater than $N/(k+1)$, where N is the number of ballots and k the number of seats to fill.²⁸ In each count one of the following things happens:

- **A candidate’s first-place votes reaches or exceeds the quota.** In this case, that candidate is declared elected. If the candidate received M first-place votes, then $M - D$ “surplus” ballots are randomly selected to be *transferred* (see below). If more than one candidate satisfies this condition, then this process is applied to all of them.
- **No candidate has first-place votes total greater than or equal to the quota.** The candidate with the least votes is eliminated from contention and all their ballots are *transferred* (see below).
- **The number of candidates remaining equals the number of seats left to fill.** All remaining candidates are elected.

The *transfer* process consists of taking the ballots to be transferred, erasing their first choices and considering as their first choice the next ranked candidate not yet eliminated.

Some minor details of the Cambridge system which deserve mention, but which we do not implement in our code are as follows:

- The selection of surplus ballots is not truly random. The precincts are randomly ordered, then the ballots arranged by precinct in the order counted. To select the surplus ballots, every n^{th} ballot is selected, where n is the candidate’s total first-choice ballots divided by the number of ballots to be transferred (rounded off), except that ballots which do not have an uneliminated candidate to transfer to are not selected. If after one such pass too few surplus ballots have been selected, then the process is repeated with the remaining ballots until the right number have been selected.
- Any candidate with fewer than 50 first-place votes is automatically eliminated.

²⁷Source: https://www.cambridgema.gov/~link.aspx?_id=D58142E0FBC64BEDB4B0F92D957FD453&_z=z

²⁸This is the so-called Droop quota. The Hare quota is N/k .

B Notes on the data

B.1 Demographic data

In this report we cite numbers for total population (POP), voting age population (VAP), and citizen voting age population (CVAP). POP and VAP at the block level are taken from the 2010 Census, while CVAP is taken from the 2014–2018 ACS five year rolling average (which has considerably larger margins of error than the census data). To obtain block-level CVAP data for district analysis, we disaggregate the CVAP data by prorating based on VAP using the `maup` package (github.com/mggg/maup).

B.2 Election data

Precinct geography for the RPV analysis was compiled and cleaned by MGGG and is available for Jones County at github.com/mggg-states. See the references there for the original sources. Cleaned precinct geography for Terrebonne Parish and Cincinnati are available on GitHub²⁹. Terrebonne Parish election results were obtained from the Louisiana Secretary of State and precinct boundaries were obtained from the Louisiana House of Representatives. Cincinnati results and precinct boundaries were obtained from the Hamilton County Board of Elections. For Pasadena we relied on expert reports for RPV results.

B.3 Demographic data for RPV

By-precinct estimates of voter demographics are required for RPV analysis. When possible, we obtained a voter file to get precise numbers (for states where race is on the voter file). In some other cases, we used the proportion of VAP by precinct. This is more stable in the case of Black residents than Hispanic residents, since the latter tend to have lower citizenship rates. CVAP, while a better estimate than VAP, cannot be computed exactly for precincts since it is not available at the block level. In cases when a voter file was difficult to obtain and VAP was expected to be a bad estimator, we used numbers reported by expert witnesses in court.

C Proof of Theorem 1

Before proving Theorem 1, we establish the following lemma which handles the case where a solid coalition of size kt focuses its support on no more than k candidates.

Lemma 1. *Fix a quota t to be used for electing candidates, independent of the number of voters. Consider a solid coalition of at least kt voters supporting a set of candidates C' with $|C'| \leq k$. Then all candidates in C' are elected.*

Proof. We will prove this by induction on $|C'|$. When $|C'| = 0$ the statement holds trivially. When $|C'| = 1$, the candidate is ranked first by every voter in the coalition, meets the threshold for election, and is elected.

²⁹github.com/mggg/minority-RCV

Assume that the statement holds for all integers $|C'| = 1, \dots, k-1$. Let $|C'| = k$. All kt voters in the solid coalition fill their first k positions with candidates from C' . By the pigeonhole principle, there must be at least one candidate that appears in the first position on at least t ballots. This candidate is elected. This consumes at most t votes from the solid coalition, leaving a situation where a set of $k-1$ candidates are supported by a solid coalition of at least $(k-1) \cdot t$ votes. By the induction hypothesis we are done. \square

The only remaining obstacle to proving Theorem 1 is to show that when there are more candidates supported by the solid coalition than they are able to elect on their own, this does not dilute the support in such a way that fewer than α of them get elected. Recall that m denotes the magnitude of the STV election, i.e., the number of seats to be filled.

Proof of Theorem 1. Let C be the set of candidates. Consider a quota of election t and a solid coalition of at least $\alpha \cdot t$ voters supporting a set of candidates C' , with $|C'| \geq k$. If $|C'| = \alpha$ we are done by Lemma 1, so assume $|C'| > k$.

Assume for contradiction that STV concludes with a committee E , and $|E \cap C'| = k' < k$. Because of the size of the solid coalition, at most $m-k$ candidates from $C \setminus C'$ can be elected without the support of any voters forming part of the solid coalition. It follows that all the candidates in C must have been eliminated or elected at the time that the last candidate was elected.

At every step i the STV algorithm either elects a candidate to the committee, or eliminates a candidate with the smallest number of first-place votes. Keep track of two quantities during the execution of this algorithm: the number of candidates that can be elected by the solid coalition if they coordinate $s_i = v_i/t$, where v_i is the number of unused coalition votes after step i of the algorithm, and r_i , the number of candidates in C' who are neither elected nor eliminated after step i of the algorithm.

When a candidate from C' is elected in step i , $s_i \geq s_{i-1} - 1$ and $r_i = r_{i-1} - 1$. Electing a candidate in $C \setminus C'$ does not affect s or r while $r_i > 0$. Eliminating a candidate from C' in step i means $s_i = s_{i-1}$ and $r_i = r_{i-1} - 1$. Eliminating a candidate not in C' means $s_i = s_{i-1}$ and $r_i = r_{i-1}$.

Initially, $s_0 = k$ and $r_0 = |C'| > k$. By assumption the algorithm reaches a point where where $r_i = 0$ (and $s_i > 0$). Let j be the first time this occurs, in other words, $r_j = 0$ and $r_{j-1} = 1$. At some step $j^* \in \{1, \dots, j\}$ it holds for the first time that $s_{j^*} \geq r_{j^*}$. Up to step j^* , no votes from the solid coalition could have helped elect a candidate in $C \setminus C'$, so after step j^* at least $s_0 - s_{j^*}$ candidates from C' have been elected. Applying Lemma 1 at step j^* shows that another s_{j^*} candidates from C' will be elected. We can conclude that $|C' \cap E| \geq s_0 = k$, a contradiction. \square

We conclude by showing that vote management systems have limited impact in a two-party setting.

Proof of Corollary 1. We will establish a lower bound on the number of candidates elected by each coalition. For coalition A, Theorem 1 guarantees α preferred candidates are elected. Let N_B be the number of voters in coalition B and let $N = N_A + N_B$ be the total number of voters. Also, let

$t = N/m + 1 + \varepsilon$, so that $\varepsilon \leq 1$. Then

$$\begin{aligned} N_B = N - N_A &\geq N - (k+1)t + m + 1 = N - (k+1) \left(\frac{N}{m+1} + \varepsilon \right) + m + 1 \\ &= \left(\frac{N}{m+1} + \varepsilon \right) (m+1 - (k+1)) - (m+1)\varepsilon + m + 1 \geq (m-k)t, \end{aligned}$$

since $\varepsilon \leq 1$. It follows from Theorem 1 that coalition B elects $m - k$ preferred candidates. Since there are only m open seats, we get that coalition A and B elect k and $m - k$ candidates respectively, regardless of what vote management system they use. \square

D Details of PL and BT models

D.1 Generating support vectors for PL and BT

For both the PL and BT models, we are required to turn polarization parameters like these

	Candidates $\{C_i\}$	Candidates $\{W_i\}$
POC voters	π_C	$1 - \pi_C$
White voters	$1 - \pi_W$	π_W

into two support vectors,

$$\begin{aligned} \mathbf{v}^C &= \langle v_{C_1}^C, v_{C_2}^C, \dots, v_{W_1}^C, v_{W_2}^C, \dots \rangle \\ \mathbf{v}^W &= \langle v_{C_1}^W, v_{C_2}^W, \dots, v_{W_1}^W, v_{W_2}^W, \dots \rangle \end{aligned}$$

where \mathbf{v}^C and \mathbf{v}^W encode the support *from* POC and White voters respectively. Let's look at the entries which encode POC support for POC candidates of choice as an example. These are the entries $(v_{C_1}^C, v_{C_2}^C, \dots, v_{C_k}^C)$ where k is the number of POC-preferred candidates running. Since we know the total support we get the equation

$$\pi_C = \sum_i v_{C_i}^C$$

We therefore want to generate tuples of numbers $(v_{C_1}^C, v_{C_2}^C, \dots, v_{C_k}^C)$ which sum to π_C . To do this, we sample *weights* $\mathbf{w} = (w_1, w_1, \dots, w_k)$ which sum to 1 and then define $v_{C_i}^C = w_i \pi_C$ to get our vector values. The weights are drawn from a symmetric Dirichlet distribution. This distribution has support all k -tuples of positive real numbers which sum to 1, and probability density

$$P(w_1, w_2, \dots, w_k) = B \prod_i w_i^{\alpha-1}$$

where B is a normalizing constant. The α parameter is called a *concentration parameter* and is used to control how uniform the w_i tend to be when drawn from this distribution. If $\alpha \ll 1$, then the vector \mathbf{w} is more likely to be sparse, i.e. has most of its entries near zero. If $\alpha \gg 1$, then the w_i are

more likely to have similar values. Thus by adjusting α , we can control how the support for POC candidates of choice from POC voters tends to be distributed.

What we have just described is how to generate a support vector for POC voters and POC candidates of choice. We have to do this three more times (POC for White, White for POC and White for White). Each of these is controlled by its own α . Thus we need to choose four α values to generate support vectors to feed into the Plackett-Luce and Bradley-Terry statistical models.

D.2 Plackett-Luce model

We refer the reader to [7] as a standard reference for this and the following section. Suppose that a voter is filling out a ranked ballot one position at a time, starting with their first-choice candidate. The Plackett-Luce model posits that at each stage the voter's choice between the remaining candidates is unaffected by the order in which higher candidates have been ranked. For example, suppose a voter is twice as likely to place candidate A first than to place candidate B first. If that voter does not place either A or B first, so that they are both available, the voter under this model is also twice as likely to place candidate A second than to place candidate B second. Mathematically, this works out to mean that each voter under a Plackett-Luce model has behavior governed by a *support vector*, a list of numbers that encodes the voter's relative level of support for each candidate. Without loss of generality, we can rescale so that the values sum to one. These support values can be used to generate a ballot for each voter in a probabilistic way. Voters fill out the ballot position by position, starting with their first choice. The probability that a voter writes down the name of a candidate at some rank is equal to the support value for that candidate divided by the sum of the supports for all the candidates not yet ranked. For example, consider the following case of two voters and four candidates.

	Candidate A	Candidate B	Candidate C	Candidate D
Voter X	0.2	0.1	0.1	0.6
Voter Y	0.7	0.1	0.1	0.1

Here, Voter X has a 60% chance of voting for Candidate D first. If they do rank Candidate D first, then they have a subsequent 50% probability of putting Candidate A second, a 25% chance of putting Candidate C second and a 25% chance of putting Candidate D second. Continuing on until the fourth choice will generate a ballot for this voter.

The Plackett-Luce model has the nice property that for very large numbers of voters all voting according to the same support vector, each candidate gets a share of first-place votes roughly equal to their support value. While applications of statistical models to ranked choice voting in the literature are extremely limited, Plackett-Luce models have been used in at least one place to study Irish RCV elections [12].

D.3 Bradley-Terry model

The Bradley-Terry model falls into the Paired Comparison (sometimes called Pairwise Comparison) family of models, which model ranked choice via paired comparisons between choices. In other words, the probability of a particular ballot depends on the order in which each pair of candidates appears. Unlike the Plackett-Luce model, which generates ballots a single choice at a time,

these models assign a probability to a complete ballot as follows. For each pair of candidates A and B , let $P(A < B)$ be the probability that a voter prefers A to B in a head-to-head comparison. To calculate the probability of a ranking $C_1 < C_2 < \dots < C_k$, we multiply all the pairwise probabilities and normalize:

$$P(C_1 < C_2 < \dots < C_k) = K \cdot \prod_{C_i < C_j} P(C_i < C_j) \quad (1)$$

where K is a constant introduced to make sure all the probabilities for all the possible ballots sum to one. Sampling directly from this distribution over ballots is not feasible computationally, so instead we use Markov Chain Monte Carlo methods to generate ballots.

The Bradley-Terry model is a particular type of paired comparison model. Just like the Plackett-Luce model, the Bradley-Terry model takes as input a support vector for each voter. If the value for Candidate A in the support vector is v_A and the value for Candidate B is v_B , then the probability that the voter prefers A to B is set at

$$\frac{v_A}{v_A + v_B}$$

Thus, to calculate the probability of a ranking $C_1 < C_2 < \dots < C_k$, we multiply all the pairwise probabilities and normalize:

$$P(C_1 < C_2 < \dots < C_k) = K \cdot \prod_{C_i < C_j} \frac{C_i}{C_i + C_j} \quad (2)$$

where K is a constant introduced to make sure all the probabilities for all the possible ballots sum to one. For example, let's revisit the example we had in the previous section:

	Candidate A	Candidate B	Candidate C	Candidate D
Voter X	0.2	0.1	0.1	0.6
Voter Y	0.7	0.1	0.1	0.1

If we compute the probabilities using Equation 2 above, we see that Voter X has just a 0.2% chance of filing the ballot ABCD, which is not surprising given that they support Candidate D so highly. On the other hand, Voter Y has a 14.3% chance of filing the ballot ABCD. This may seem low, but remember there are 6 different ballots Voter Y can submit that begin with Candidate A, all of which are equally likely since they support B, C and D equally. So Voter Y has an 85% chance of voting for Candidate A first. Note that this is higher than their support value viewed as a percentage (70%). This illustrates that, unlike Plackett-Luce, the Bradley-Terry model does not generate first choices in proportion with the support vector.