


Modeling the Fair Representation Act



MGGG Redistricting Lab

Abstract

The Fair Representation Act (FRA) proposes a fundamental change to the traditional electoral system used to vote members into the United States House of Representatives. Using a combination of multi-member districts and ranked-choice voting, the FRA has been claimed by advocates to promote fair electoral outcomes. Here, we investigate the potential effects of the FRA by linking several statistical models to (a) generate large sets of valid multi-member districting plans and (b) estimate the likely effectiveness of a given plan in providing electoral opportunities to communities of color. We find robust evidence that the FRA promotes far more proportional outcomes for underrepresented racial and ethnic groups, without the need for race-conscious line-drawing.

Contents

1	Introduction	2
2	Methods	3
2.1	Building multi-member plans	3
2.2	Simulating elections	4
2.3	National execution and focus states	5
2.4	Limitations of approach	5
3	Detailed findings from focus states	6
3.1	Florida	7
3.2	Illinois	8
3.3	Maryland	10
3.4	Massachusetts	11
3.5	Texas	13
4	National comparison and conclusions	15

1 Introduction

Since 1967, the United States has required that voters across all states elect members to the House of Representatives from *single-member districts*, or *SMDs*.¹ Despite their frequent use as a litigation remedy for minority voters under the Voting Rights Act (VRA) of 1965, it is well understood that SMDs can also effectively dilute the voting power of racial or ethnic blocs. Especially with the future of the VRA in doubt, civil rights thinkers are increasingly looking to alternative voting systems.

Ranked-choice voting, or *RCV*, is a family of electoral systems in which voters rank candidates in order of preference instead of selecting only one candidate. Though the name "ranked choice" only refers to the voter's ballot, the term RCV typically also refers to a general algorithm which uses multiple rounds of tabulation to convert the ranked results to one (or several) winners. If one winner is to be selected, this is typically called instant runoff voting, or *IRV*; if multiple winners are desired, the same algorithm is now called single transferable vote, or *STV*. In brief, the process works by immediately electing any candidate who was ranked first by some minimum threshold of the electorate, eliminating the candidate with the fewest #1 rankings, and then transferring support from the elected or removed candidates to the next candidate listed on the ballot, so long as that candidate is still in contention.² Prior research by members of our lab and collaborators has shown that STV elections can systematically secure fair representation for groups who are typically locked out by SMDs [2]. The Fair Representation Act (FRA) would implement this electoral system for Congressional representation across the United States, adding the additional requirement that all districts be drawn to elect three, four, or five members of the House [1].

In this report, we investigate potential representational outcomes under the FRA, focusing on the potential for members of racial and ethnic minorities to elect candidates of choice. Using industry-leading randomized techniques for generating districting plans and ranked-choice ballots, we model RCV outcomes across the country. Below, we highlight the findings for POC representation under the FRA in five key states: Florida, Illinois, Maryland, Massachusetts, and Texas. The remaining 45 states are also modeled, and all results are summarized in §4, with supporting data in the supplement.

Note on terminology. Throughout this report, we discuss how the electoral system implemented by the FRA may change the representational landscape for people who have systemically been denied equitable political representation. Below, we broadly refer to "POC" (people of color) and "White" subgroups, where White refers to those whose census response lists them as non-Hispanic single-race White, and POC is the complement. It is important to remember that the models in this report do not predict how many representatives will be people of color themselves, but rather how many will be POC-preferred. The methods, models, and model assumptions are explained in the following section.

Contributors

Anthony Pizzimenti and Moon Duchin led the study design and the writing of this report, with support from Gabe Schoenbach. The material here draws on earlier work by many current and former members of the MGGL Redistricting Lab, including Duchin and Schoenbach, with major contributions from Amy Becker, Dara Gold, and Thomas Weighill. This study was supported by funding from FairVote, a nonpartisan advocacy group associated with ranked choice and other electoral reforms.

¹When a free-standing law on congressional districts failed, the SMD requirement became a rider to an unrelated bill, ultimately passed as *An Act for the relief of Dr. Ricardo Vallejo Samala (and provide for Congressional redistricting)*, Pub. L. No. 90-196, 81 Stat. 581 (1967).

²A parallel can be drawn to Iowa's Democratic and Republican Caucus procedures: candidates must meet a minimum "viability" threshold, and supporters of non-viable candidates are given the chance to re-allocate their votes after each viability round.

2 Methods

We conduct this study by linking together districts and ballots to estimate outcomes. First, we generate large sets, called *ensembles*, of valid multi-member districting plans for each state (see §2.1 for details). Second, we generate *simulated elections* using the RCV models introduced in Benade et al. [2]. Combining these, we can then run STV algorithms to estimate projected representation for POC candidates of choice under the Fair Representation Act (see §2.2).




Figure 1. High-level view of methods. We generate an *ensemble* of random multi-member districting plans; we generate *simulated elections* based on voting history in each state; then we run the STV algorithm to combine the districts and votes into outcomes. Figure 7 shows the results for the whole nation, with yellow boxes for the statewide share of minority population and blue circles for the projected share of minority-preferred representation.

2.1 Building multi-member plans

To create large batches of multi-member districting plans, we use GerryChain, an open-source software library developed by the MGGG Redistricting Lab that implements a "recombination" (or ReCom) algorithm for district generation [9, 4]. ReCom lets us randomly "walk" through the space of possible districting plans, adding one new plan at each step in the process until we have created a large ensemble of plans. To generate each successive plan, ReCom merges two districts in the current plan and re-splits them into a different pair of districts. (See Appendix B for more discussion.)

For this study, we have developed a modification for the ReCom algorithm to allow the creation of districts of different sizes. This allows us to generate, for each state, two ensembles of 100,000 valid multi-member redistricting plans. The first ensemble is a *neutral* ensemble, in which the chain run always accepts any valid redistricting plan that is proposed. The other is a (heuristically) *optimized* ensemble based on the idea that districts have a *threshold* for electing candidates—in districts with 3, 4, or 5 members to be elected, that threshold is 1/4, 1/5, or 1/6 of the electorate, respectively. The optimized ensemble seeks plans that distribute the minority population to maximize the number of threshold increments. If voting were completely polarized by race and ethnicity, this would get the best possible representational outcome.

Statistics over these ensembles help us to survey the redistricting landscape of a given jurisdiction, and to investigate the claims that STV is likely to lead to more proportional representational outcomes. Comparing "optimized" to "neutral" ensembles allows us to probe how outcomes are sensitive to the demographic composition of the districts.

2.2 Simulating elections

The political science literature contains many models of voter behavior in plurality elections, but almost nothing to model voter ranking behavior. Moreover, the shortage of currently implemented RCV voting systems nationwide means researchers do not yet have access to a large dataset with real ranked-choice ballots that could help fit a model to observed voting behavior. Here, as in earlier work, our solution is to use and compare multiple ballot generation models. All of these choices are customizable.

Ranking models

We use four probabilistic models to generate ballots: Plackett-Luce (where voters have an overall preference between two slates and then flip a weighted coin to choose from each); Bradley-Terry (where the likelihood of a given ballot is based on how it ranks the candidates pairwise); Alternating Crossover (where every voter is either a bloc voter whose ballot type puts one slate entirely above the other or an alternating voter who trades off between the two slates); and the Cambridge Sampler (where ballot types are chosen at random from actual historical RCV elections in Cambridge, MA). Instead of choosing between these models of voter behavior, we run them all and report the results split out by model before aggregating. The models are described in moderate depth in the supplementary material (Appendix C) and in greater depth in Benade et al. [2].

Parameters for generating ballots

- **District specs, voters, and candidates.** The number of representatives to be elected, the number of White-preferred and POC-preferred candidates running for office in each district (referred to as *candidate pools*), and the share of minority voting age population in each state (POCVAP share).

To account for the variety in expected candidate availability, we construct four candidate pools for each district. In the results below, we use candidate pools that have either 1.5 times or 2 times the number of candidates as seats to be filled, rounded down. Then we either assume that half the candidates will be POC-preferred, or that the POC candidate share will reflect the POC population share.

Example: Suppose a given district has five seats to be filled and 30% of the electorate is people of color.

Pool 1	Seven total candidates: $4C + 3W$
Pool 2	Seven total candidates: $3C + 4W$
Pool 3	Ten total candidates: $5C + 5W$
Pool 4	Ten total candidates: $3C + 7W$

- **Polarization between the slates.** For each group, how strongly do White and POC voters diverge on their preferred slate of candidates? We begin by supposing that there are minority-preferred candidates (denoted C) and White-preferred candidates (denoted W). We then introduce parameters to record the level of cohesion—for instance, $\pi_C = .70$ indicates that POC voters tend to support the slate that is POC-preferred overall at a rate of 70%. Likewise, $\pi_W = .60$ would tell us that 40% of White

voters tend to "cross over" to support the minority-preferred slate (since 60% cohesively support the White-preferred slate). The estimated π_W and π_C for all states can be found in Table 1.

- **Consensus within the slate.** These parameters describe how much consensus each group has on the preferred ordering within each slate of candidates. For instance, it is important to distinguish whether there is a single "strong candidate" or whether voters spread their support evenly among those on the ballot within their preferred slate.

Polarization parameters for each state and more detailed explanations can be found in Appendix C.

2.3 National execution and focus states

Using the input data described above, we generate neutral 100,000-plan ensembles in all 50 states, and optimized 100,000-plan ensembles in each focus state. Then, in each focus state, we sub-sample 50 plans from both the neutral and optimized ensembles; in all other states, we sample five plans from the neutral ensemble. For each district in these sub-samples, we simulate five ranked-choice voting elections for every combination of model, candidate pool, and support scenario. Furthermore, each election sees one ballot cast for every 45 voters in the district: for example, if a 5-member district contains 1,580,400 voters, then 35,120 ballots are cast in each simulated election.³ We then aggregate these per-district election results to make comparisons across models, support scenarios, candidate pools, and evaluate outcomes at the state level.

Because the FRA calls for multi-member districts and STV nationwide, we have compiled data in all fifty states. For the purposes of exposition, we present a more detailed discussion for the following five states.

1. **Florida**, which has significant Black and Latino populations, and was apportioned an additional (28th) Congressional seat following the 2020 Decennial Census.
2. **Illinois**, which saw modest population decline between 2010 and 2020, but has a large Black population and rapidly growing Asian and Latino populations.
3. **Maryland**, which gained more than 7% population from 2010 to 2020, and has growing Black and Latino populations.
4. **Massachusetts**, which has a similar size to Maryland but a much smaller POC population share (29.56% POCVAP in MA versus 50.13% in MD).
5. **Texas**, which has gained nearly 16% population from 2010 to 2020 (and was consequently apportioned two additional Congressional districts), largely driven by marked increases in Black, Latino, and Asian population growth over that span.

Given their variation in Congressional delegation size, POC population size, partisan lean, and district magnitudes, we believe that these five states have strong illustrative value. We refer to these states collectively as *focus states*.

2.4 Limitations of approach

The study design centers the important reality that racial and ethnic groups are highly ideologically diverse, by allowing flexibility for voters to cast a wide variety of ballots without assuming universal bloc voting. However, there are several limitations of our approach that are important to flag. First, we consolidate multiple racial and ethnic groups as "POC" and assume a modest level of polarization, with at least 60% of people of color supporting one slate of candidates and at least 60% of White voters supporting a disjoint slate (see Table 1). Though we think that this is a reasonable assumption in general, the different groups

³The ballot-to-voter ratio of 1 : 45 was determined by repeatedly conducting simulated RCV elections at different scales, settling on a ratio that produced robust results and tolerably short computation times.

may vote less coalitionally in some parts of the country, and the overlap of preferred slates will surely be messier in practice. Second, our candidate pool construction assumes that POC-preferred candidates are available *at least* in proportion to the POC population share; this may be overly optimistic, especially in some Northern states like Pennsylvania and Massachusetts, where POC candidates for statewide and federal office are rare. It is certainly possible that, as some advocates have argued, the shift to an RCV system would itself encourage more candidates of color to run for office—or, in the case of partisan elections, would encourage the parties to be less exclusionary in their candidate support. But this remains an open question and a major factor in the ultimate success of a new system. Finally, by mostly presenting results in an aggregated fashion below, we lose some ability to discern which combinations of parameter settings and models are producing which projections.

In sum, we feel that this approach is very sound to highlight trends and the overall promise of the FRA-style electoral system, but that users and analysts should re-run the models with assumptions and settings best suited to their state or locality. All replication materials are available to the public at [7].

3 Detailed findings from focus states

For the five focus states, we present a short discussion to accompany a set of figures that show the results of our modeling. Each figure contains four plots.

- **Minority share across districts.** In a 3-member district, the STV electoral representation threshold is 25%—any candidate with this much first-place support is sure to be elected. In a 4- or 5- member district, the threshold is 20% or about 16.7%, respectively. This plot shows statistics for an ensemble of multi-member districts made to be equipopulous, compact, and connected, but without any attention to race. Colored lines show the thresholds for each type of district.
- **Simple seat projections.** How many thresholds were crossed overall, summed over the districts? The darker green shows this for the neutral ensemble; the lighter green shows an alternative ensemble that is heuristically optimized to cross more thresholds. That is, the lighter green districts are drawn in a race-conscious way to try to optimize performance for people of color.
- **Detailed seat projections, overall.** We then run samples from both sets of districts through the simulated election models. Strikingly, we find overall that the districts drawn to optimize performance do *not* out-perform the ones drawn with no attention to race. In this plot, we have marked the status quo level of POC representation in each state for comparison purposes.⁴
- **Detailed seat projections, by model.** In this plot, we include a marker for the POC proportion of voting age population (yellow) and for the state's level of support for the Biden-Harris ticket (gray). This is included because the polarization parameter in the voter models was set with reference to how POC and White support distributed between the Biden-Harris ticket and the Trump-Pence ticket in the 2020 presidential election. By including those markers, we allow readers to compare the representation projections to both simple demographics and to estimated overall support for a more racially diverse Democratic ticket—as a check that the models are appropriately complex, we want to be sure that neither demographics nor partisan politics are singlehandedly dominating our findings. In order to break down the overall findings, we separate them out according to the four models of voter ranking behavior used to create simulated ballots. We find overall that the Cambridge model tends to have the lowest projected outcomes, which is consistent with findings in [2]—that truncated ballots and different levels of candidate strength between the two slates can combine to cause the POC slate to underperform. This is discussed further in the conclusion.

⁴Current representation numbers were provided to us by FairVote staff.

3.1 Florida

As of the 2020 Decennial Census, Florida's 21,538,187 residents will be represented by 28 House members. Under the FRA, those 28 members would be grouped into one 3-member district and five 5-member districts. Buoyed by fast-growing Latino communities—which jumped in size by over a third since the 2010 Census—Florida's voting age population is roughly 45% POC.




Figure 2. Florida's 28 seats are divided into one 3-member district and five 5-member districts. Florida has 45% POCVAP statewide and Biden earned 48% of the major-party vote in 2020.

Figure 2 shows the demographic distributions observed in our neutral ensemble. Note that all of the five-member districts, and likely all of the three-member districts, have POC population exceeding the threshold for election. This suggests that this election system would likely result in at least one POC-preferred candidate elected in each district; that is, every person in Florida would be represented by a POC candidate of

choice in their own district, a major departure from the status quo.

Using the models of voter behavior, we can predict potential representational outcomes under an STV system. Two models—Plackett-Luce and Bradley-Terry—predict that Floridians, across both ensembles, all candidate pools, and all support scenarios, typically have the opportunity to elect 12 to 14 POC candidates of choice statewide in a delegation of 28 members, which is in line with POC proportionality in voting age population. The Alternating Crossover model skews slightly lower. The Cambridge Sampler, which takes bullet voting into account, has a "long tail" to the left, but still puts the likely range at 9-13, with 11-13 seats most likely. For all models and all situations taken together, POC-preferred candidates most often secure 11-14 seats, a likely improvement on the status quo POC representation of 11.

3.2 Illinois

Illinois has a reported 2020 Decennial Census population of 12,812,508, and lost a seat to drop to 17 Congressional districts. Using the FRA guidance for selecting an optimal district configuration, Illinois would have four 3-member districts and one 5-member district. The state has substantial Latino, Black, and Asian populations, and has nearly 40% POCVAP overall.

In Figure 3, we show the distribution of POCVAP shares from our neutral districting ensemble. Across all plans in the neutral ensemble, the district with the lowest POCVAP share often fell short of hitting the STV representation threshold of 25%, but the other districts all tend to cross the threshold at least once. The number of threshold increments hit overall by all districts is 5-7 in the neutral ensemble and can easily achieve 8 if plans are selected with this priority. However, the simulated election analysis predicts higher shares significantly higher than the simplified threshold count: 8 to 11 POC-preferred candidates rather than 6-8. Furthermore the "optimized" districts perform no better than the neutral ones. The elevated outcomes in the detailed model relative to the simple model are due to substantial crossover support estimated from White voters. As usual, the projection is most uncertain in the Cambridge Sampler model, where voters can and often do cast a short ballot.




Figure 3. Illinois' 17 seats are divided into four 3-member districts and one 5-member district. Illinois has nearly 39% POCVAP overall, and supported the Biden-Harris ticket at 58.66%.

3.3 Maryland

As of the 2020 Decennial Census, Maryland's total population was 6,117,224, centered in the Washington-Baltimore metropolitan area. With the apportionment of eight Congressional districts, the FRA-optimal districting configuration consists of one 3-member district and one 5-member district. The state is nearly one-third Black and 11.8% Latino—a significant increase from the previous Census—and more than half of its voting-age residents are people of color, making it one of only six majority-minority states.⁵




Figure 4. Maryland's 8 seats are grouped into one 3-member district and one 5-member district. The state has just over 50% POCVAP and supported Biden-Harris at roughly 67%.

Figure 4 shows that the composition of the larger (5-member) district mirrors that of the state, while the

⁵The other states with more than half POCVAP are California, Hawaii, Nevada, New Mexico, and Texas—see Table 1.

smaller district is similar but with larger variance.

As of early 2022, the newly-enacted single-member Maryland Congressional districting plan has three districts out of eight with a greater-than-statewide POCVAP share: district 4 at 80%, district 5 at 62%, and district 7 at 65%. Each of these districts draw from sizeable POC populations near Northeast DC, Southeast DC, or Baltimore, but have starkly elevated POC shares compared to the state at large—leading to depressed minority shares in the other five districts. This plan would therefore be highly likely to have three POC-preferred representatives and five who are not. The Maryland plan, however, is still being contested.

Ranked choice projections, instead, show that 5-6 seats out of eight would be likely to align with POC preferences. This is because ranked choice allows for White crossover support to coalesce with POC voting more easily than in a single-member district plan.

3.4 Massachusetts

As reported by the 2020 Census, Massachusetts' statewide population is 7,029,917, a 7.4% increase over ten years. Overall, the Massachusetts population is just over 32% POC, and its voting-age population is nearly 30% POC.

In Figure 5, we visualize the POC voting-age population shares in the Massachusetts neutral districting ensemble, which are extremely tightly clustered, with roughly 30% in each multi-member district, as in the state as a whole. This reflects a phenomenon on the level of race that is well documented on the level of party preference: Massachusetts has a fairly "rigid" districting problem, with not a great deal of downstream variation in the properties of districts as the lines are varied. (See also [3], which studies this from the point of view of political competitiveness.)

With an alternative ensemble that is optimized to cross more threshold increments, we can drive up the share of the ensemble that crosses the threshold once per district, but we never get a fourth opportunity. (We note that this might be a reason that ranked choice advocates should potentially prefer a 5/4 magnitude split to a 3/3/3 split.) This means that if voters in Massachusetts voted strictly by race with no crossover voting—the simple seat projection—then we would expect three seats for POC candidates of choice.

However, all four RCV models predict that voters of color are likely to elect more than three candidates of choice, often four or five out of nine representatives. For those who know Massachusetts electoral history, in which Ayanna Pressley bucked the Democratic establishment to challenge Mike Capuano and secure the first-ever MA Congressional seat for a candidate of color, this will seem overly rosy. The reason for this high prediction is twofold: first, our models assume that many viable POC-preferred congressional candidates will be present on the ballots, which is aspirational in a state whose Democratic bench is dominated by White politicians. Besides Pressley, Leland Cheung and Jay Gonzalez were the only plausibly viable POC candidates for statewide office in the last ten years. Both were blown out by more than 20 points in their respective elections. It is clear that candidate availability will be an impediment to realizing the prediction made by this model.

A second significant issue is related to the first: because of the lack of POC-preferred candidates in recent history, it is difficult to reliably estimate the levels of voting polarization. In Massachusetts, the overwhelmingly Democratic lean in presidential elections means that support for the Biden-Harris ticket may be especially weak as an indicator of White willingness to support POC candidates of choice.




Figure 5. Under the FRA, Massachusetts’ nine representatives would be grouped into three 3-member districts. Nearly 30% of the state’s voting-age population are people of color.

3.5 Texas

Texas has grown rapidly since the 2010 Census, adding nearly four million people and gaining two Congressional districts, bringing the number to 38; under the FRA, the optimal districting configuration is one 3-member district and seven 5-member districts. This population growth can be attributed to burgeoning Black and Latino populations, which grew by more than 25% and 20%, respectively, so that Texas is now a majority-minority state overall, with more than 56% POC VAP share.

Figure 6 shows the POC VAP distributions for the Texas ensemble of multi-member districts. The typical plan will likely allow for POC voters to be represented at least once in each multi-member district.

Unlike the other focus states, there is a reasonable chance that outcomes could underperform POC population proportionality, as well as underperforming the simple seat projection made by counting thresholds. When the simulated RCV election outputs are disaggregated by model, each model echoes this shortfall. The classical probabilistic models (Plackett-Luce and Bradley-Terry) split the difference between population proportionality and partisan lean. Here, the Alternating Crossover model shows a dichotomy, with one frequent outcome at 18 seats and another at 22. This *bimodal* property in the AC model can also be observed at a lower level in other states, such as Florida (Figure 2.) It indicates that the outcomes can depend on other parameters not broken out here, such as whether there are particularly strong candidates in the White- and POC-preferred slates.

The estimated likelihood that any voter of color prefers the POC slate of candidates sits at 70%, while the estimated cohesion for White voters is higher, at 75% (Table 1). While this still reflects a clear and stark assessment of voting polarization, this makes Texas one of only a handful of states to be modeled with higher White than POC cohesion—indeed, Texas and Oklahoma are the only two with significant POC population to have this property. This accounts for POC seat projections falling short of population proportionality.

The full sweep of statewide results shows a wider range of possible outcomes than in many other states: it is reasonable to project that Texan voters of color may end up with 16-21 representatives of their choice, with most of the probability in the 17-20 range. These results should be compared to the current status quo, marked at 12 POC representatives.




Figure 6. The 38 Texas representatives are divided between one 3-member district and seven 5-member districts. More than 56% of the state’s voting-age population is people of color, while the Biden-Harris ticket had roughly 47% support.

4 National comparison and conclusions

In addition to our focus states, we created randomized multi-member districts and conducted simulated RCV elections in the remainder of the states that have a Congressional apportionment of three or more. Figure 7 shows nationwide results, first alphabetically by state name and then re-organized in ascending order of POC population share. Generally, these results confirm that **predicted RCV outcomes tend to track with proportionality**, as advocates have claimed.⁶

A second finding emerges from the detailed state-by-state analysis seen in the focus states above. In each state, we started by simply counting thresholds crossed in the districts generated at random.⁷ It is well-known that two perfectly bloc-voting groups (called "Solid Coalitions" in the social choice literature) will tend to get representation proportional to their voter numbers in STV systems. This means that a plan that maximizes threshold-crossing would maximize representation if everyone was bloc-voted along racial and ethnic lines. But interestingly, we found that when more realistic crossover voting and polarization is assumed, the plans with higher threshold-crossing numbers perform no better than the neutral ones. This is important for people engaged in the reform effort: it suggests that STV voting is far less sensitive to exactly how the districts are drawn—**ranked choice can secure proportionality without race-conscious line-drawing**.⁸

The tools and results here also allow us to drill down and find other actionable issues for reformers, such as by spotlighting which combinations of voter behavior and candidate slates can lead to anomalous underperformance. The dynamics around ballot truncation and strong candidates are particularly worthy of further study.⁹ In addition, we find that the results are fairly robust against the possibility of systematically lower turnout by people of color (see Supplement D).

The Fair Representation Act aims to establish an electoral system which provides all voters—and in particular, voters of color—with fair representation. Having modeled potential ranked choice outcomes across the country, with a careful look at five focus states, we find extremely encouraging results. Single transferable vote in multi-member districts can secure proportional representation for minorities without a race-conscious line-drawing process. In a time when our courts are increasingly and demonstrably averse to race-conscious policy-making, this becomes a very attractive system to help the United States live up to the ideals of its representative democracy.

⁶See also [5], a 2021 national study by Garg et al. that finds similarly encouraging proportionality results under the more restrictive assumption of solid bloc voting. They also find in that setting that the dependence of outcomes on district design is reduced as the district magnitude increases from single-member to larger multi-member districts.

⁷For instance, if a Maryland plan had 42% POC population in its 3-member district and 52% POC population in its 5-member district, then since the election threshold is 25% for 3-member districts and 16.67% for 5-member districts, we'd pass one threshold in the smaller district and 3 thresholds in the larger district, adding up to four in all. This produces a "simple seat projection" of 4 seats out of 8 for POC-preferred candidates.

⁸And indeed, race-conscious line-drawing, at least with the simple heuristic optimization used here, does not make any appreciable difference in the modeled outcomes.

⁹These effects show up in the current project with the Cambridge Sampler model of voter behavior. The previous paper [2] found the same phenomenon, and isolated some of the parameter settings where its effects are most severe. In short, ballot truncation can exaggerate the effects of subtly different voter behavior between the two blocs.




Figure 7. Nationwide projections—using random districts, varied voting scenarios, and simulated election results—for the 37 states with at least three U.S. House members. States are ordered either alphabetically (top) or by ascending share of voting age population for people of color (POCVAP, bottom). Larger circles indicate greater frequency. To facilitate comparisons, the yellow squares show the statewide POCVAP share and the red dots show the status quo (2021) POC share of Congressional representation. Compare Figure 11, which assumes systematically lower turnout for people of color.

References

- [1] H.R. 3863 – Fair Representation Act.
- [2] Gerdus Benade, Ruth Buck, Moon Duchin, Dara Gold, and Thomas Weighill. Ranked Choice Voting and Proportional Representation. *Available at SSRN*, 2 2021. <https://papers.ssrn.com/abstract=3778021>.
- [3] Daryl DeFord, Moon Duchin, and Justin Solomon. A computational approach to measuring vote elasticity and competitiveness. *Statistics and Public Policy*, 7(1):69–86, 2020.
- [4] Daryl DeFord, Moon Duchin, and Justin Solomon. Recombination: A family of Markov chains for redistricting. *Harvard Data Science Review*, 3 2021. <https://hdsr.mitpress.mit.edu/pub/lds8ptxu>.
- [5] Nikhil Garg, Wes Gurnee, David Rothschild, and David Shmoys. Combatting gerrymandering with social choice: the design of multi-member districts. *arXiv preprint arXiv:2107.07083*, 2021.
- [6] MGGG Redistricting Lab. Analysis of county commission elections in Yakima County, WA. <https://mggg.org/uploads/Yakima.pdf>.
- [7] MGGG Redistricting Lab. FairVote replication repository. <https://github.com/mggg/FairVote-FRA/>.
- [8] MGGG Redistricting Lab. Findings on the City of Lowell’s Election Systems. <https://mggg.org/publications/Lowell-Detailed-Report.pdf>.
- [9] MGGG Redistricting Lab. GerryChain. <https://github.com/mggg/GerryChain>.
- [10] U.S. Census Bureau. Decennial Census P.L. 94-171 Redistricting Data. <https://www.census.gov/programs-surveys/decennial-census/about/rdo/summary-files.html>.
- [11] U.S. Census Bureau. Standard Hierarchy of Census Geographic Entities. <https://www2.census.gov/geo/pdfs/reference/geodiagram.pdf>.
- [12] US Census Bureau. Post-Enumeration Survey and Demographic Analysis help evaluate 2020 Census results, 2022. <https://www.census.gov/newsroom/press-releases/2022/2020-census-estimates-of-undercount-and-overcount.html>.
- [13] A.P. VoteCast. National Exit Polls: How Different Groups Voted. <https://www.nytimes.com/interactive/2020/11/03/us/elections/exit-polls-president.html>.
- [14] A.P. VoteCast. National Voter Surveys: How Different Groups Voted. <https://www.nytimes.com/interactive/2020/11/03/us/elections/ap-polls-national.html>.

A Data preparation

For each of the five focus states, we create a geometric and demographic data product using a combination of 2020 Census demographic, 2020 Census geometric, and state-distributed geometric data.¹⁰ These data products consist of 2020 population and demographic data adjoined to a state’s 2020 Census *voting tabulation districts*, or *VTDs*, which correspond closely to the voting districts used in the state’s most recent election.¹¹ To create these population- and demographic-augmented VTDs, we follow three general steps:

1. **Retrieve Census geometries.** The U.S. Census reports its data at varying levels of granularity, ranging from census blocks (finest, sometimes as small as a city block) to block groups, tracts, counties, and so on. Population and demographic data are reported at the Census block level and blocks nest in all other geographic units.
2. **Join data from the 2020 Decennial PL94-171.** We download Census block-level data from tables P1 and P4, respectively [10]. We then attach these population and demographic data to the block geometries from the previous step.
3. **Aggregate block-level data to VTDs.** Because blocks tile VTDs—that is, each block belongs to exactly one VTD, and no two blocks overlap—we assign each block to a VTD and aggregate population and demographic data from the block level up to the VTD level [11].

VTDs are small enough to let us create diverse sets of districts, but can still be handled in a computationally efficient way. In addition, they can often be joined to historical election data.

B Ensemble methods

B.1 Basic ReCom

The basic recombination step fuses two districts at a time (in Figure 8, red and blue are chosen) and then re-splits the combined area into two new districts. The splitting is accomplished by selecting a random spanning tree of the double-district and cutting an edge that leaves balanced complementary components. See [4] for details.




Figure 8. A single ReCom step on a four-by-four grid. In step 1, we are given an initial, population-balanced districting plan. In step 2, we choose the **red** district and the **blue** district, then merge them together to form the **purple** district. In step 3, we split the **purple** district into two population-balanced pieces.

¹⁰Though the Decennial Census data are authoritative, the dataset is known to have some systematic error. According to the Census Bureau’s post-enumeration survey (PES), the Black, Hispanic, and Native American (on reservation) populations were undercounted by 3.30%, 4.99%, and 5.64% respectively, while the non-Hispanic White alone population was overcounted by 1.64% [12].

¹¹A handful of states did not participate in the 2020 Census VTD collection program. In these states, we use Census block groups in place of VTDs.

B.2 Multi-member ReCom

With this modified splitting step, we can fuse adjacent multi-member districts and divide the resulting super-district into new districts of the appropriate sizes. For example, suppose we have two adjacent multi-member districts; district A is represented by five representatives, district B by three. If a ReCom step selects A and B to be merged and re-split, we draw a spanning tree of the union and search for an edge to delete so that the complementary components retain approximately the 5-to-3 population ratio.

For each plan encountered by the chain, we calculate and store the following statistics.

- Total population of each district and the corresponding number of representatives;
- The POC voting-age population (POCVAP) and total voting-age population (VAP) of each district;
- A "tilt" score of each plan (described in the next section) that attempts to optimize for POC representation; and
- An overall measurement of compactness called the *cut edges* score [4].

B.3 Preferential plan acceptance

At each step in the Markov chain, a new districting plan P is proposed. In a neutral chain, every valid proposed plan is accepted regardless of its characteristics (beyond the requirements of contiguity and population balance). In heuristic optimization chains, on the other hand, we preferentially accept a proposed plan based on how it compares to the current plan. Here, let $P = (P_1, \dots, P_k)$ represent a districting plan with k districts; let m_i be the magnitude of district P_i , and set $t_i = 1/(m_i + 1)$ to be the corresponding electoral threshold. Let Q count the number of times that the threshold is exceeded in a given district:

$$Q(P_i) = \left\lfloor \frac{\text{POCVAP}_i}{t_i} \right\rfloor,$$

where POCVAP_i is the POCVAP share in district P_i . Then $Q(P)$ is the sum of the district thresholds surpassed in its districts. For example, if a district has 3 members and 48% POCVAP share, we have

$$Q(P_i) = \lfloor 0.48/0.25 \rfloor = 1.$$

By using a chain run that preferentially accepts proposed plans with a higher Q score, this plan generation method encourages the creation of plans that score higher on this statistic. If the groups voted as solid coalitions, this would correspond to the highest possible representation.

C Ranking models

C.1 Polarization and concentration parameters

C.1.1 Polarization parameters

We begin with the assumption that there is a slate of candidates that are White-preferred and a disjoint slate of candidates that are POC-preferred overall. The *polarization parameters* $0 \leq \pi_C, \pi_W \leq 1$ encode the strength with which a POC voter prefers the POC slate or a White voter prefers the White slate, respectively. This is used somewhat differently depending on the model, but in all four models this will mean that if, for instance, $\pi_C = .75$, then voters of color will rank a POC-preferred candidate at the top of their ballot 75% of the time, in expectation.

State	POCVAP%	Biden Share	POC Cohesion			White Cohesion	
			Reported	Inferred	π_C	Reported	π_W
Alabama	34.53%	37.09%	0.81	0.83	0.9	0.77	0.85
Alaska	38.52%	44.74%	0.44	0.47	0.6	0.53	0.6
Arizona	41.93%	50.16%	0.59	0.62	0.7	0.52	0.6
Arkansas	28.48%	35.79%	0.74	0.8	0.9	0.72	0.75
California	62.13%	64.91%	0.75	0.78	0.8	0.47	0.6
Colorado	31.35%	56.94%	0.55	0.57	0.6	0.41	0.6
Connecticut	33.3%	60.2%	0.75	0.82	0.9	0.43	0.6
Delaware	37.56%	59.63%	0.84	0.86	0.9	0.48	0.6
Florida	45.35%	48.31%	0.67	0.67	0.7	0.62	0.75
Georgia	47.18%	50.12%	0.81	0.82	0.9	0.69	0.75
Hawaii	76.16%	65.03%	0.64	0.66	0.7	0.35	0.6
Idaho	18.52%	34.12%	–	0.56	0.6	0.64	0.75
Illinois	38.76%	58.66%	0.81	0.88	0.9	0.5	0.6
Indiana	21.7%	41.8%	0.72	0.81	0.9	0.61	0.75
Iowa	14.54%	45.82%	0.64	0.78	0.8	0.55	0.6
Kansas	24.46%	42.51%	0.62	0.73	0.8	0.59	0.6
Kentucky	16.6%	36.8%	0.62	0.65	0.7	0.66	0.75
Louisiana	41.69%	40.54%	0.79	0.82	0.9	0.77	0.85
Maine	9.21%	54.67%	0.47	0.86	0.9	0.44	0.6
Maryland	50.13%	67.03%	0.88	0.92	0.9	0.46	0.6
Massachusetts	29.57%	67.12%	0.8	0.89	0.9	0.37	0.6
Michigan	25.09%	51.41%	0.8	0.83	0.9	0.55	0.6
Minnesota	20.05%	53.64%	0.66	0.71	0.8	0.47	0.6
Mississippi	42.24%	41.62%	0.9	0.94	0.9	0.81	0.85
Missouri	22.03%	42.16%	0.71	0.8	0.9	0.62	0.75
Montana	14.42%	41.6%	0.45	0.61	0.7	0.58	0.6
Nebraska	20.79%	40.22%	0.58	0.69	0.7	0.61	0.75
Nevada	50.04%	51.22%	0.62	0.66	0.7	0.56	0.6
New Hampshire	11.15%	53.75%	0.56	0.74	0.8	0.46	0.6
New Jersey	45.53%	58.07%	0.76	0.78	0.8	0.49	0.6
New Mexico	59.48%	55.52%	0.61	0.63	0.7	0.48	0.6
New York	45.54%	61.73%	0.77	0.79	0.8	0.48	0.6
North Carolina	36.36%	49.32%	0.79	0.8	0.8	0.66	0.75
North Dakota	15.59%	32.78%	–	0.67	0.7	0.66	0.75
Ohio	21.6%	45.92%	0.79	0.82	0.9	0.6	0.6
Oklahoma	34.77%	33.06%	0.49	0.51	0.6	0.71	0.75
Oregon	25.05%	58.31%	–	0.6	0.7	0.39	0.6
Pennsylvania	23.85%	50.59%	0.82	0.87	0.9	0.57	0.6
Rhode Island	27.62%	60.6%	–	0.81	0.9	0.41	0.6
South Carolina	35.14%	44.07%	0.78	0.79	0.8	0.73	0.75
South Dakota	16.91%	36.57%	–	0.61	0.7	0.63	0.75
Tennessee	26.49%	38.17%	0.8	0.86	0.9	0.69	0.75
Texas	56.84%	47.17%	0.67	0.68	0.7	0.66	0.75
Utah	22.64%	39.31%	0.46	0.6	0.7	0.59	0.6
Vermont	9.6%	68.3%	–	0.99	0.9	0.31	0.6
Virginia	38.67%	55.15%	0.76	0.76	0.8	0.53	0.6
Washington	32.64%	59.93%	0.49	0.5	0.6	0.36	0.6
West Virginia	9.29%	30.2%	–	0.91	0.9	0.71	0.75
Wisconsin	18.17%	50.32%	0.73	0.75	0.8	0.52	0.6
Wyoming	16.15%	27.52%	–	0.24	0.6	0.69	0.75

Table 1. Polarization levels inferred from A.P. VoteCast exit polls [13] and surveys [14] of the 2020 Presidential election (Biden-Harris vs. Trump-Pence). Each group's "Reported" column reports raw polling data, if available. We compare the reported POC support for Biden-Harris to the level inferred from reported White support, which typically has a larger sample size. Finally, we report the parameters π_C and π_W used in the model runs.

Typically, racially polarized voting (or RPV) techniques are used to measure the level of support for different racial and ethnic groups for a set of candidates. However, with a national scope for this report, state-by-state RPV is impractical. Therefore we will use a combination of exit polls and voter surveys as a guide to the state-by-state levels of polarization [13, 14].

The MGGG Lab has conducted RPV studies in numerous states, from Alabama to Wisconsin, and we note that in recent elections, it is not unusual to have extremely cohesive Black support for candidates of choice, at well over 90% levels. Latino and Asian polarization is typically somewhat less pronounced, and White cohesion varies considerably around the country. Based on this experience, we have selected tiers for the polarization parameters:

$$\pi_C \in \{0.6, 0.7, 0.8, 0.9\}; \quad \pi_W \in \{0.6, 0.75, 0.85\}.$$

We believe that White support for the Biden ticket will systematically upper-bound White support for general POC candidates of choice, so π_W (the support of White voters for the White-preferred slate) will be somewhat higher than the reported Trump vote. Likewise, we conclude from comparison to RPV in the states where we have conducted it that the POC support for Biden lower-bounds POC support for a POC-preferred slate within the state—in particular, $\pi_C \geq 0.5$ by definition of "candidates of choice." Table 1 reports the raw numbers and the selected polarization parameters.

C.1.2 Concentration parameters

The *concentration parameters* control the consistency of candidate order among each group's voters. These are four positive numbers $\alpha_{cc}, \alpha_{cw}, \alpha_{ww}, \alpha_{wc}$. The higher one of these values is, the less agreement on rank-order for the specified voters on the specified slate. For instance, if $\alpha_{cc} = 2$, then POC voters will have greater variation in the way they order the POC slate; on the other hand, $\alpha_{cc} = 1/2$ indicates that POC voters broadly agree on the preference order within that slate. Following [2], we describe four broad scenarios that help us compare models.

$(\alpha_{cc}, \alpha_{cw}, \alpha_{wc}, \alpha_{ww})$	The support scenario for α .
$(1/2, 1/2, 1/2, 1/2)$	There are strong candidates from the perspective of both voter groups within both candidate slates.
$(2, 1/2, 1/2, 1/2)$	POC voters have more variation in rank ordering the POC slate, but agree on the ranking of the White slate; White voters have more consensus on both slates.
$(2, 2, 2, 2)$	Preferences vary for both groups and both slates.
$(1/2, 1/2, 2, 2)$	POC voters have greater consensus for both slates of candidates, while White voters have less consensus.




Figure 9. Candidate support values are drawn from symmetric Dirichlet distributions parameterized by a single concentration parameter α . Above are three of these distributions and their respective α values, from which we sample support values for a set of three candidates C_1 , C_2 , and C_3 ; brighter colors indicate greater sampling likelihood. When $\alpha < 1$, the support values sampled are likely to heavily favor one candidate; when $\alpha = 1$, all support values are equally likely to be sampled; when $\alpha > 1$, support is more evenly distributed among the candidates.

C.2 Plackett-Luce

The Plackett-Luce (PL) model is a standard method of generating a random permutation. Here, given the polarization parameters π_C and π_W and concentration vector α , we generate a list of k numbers summing to 1, each of which encodes the likelihood of voting for a particular candidate. This list of numbers, called a *support vector*, underpins the PL model. In our setting, we build one support vector for POC voters and one for White voters. We then construct detailed ballots for each voter: one by one, candidates are randomly selected according to the probability distribution specified by the support vector until the ballot is complete.

C.3 Bradley-Terry

In the Bradley-Terry (BT) model, we calculate *head-to-head* preferences between candidates, rather than choosing candidates one at a time. These head-to-head preferences are determined by each voter’s support vector, but instead of sampling individual candidates one-by-one, we sample entire rankings; the probability that any individual full candidate ranking appears is determined by the order in which each head-to-head pairing of candidates appears. Because the number of possible rank-orderings increases extremely quickly as more candidates are added, it is infeasible to generate all such rankings. Instead, we use an MCMC method to sample rankings.

C.4 Alternating Crossover

Introduced by the Lab in previous studies [6, 8], the Alternating Crossover (AC) model categorizes voters into one of two types: *bloc* or *crossover*. Bloc voters first rank their group’s candidates of choice, then other candidates; crossover voters first select an out-group candidate, then alternate between candidate groups. For example, if there is a race with three POC-preferred candidates and three White-preferred candidates, then all voters select from four ballot types.

After determining each voter’s ballot type, the rank-order of the candidates is sampled based on the voter’s group’s support vector α .

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

Table 2. Ballot types in the Alternating Crossover model.

C.5 Cambridge Sampler

The Cambridge Sampler (CS) model is the only model used which relies on existing ranked-choice ballot data to predict voter behavior [2]. The underlying dataset is ten years' worth of ranked-choice, at-large ballots cast in city council elections in Cambridge, Massachusetts. For each voter, the polarization parameters π are used to randomly assign the first choice on the ballot as C or W; then, the rest of the ballot type is randomly sampled from historical Cambridge ballots with a matching first entry. That historical distribution of ballot types is shown in Figure 10.




Figure 10. The top 30 ballot types in the Cambridge ballot dataset.

For example, if a voter chooses C at the top of the ballot, and if 5% of all Cambridge ballots that start with C are of the form CWC, then there is a 5% chance that the voter fills out a CWC ballot. Once the ballot types are complete, candidate selections are filled into the ballot according to the group support vectors. Importantly, this is the only one of the four models used here that can construct *incomplete ballots*, or ballots which do not rank enough candidates to fill out the ballot. A common type of incomplete ballot is the *bullet vote*, where only one candidate is ranked; the Cambridge Sampler accurately reflects the prevalence of bullet votes, for example.

D Turnout Effects




Figure 11. Low turnout. We recapitulate Figure 7, this time with the assumption that people of color turn out to vote at only 75% the rate of their White counterparts. As before, blue dots (sized by frequency) show the projected representation for POC-preferred candidates; yellow squares show the statewide POCVAP share; red dots show the status quo (2021) POC share of Congressional representation. Even under low turnout, the proportionality trend is robust.