Black Representation and District Compactness in Southern Congressional Districts

Nicholas Goedert
Department of Political Science
Virginia Tech

Robert Hildebrand
Grado Department of Industrial Engineering
Virginia Tech

Matt Pierson
Center for Geospatial Information Technology
Virginia Tech

Laurel Travis
Department of Business Information Technology
Virginia Tech

Jamie Fravel
Grado Department of Industrial Engineering
Virginia Tech

February 2024

Abstract: This paper explores the assumed trade-off between district compactness and Black representation in legislative districts in the American South. We perform analysis both on heuristically generated districts using current US demographics, and on historical congressional maps since the 1970s. Computations are performed using an iterative heuristic to find feasible solutions guided by multiple objectives. We find that while the trade-off has been strongly observed historically, it is possible to effectively address both goals simultaneously in most cases. We are able to demonstrate maps substantially superior to the present enacted maps on both dimensions in at least seven of nine states analyzed. Nevertheless, the trade-off appears more necessary in states with larger and/or more heavily rural Black populations than in more urbanized states, where the drawing of compact Black-influence districts is easier.

I. Introduction

The process of legislative districting is inherently one involving trade-offs between democratic normative values. Among the most frequently cited are: partisan fairness; competitiveness; representation of ethnic, religious, or social minorities; population equality; preservation of "communities of interest"; and "traditional districting principles" like contiguity, compactness, and preservation of political boundaries (see Butler and Cain (1992) for an overview). A map designed to assure representation for a politically cohesive minority may be forced to draw districts that are less competitive; while a map seeking to preserve existing county lines may be forced to split geographically dispersed communities of interest. In this article, we explore one particular trade-off that has been prominent in the political and legal battles surrounding congressional districting for the last four decades: the balance between district compactness and the representation of Black voters.

This article primarily evaluates the trade-off between these two goals through heuristic optimization. We simulate district maps in nine Southern states using 2020 Census data and several different weightings for compactness and Black descriptive representation to estimate the degree to which compactness must be sacrificed to achieve a proportionate Black delegation, and vice versa. We simulate these maps by introducing an algorithmic hill-climbing method of automated districting that allows users to specify multiple objectives in any combination of weightings, and draw districts heuristically optimized to that combination of objectives. While there are notable differences across states, we find that it is consistently possible to simultaneously come close to maximizing one value at relatively modest cost to the other. We also look back at congressional district maps since the 1970s to establish the historical tension between these two norms.

While the trade-offs between many competing districting norms present several interesting and substantively important questions, the specific trade-off between Black representation and compactness is of tangible and urgent legal significance at the American federal level in 2024.

Although the U.S. Supreme Court has taken major steps over the past decade to pull back from adjudicating partisan gerrymandering (e.g. *Rucho v. Common Cause* (2019)) and enforcement of preclearance under §5 of the Voting Rights Act (VRA) (e.g. *Shelby County v. Holder* (2013)), it has simultaneously remained significantly involved in claims of racial gerrymandering. Since 2016, the Court has upheld lower court rulings (sometimes on procedural grounds) in several cases striking down congressional and state legislative maps in Virginia and North Carolina as unconstitutional racial gerrymanders (*see Whittman v. Personhuballah* (2016); *Cooper v. Harris* (2017); *Bethune Hill v. Virginia State Board of Elections* (2017)). And most recently, the Court's decision in *Allen v. Milligan* (2023) mandated a redrawing of Alabama's congressional map to include an additional Black opportunity district. The outcome of this case was surprising to many who expected a conservative court to reject past precedent on racial gerrymandering. Instead, it reinforced the need to draw Black opportunity districts in many Southern states by upholding the "Gingles test", explicitly conditioning the creation of minority opportunity districts on compactness. Beyond the immediate impact on Alabama, *Milligan* will likely lead to changes in congressional or legislative maps in Louisiana and Georgia, and potentially enhance bulwarks against racial vote dilution in many states going forward. Thus, the extent to which facilitating minority representation through districting is conditional on compactness remains crucial to future interpretations of the VRA and Equal Protection Clause for the foreseeable future.

Our approach yields three main substantive contributions. First, we find that in every state simulated, it is possible to generate maps providing close to our best simulated Black representation while still drawing reasonably compact districts. Second, we find that where a trade-off must be made between these two norms, it is more severe in states where the Black population is less urbanized. Third, we find that in at least seven of the nine states simulated, our algorithm generates maps that are substantially superior in *both* compactness and Black representation to the present enacted maps. We present these contributions with the caveat that our algorithm is *not* designed to draw precisely legal maps with any specific state requirements or constitutional test in mind. As background, we also observe a negative correlation between Black representation and compactness in historical maps enacted since the 1970s. This paper additionally represents a methodological contribution in adapting an existing method (Duchin 2018) to create a hill-climbing algorithm to heuristically optimize different objective functions simultaneously. This approach can be used to explore other tradeoffs in metrics with regards to redistricting.

The structure of this paper proceeds as follows. Part II provides background on two subjects: the historically intertwined roles of compactness and Black-influence districts in legal battles involving the VRA and racial gerrymandering, and the use of computer algorithms and simulations in recent districting scholarship and litigation. Part III describes our methodology for selecting states and objective functions, and our heuristic-based approach to optimization for generating district maps. Part IV explores the trade-off between compactness and Black representation from a historical perspective, with analysis of Southern congressional maps since the 1970s. Part V describes the results of our approximately optimized districts in nine Southern states. Part VI analyzes the role of urbanization in the compactness/Black representation trade-off and Part VII concludes with discussion of limitations and potential future research.

II. Background

Legal Background: District Shape and Black Representation in the South

In addition to its role as evidence of racial or partisan gerrymandering, district compactness is viewed by many as a worthy goal in its own right, and is required for congressional districts in the state constitutions of 18 states (Kaufman et al. 2021). Butler and Cain (1992) provide an overview of the normative benefits of compactness, particularly "representational convenience and propinquity" (72). On the first point, Butler and Cain assert "contorted and sprawling districts make it harder for representatives to do their jobs…Representatives may even feel deterred from visiting the remote sections of their districts (Id.). On the second, they describe, "The voter's interest in services is shared with others in the same geographic area;…the inability to exclude local public goods forges a common interest among people who live in a common geographic area" (73). More concretely, recent research has found that improved district compactness may increase both voter turnout (Ladewig 2018; Hayes and McKee 2009) and citizen engagement with their government (Bowen 2014).

But district compactness has also been an important if sometimes ambiguous factor in multiple areas of racial gerrymandering jurisprudence, including establishing both VRA §2 vote dilution and Equal Protection claims. Prior to the 1980s, congressional districts in Southern states were almost never drawn to promote the election of Black representatives. Indeed, Black populations were frequently split across districts to reduce the possibility of a Black candidate winning, while improving the prospects of white conservative Democrats. This political and legal landscape was disrupted by amendments to the VRA in 1982 providing that vote dilution claims could be proven through evidence of discriminatory effect alone, and the subsequent Supreme Court decision *Thornburg v. Gingles* (1986), establishing the conditions under which a minority group could assert such a claim (see Canon 2022 for a review of the role of race in districting and associated research since the *Thornburg* decision). The first prong of the "Gingles test" was that the racial minority group be "sufficiently large and geographically compact to constitute a majority in a single-member district" (478 U.S. at 51), for the first time asserting compactness considerations into the law surrounding racial gerrymandering.

The reaction to the *Gingles* decision, interpreted to require drawing districts to guarantee Black representation, was dramatic, as eleven Southern states drew new Black-majority districts prior to the 1992 elections. However, several new majority-minority districts were extremely strange in shape, linking distant pockets of Black populations by very narrow threads to establish contiguity, or choosing to include or exclude fragments of cities or counties based on their racial composition. These maps suggested that it was not feasible to draw compact districts that would sufficiently represent minority voters to satisfy the VRA and the Gingles test. This provoked both a political and legal backlash, beginning with *Shaw v. Reno* (1993), in which the Supreme Court overturned the congressional map in North Carolina under the Equal Protection Clause. The map at issue in *Shaw* and several other recent North Carolina congressional maps are shown in Figure 1.

[Figure 1 about here]

The *Shaw* plaintiffs made district compactness a central element to their claim, objecting to "redistricting legislation that is so extremely irregular on its face that it rationally can be viewed only as an effort to segregate the races for purposes of voting, without regard for traditional districting principles" (509 U.S at 642). And indeed, the Court found the districts in *Shaw* so bizarre that their shape was dispositive in elevating the Equal Protection analysis to strict scrutiny, holding, "in some exceptional cases, a reapportionment plan may be so highly irregular that, on its face, it rationally cannot be understood as anything other than an effort to "segregat[e] ... voters on the basis of race" (*Id.* at 646-7). Succeeding litigation placed slightly less emphasis on compactness as dispositive, instead interpreting district shape as one element among several that could contribute to proving predominant racial motivation. In overturning a district map in Georgia with slightly less irregular districts than the *Shaw* map, the Supreme Court held in *Miller v. Johnson* that:

> Shape is relevant not because bizarreness is a necessary element of the constitutional wrong or a threshold requirement of proof, but because it may be persuasive circumstantial evidence that race for its own sake, and not other districting principles, was the legislature's dominant and controlling rationale in drawing its district lines. (515 U.S. at 913)

And in subsequent cases in North Carolina, the Court held that non-compact districts could be defended under alternate justifications, including partisan gerrymandering. For example, while the Court overturned the North Carolina map in *Cooper v. Harris* (2017), they downplayed the role of compactness in their decision, stating:

> But such evidence loses much of its value when the State asserts partisanship as a defense, because a bizarre shape—as of the new District 12—can arise from a political motivation as well as a racial one. (citing *Cromartie I*, 526 U. S., at 547) (581 U.S. at 19)

Furthermore, extremely non-compact districts may have become less necessary in the 2000s as the federal government and courts became both more flexible in their acceptance of majority-influence districts (e.g. *Georgia v. Ashcroft* (2003)), and more stringent in their interpretation of the Gingles test (e.g. *Bartlett v. Strickland* (2009)). At the same time, empirical evidence suggested that supermajorities were not necessary to produce very high likelihoods of minority candidate success (e.g. Lublin 1997). Thus, several states since 2000 have drawn fewer extremely intricate districts with 60% or more Black population, instead opting for districts with 40%-50% Black population that will usually still elect Black candidates with the help of Latino or white Democratic voters. As we explore more thoroughly in Section IV, this has had inconsistent effects on district compactness, as it overlapped with a period of increasingly aggressive partisan gerrymandering.[1]

Most recently, a federal district court overturned a congressional map in Alabama as unconstitutionally packing Black voters into a single district (*Singleton v. Merrill* (2022)), inviting analogous claims against the lone Black-majority districts in Louisiana and South Carolina. In this case, the district court affirmed the central role of compactness in adjudicating vote dilution claims, stating "compactness is critical to advancing the ultimate purposes of §2, ensuring minority groups equal opportunity" (Id. at 45), and ultimately holding that "Black voters as a group are sufficiently large and geographically compact to constitute a majority in a second congressional district [in Alabama]" (Id. at 147). And while the U.S. Supreme Court left the single-Black district map in place for the 2022 federal elections, the Court's June 2023 decision in *Allen v. Milligan* (599 U.S. 1)

[1] Note that the map in the bottom-right corner of Figure 1 was overturned in *Cooper v. Harris*. While this map contains many extremely noncompact districts, it was replaced by a much more aesthetically pleasing map in the bottom middle of the figure, which nevertheless resulted in a delegation with identical racial and partisan composition during its two cycles in effect.

reinstated this decision in holding that Alabama's existing districts illegally dilute votes under the VRA, requiring the state to redraw the map with a second Black opportunity district. After resistance from the Alabama state legislature, the District Court installed a map incorporating a 52% BVAP 7th District and a 49% BVAP 2nd District.

The 5-4 decision in *Milligan* came as a surprise to many given the conservative leaning of the Court, and largely upheld existing precedent as it related to the application of VRA §2 and the Gingles test. The Court very clearly leans on the *stare decisis* principle of upholding past precedent, stating "for the last four decades…courts have repeatedly applied the effects test of §2 as interpreted in Gingles and, under certain circumstances, have authorized race-based redistricting as a remedy" (599 U.S. 1 at 33). Justice Kavanaugh's concurrence elaborates on this principle, stating

> "[T]he stare decisis standard for this Court to overrule a statutory precedent, as distinct from a constitutional precedent, is comparatively strict. Unlike with constitutional precedents, Congress and the President may enact new legislation to alter statutory precedents such as Gingles. In the past 37 years, however, Congress and the President have not disturbed Gingles, even as they have made other changes to the Voting Rights Act." (Id. at 35)

In other words, it would be up to Congress, not the Court, to replace the Gingles test if they so choose.

Most significant to our project, the *Milligan* decision reinforced both the continued relevance of minority opportunity districts in voting rights law (particularly Black opportunity districts in the American South) and the central role of district compactness in evaluating when such districts should be required. The opinion emphasizes that the first Gingles prong, "focused on geographical compactness and numerosity, is needed to establish that the minority has the potential to elect a representative" (Id. at 10). It also highlights the relative compactness of proposed districts as persuasive, describing, "with respect to compactness...the maps submitted by one of plaintiffs' experts...perform[ed] generally better on average than did HB 1." (Id at 12). The Court also differentiates this case from adverse precedent like *Shaw* where "we relied on the fact that the proposed district was not reasonably compact." (Id. at 19)

These recent developments, such as the maps proposed by plaintiffs in *Milligan* and parallel lawsuits in Louisiana and Georgia, have called into question the extent to which minority representation and compactness must act as substitutes for each other, demonstrating that in many cases it is possible to promote ample Black representation while still drawing reasonably shaped districts. While past cases and articles have explored the trade-off between compactness and racial opportunity districts in one state with respect to a single enacted or proposed map (e.g. Cirincione et al. 2000 in South Carolina; Altman and McDonald 2018 in Ohio), this article uses a new method of algorithmic optimization to define this trade-off more robustly.

Literature Review: Use of Algorithmic Districting

The past decade has seen a proliferation in the use of automated or algorithmic districting in both scholarship and litigation involving legislative gerrymandering (see Becker and Solomon 2021 for an overview of algorithmic methods used in districting). Typically, algorithmic districting has been used in litigation and political science literature to create a set of plausible maps meeting only very broad guidelines of contiguity and population equality (and sometimes majority-minority districts). This set of plausible maps is then used as a counterfactual to compare against a real-world map. If the real-world map is shown to be an extreme outlier against the set of automated maps on a particular measure (for example, partisan bias), this may be evidence that some additional factor (e.g. race or partisanship) was a dominant factor in drawing the map (e.g. Tam Cho and Liu 2016; Cirincione et al. 2000). In other instances, measuring the distribution of some factor in the set of plausible maps may be used to show that bias may be produced without intentionality (e.g. Chen and Rodden 2015). In some cases, these comparisons might be used to rebut claims that a map is biased under a universal standard (Stephanopolous and McGhee 2015; McDonald et al. 2018, 2015). Simulations as counterfactuals became essential tools in several of the gerrymandering cases with one simulation expert, Jowei Chen (2017), cited in congressional redistricting cases in Pennsylvania, North Carolina, and Michigan.

A variety of methodologies have been used for computer generation of redistricting plans. Most fall into one of two categories: optimization and Markov Chain Monte Carlo (MCMC). Optimization considers an objective and tries to find the best plan with respect to that objective. Optimization techniques have been applied to redistricting for some time, notably Garfinkel and Nemhauser (1970). In contrast to our "heuristic optimization" approach (described below), exact optimization techniques that prove optimality of a single plan have trouble converging due to the large number of possible plans available unless an unrealistically large spatial grain is used (e.g. Swamy et al. 2019). Buchanan (2023) surveys both recent and past results in these optimization areas.

MCMC is a descriptive tool that yields a collection of plans that fall into a category of a "reasonable" plan. This ensemble (usually at least thousands of plans) then provides a distribution of different scores to describe the variations in a state. In contrast to the MCMC approach, our goal is not to produce a set of possible maps given only loose constraints for constitutionality. Rather, we are interested in what happens when optimization techniques are applied to produce maps that maximize a specific objective or set of objectives. For example, what are the consequences in terms of partisan balance or racial representation of a map that attempts only to maximize compactness? Some recent work, much of it within the computer science literature, has optimized districts along a single dimension while constraining on others (e.g. Saxon 2020; Validi et. al*.* 2022; Borodin et al. 2022). Other recent articles generate random maps emphasizing compactness, but use other substantive criteria for the inclusion or exclusive of these maps into larger ensembles (e.g. Becker et al. (2021) generate ensembles meeting certain VRA criteria of Latino representation in Texas). But many works on redistricting have spoken of trade-offs inherent in balancing *multiple* normative goods in drawing maps. Our work is intended to more rigorously explore these trade-offs.

III. Methodology

Hill Climbing Methodology

We use hill climbing, a heuristic optimization approach to generate good maps under certain objectives. A "heuristic" is a reasonably intelligent algorithm that searches for excellent plans, but is not guaranteed to find optimal solutions and does not provide a formal proof of optimality of solutions that it finds. The hill climbing heuristic evaluates neighboring solutions to a current solution and moves to the new solutions when the objective value is improved. To enable the hill climbing approach, we must first construct an initial "feasible" plan that satisfies some constraints that we specify. We accomplish this initial construction and hill climbing parts of the algorithm by adapting the MCMC approach of Duchin (2018). We define a feasible plan to have districts that are contiguous and nearly equal population, and we consider only feasible plans at each step of the algorithm. The methodological innovation in Duchin (2018), also used here, is the mixture of a variety of types of steps that *maintain feasibility*. This distinguishes our approach from Herschlag et al. (2017) as does our aggressive use of large steps at the beginning of the algorithm.

The algorithm proceeds as follows:

1. Construct an initial plan as described below, and use it as the "incumbent plan."
2. Generate a neighboring candidate plan using either a "recom" or "flip" step, described below.
3. Compare the candidate plan to the incumbent plan using the objective function and declare the candidate to be the new incumbent if it strictly improves the objective.
4. Either stop, if the maximum number of improving steps has been exhausted, or return to step 2.

The primary difference between this approach and MCMC is in step 3, where the change of incumbent is made *if and only if the objective is improved*, rather than probabilistically. [2]

[2] While our approach is similar to simulated annealing, it relies on a more extensive use of large steps to avoid local optima. These large steps give wide variation in the maps produced, rendering the more computationally-intensive approaches of tabu search and genetic algorithms unnecessary.

Generating Initial Maps

To generate initial plans, we follow Duchin (2018). We encode the structure of the redistricting problem as an mathematical network with nodes, a.k.a. graph, that represent each census block group and then edges that connect nodes whenever the corresponding block groups are adjacent.[3] Thus the graph encodes a list of each block group and also a list of any pair of block groups that touch. We then use a randomized algorithm (where some decisions are made probabilistically) to find a "spanning tree" of the graph – this spanning tree encodes a minimal set of edges (pairs of block groups that touch) that allow the graph to be connected, keeping just enough edges that allow one to traverse all the nodes in the graph from any starting point. We then try to break this spanning tree into smaller connected graphs that each satisfy our population criteria for a district. This is done with a random search approach that traverses the spanning tree from a random point until it has seen enough nodes to satisfy the population of a district. Traversed nodes are then separated from the spanning tree and labeled as an initial district. The process of searching for cuts is repeated until the graph has been cut into equal size districts. If no feasible partition of the spanning tree is found, a new tree is generated, and the process is repeated.[4]

Permutations

Once this initial feasible plan has been generated, the hill-climbing algorithm searches for improvements using small flip and large recom steps. Small flip steps involve finding census block groups on the edge of a district, flipping them over into the adjacent district, and seeing if feasibility is

[3] Census block groups are geographical units used by the US Census containing an average of 39 census blocks, with populations between 600 and 3,000 people. Block groups are the smallest unit for which we have found it possible to effectively and feasibly generate maps. Other works generating maps by algorithm generally either use block groups (e.g. Cirincione et al. 2000), or somewhat larger precincts (E.g. Chen 2017; Duchin 2018; Becker et al. 2021; Borodin et al. 2022) or census tracts (Saxon 2020).
[4] A detailed description of this can be found in Duchin (2018), and open-source implementation of the code is available in GerryChain (github.com/mggg/GerryChain).

maintained. Large recom steps involve taking a random district and completely recombining it with an adjacent district, and then using the same approach used to generate initial plans to re-split the two combined districts. The important innovation here is that this recom step maintains feasibility, so the algorithm does not need to spend computational time to return to the feasible set.

Another distinguishing attribute of our approach is the aggressive use of large recom steps at the beginning of the algorithm. At each step of the algorithm, the decision of which type of step to use is made randomly, with the probability of taking a small step equal to the percent of the total steps that have been completed. So, for example, if we are only 5% of the way through the algorithm's iterations, then, on average, 95% of our steps are large recom steps.

This process gives an approximation to the objective function values of the optimal solution. While there are no guarantees about the accuracy of this approximation, two results are somewhat reassuring. Firstly, clustering of the objective function values is observed in the results for any one objective function used in the algorithm (See, for example, Figure 7). Recall that each of these replications is started from a different randomly generated redistricting plan. Secondly, the figure in Appendix B shows objective function improvement as a function of iterations. Since the heuristic includes a substantial percentage of recom steps, reducing the tendency to become stuck in local optima, the taper-off of objective function improvement over time is reassuring.

Eligible States

As this article explores the potential for districts to specifically maximize the election of Black representatives, we only simulate maps in Southern states that would meet the *Gingles* test for drawing Black majority or influence districts, but would not meet this test for any other racial minority group (as the potential to draw Latino-majority or influence districts would significantly complicate our objective functions). Among 16 states in Southern region as broadly defined by the US Census, this limitation eliminates five states as not having sufficient Black population, as well as Texas and Florida (as meeting the test for Latinos), leaving us with a group of nine states that will constitute our data set in Section V: Alabama, Georgia, Louisiana, Maryland, Mississippi, North Carolina, South Carolina, Tennessee, and Virginia. While Maryland is not typically grouped in the South culturally or politically, it does contain multiple Black opportunity districts that were expanded under the 1982 Amendments, and thus represents a good comparison with states more universally classified as Southern. Fourteen of these sixteen states (excluding Delaware and West Virginia as too small) are included in the historical analysis in Section IV.

Measuring Compactness

This paper evaluates the compactness of districts using the Polbsy-Popper measure, defined as the area of a district divided by its perimeter squared. As is common, for district k with area A_k and perimeter R_k, we define the Polsby-Popper score of a map to be:

$$\text{PP-score} = \text{average PP-score per district} = \frac{4\pi}{|N|}\sum_{k\in N}\frac{A_k}{{R_k}^2}$$

This has a nice geometrical interpretation showing how the perimeter of a circle with the same area compares to the perimeter of a district. Along with Reock compactness, this is one of two measures of district compactness most often used in districting scholarship and litigation. For example, discrepancies in Polsby-Popper compactness scores were seen as decisive by the Pennsylvania Supreme Court in striking down the state congressional map in 2018 in *League of Women Voters vs. Pennsylvania* (2018), and are the primary measure of compactness cited by the lower court in the aforementioned *Milligan* case.[5]

[5] Though not yet prominent in political science scholarship or jurisprudence, the graph-theory based "cut edges" measure of compactness has also become increasingly common in the computational literature. Though we do not directly optimize on this measure, the within-state correlation between cut edges and Polsby-Popper compactness in our simulation data set is .96, and thus maps that perform well on one measure almost universally also perform well on the other. While Reock compactness and other more recently proposals such as population-weighted measures cannot be as efficiently evaluated computationally by our algorithm, and thus are not feasible objectives at our scale, we also believe our

Estimating Black Representation

In contrast to compactness, measuring expected descriptive Black representation in hypothetical districts is less straightforward, as it is potentially the function of many political and demographic variables. Nevertheless, scholarship has consistently found that the Black population is overwhelmingly the most important explanatory variable in predicting Black representation, typically modeled through a probit or logit curve (e.g. Lublin et al. 2020; Hicks et al. 2017; Casellas 2011; Epstein and O'Halloran 1999; Lublin 1997; Cameron et al. 1996). Additionally, Latino population is often found to be significantly positively correlated with Black descriptive representation, and important differences are seen between Border/Rim South and Deep South states.

Our model adheres closely to this line of scholarship, most recently Lublin et al. 2020, Hicks et al. 2017, and McKee and Springer 2015.[6] We estimate a probit curve for the probability of electing a Black representative using data from all U.S. House elections in the broadly defined South from 1992-2020 with Black voting-age population (BVAP) and Hispanic voting-age population (HVAP) as independent variables and the probability of electing a Black representative as the dependent variable, separately estimating the intercept and slope of Black population for Deep and Border/Rim South states. This data set is composed of 2,260 Southern House races, among which a Black candidate was

---

compactness objective has face validity in generating consistently compact districts even when judged subjectively as shown in the Objective 1 and 2 maps for all states in Figure 8.

[6] These models also idiosyncratically incorporate other demographic variables (e.g. % of government workers in Lublin et al.), but none are found to have consistently significant effects on Black representation. Some models also incorporate campaign-specific variables (e.g. open seats in Hicks et al.) that would be inappropriate or impractical to include in the context of generalizing about hypothetical future elections. Also, Becker et al. (2021) recently take issue with this canonical method, instead using ecological inference and precinct-level election data to estimate which candidates in statewide elections in Texas were the preferred choice of Black and Latino voters. While this method more directly measures minority voter preferences among the existing candidate field, it is not itself a measure of Black descriptive representation and does not address supply-side conditions of when Black candidates choose to run.

elected 309 times (13.7%)[7]. This yields a probit function, used in this paper to estimate the probability of a district electing a Black representative, with features in terms of the ratio of Black voters to the voting population and the ratio of Hispanic voters to the voting population as

$$Pr\ (Black\ Representative)\ =\ \Phi(\beta_0 + \beta_1 \frac{BVAP}{VAP} +\ \beta_2 \frac{HVAP}{VAP})$$

For the Border/Rim south states, including MD, NC, TN, and VA, we estimate $\beta_0 = -4.194, \beta_1 = 9.75, \beta_2 = 3.00$, whereas for the Deep south states, including AL, GA, LA, MS, and SC, we estimate $\beta_0 = -4.729, \beta_1 =\ 10.44,\ \beta_2 = 3.00.$

[Figure 2 about here]

Figure 2 shows the resulting probability of electing a Black representative given a district's Black population for each state. While our simulations use the actual Black and Latino population of each individual district to generate the probability of Black representation, in this figure we *impute a statewide average Latino population* for each state for ease of visualization. We see three clusters of state curves: three Border/Rim South states (MD, NC, and VA) with higher Latino populations (around 10%) at the top; one low-Latino population Rim South state (TN) and one high-Latino Deep South state (GA) in the middle; and four low-Latino Deep South states (AL, LA, MS, and SC) at the bottom. In order to elect a Black representative a majority of the time, a district must exceed Black population of around 40% in the top three states, 43% in the middle two states, and 45% in the bottom four states. This difference across states is greatest around this inflection point: a district with 35% Black population has a 30% chance of electing a Black representative in Maryland by only a 10% chance in Alabama (again imputing statewide average Latino populations). The curves converge in all states in districts with clear Black majorities: a 60% Black district will elect a Black representative 96% of the

[7] To account for potential serial correlation over time in a given district, House races are clustered by district crossed with decade, though this affects only standard errors of the estimates and not the coefficients.

time in Maryland and 91% of the time in Alabama.[8] Our Black Representation objective seeks to maximize the sum of the probability of Black representation across all districts in the state.[9]

It is worth noting that, based on the data used to generate our Black representation functions, representatives elected from Black opportunity districts will virtually always be Democrats. While our algorithm allows for a modest probability of electing a White representative in a Black opportunity district (e.g Steve Cohen in TN-09) or a Black representative in a heavily White and Republican district (e.g. Byron Donalds in FL-19), and that Black members may be elected from either party, no Black Republican has actually been elected from a Black opportunity district. There have been ten Black Republicans elected to Congress since the passage the 1982 VRA amendments (two newly elected in 2022) compared with over 100 Black Democrats; none of the ten have been elected from a district with greater than 20% BVAP, and only one of the ten, Tim Scott (SC), has been elected within the nine states we simulate. Whether the creation of Black opportunity districts results in an *overall* increase in Democratic representation is a question with a long and controversial history. Although this question too complex to sufficiently address in this paper, our simulations suggest that Black representation and Democratic representation are positively correlated in most states in our data set, with the exceptions of Maryland and Virginia (see Appendix D).

[8] Throughout this paper, we refer to districts with at least 50% BVAP as "Black-majority" and districts with at least a 50% probability of electing a Black representative (often in the 40-50% BVAP range) as "Black opportunity". These definitions are not intended to reference or suggest a legal standard for application to the Gingles test or other related litigation, and are not explicitly maximized by our objective functions. Estimating Black representation probabilistically, rather than assessing districts in a binary or categorical way as Black-majority or Black-opportunity districts, allows the algorithm to account for unusual events such as the election of a Black Republican in FL-19 (7% Black) or a White Democrat in TN-09 (66% Black).

[9] Thus, in Maryland for example, it would favor two 45% Black districts (each with a 70% chance of Black representation) over one 65% Black and one 25% Black district (with 99% and 8% chances of Black representation respectively).

Legality of Algorithmic Districts

This paper is focused on the generation of hypothetical maps to meet the joint objectives of compactness and descriptive Black representation while conforming to the constraints of contiguity and approximate population equality. It generates maps based solely on the geography and demographics of the state, and is not strictly focused on the drawing of legal districts or conforming to any specific legal test; indeed, the districts generated by the algorithm might face a constitutional challenge for at least four reasons.

First, our code specifies a maximum deviation of 1% from ideal population in order to facilitate a broad range of alterations. While this threshold is the same or smaller than other comparable works employing districting algorithms, the Court has typically demanded almost *exact* population equality in the construction of congressional districts (e.g. *Karcher v. Daggett* (1983); *Vieth v. Jubelirer,* (2002)). As a result, almost all states draw congressional districts so that the maximum population deviation drawn from U.S. Census figures is one or two *people*. However, exceptions do exist if they serve "legitimate state objectives" (e.g. *Abrams v. Johnson* (1997)). It is unclear given scant existing precedent whether this would pass muster with the courts as a permissible state objective.

More importantly, the Court has overturned maps due to racial consideration for two basic reasons. Maps that draw districts that do not facilitate the election of large, geographically compact, and politically cohesive racial minorities may be overturned under a VRA ill claim of vote dilution. As discussed in Section II, this has sometimes been interpreted to require the creation of majority-minority districts (e.g. *Thornburg v. Gingles*), though minority-influence districts have recently become more accepted (e.g. *Georgia v. Ashcroft*). It is likely that especially the maps drawn by our algorithm that exclusively weight compactness may be vulnerable to such a claim. Additionally, districts that are drawn predominantly with racial considerations in mind may be vulnerable to an Equal Protection claim (e.g. *Miller v. Johnson* (1995); *Cooper v. Harris* (2017)). For example, maps have been overturned under this claim when districts have been drawn that are extremely irregular in shape (e.g. *Shaw v. Reno* (1993)) or when several districts are drawn to meet an exact minority population threshold (*Bethune Hill v. Virginia* (2018)). It seems likely that maps generated by the algorithm that primarily weight Black representation could be vulnerable to this claim. Finally, individual states may have districting criteria such as the preservation of county lines or communities of interest that are not incorporated into our universal constraints.

Further, this paper evaluates the relation between compactness and the easily quantifiable concept of Black *descriptive* representation, rather than more nuanced or subjective concepts of "opportunity to elect" or "candidate of choice". While descriptive representation can be more straightforwardly measured and integrated into a function weighting multiple objectives, it is less analogous to measures likely to be actually used in voting rights litigation.

IV. Analysis: Historical Compactness vs. Black Population Concentration

The first step in our exploration of the trade-off between compactness and Black representation is to ask whether this trade-off has been observed in the past, as historical background to our simulated districts using present data. In other words, did real maps used in Southern states actually find it necessary to systematically reduce compactness in order to enhance Black representation?

[Figure 3 about here]

Figure 3a plots the regional average Polsby-Popper score of CDs since the 1970s.[10] The patterns that emerge from this data are unsurprising given the recent legal history surrounding racial gerrymandering. During the 1970s and 1980s, Black population in most Southern states was diluted across many districts, but these districts were reasonably compact. Following *Gingles*, states across the

[10] Both compactness and proportionality measures for congressional maps in all fourteen Southern states since the 1970s are shown in Appendix A. Compactness data is drawn from the map used in the first election of the decade; measures for mid-decade map changes are not shown.

South were induced to draw majority Black districts in the 1990s, and district compactness dropped dramatically. But decisions like *Shaw* (1993), *Miller* (1995), and *Ashcroft* (2003) reduced both the need and permissibility of majority-minority districts. Thus, districting compactness rebounded, with trends most pronounced in the Deep South, while still remaining lower than the pre-*Gingles* era.[11]

Alongside this, we also measure the extent to which the racial make-up of congressional delegations actually reflected the state. We express this as the difference between the average Black proportion of the congressional delegation throughout a decade and the Black proportion of the statewide population.[12] Regional average proportionality is shown in Figure 3b, where we see that Black voters have been systematically underrepresented in the South throughout the last 50 years (and of course, throughout American history), with underrepresentation being consistently most severe in the Deep South. But the VRA amendments and *Gingles* case in the 1980s did have a dramatic effect, reducing average underrepresentation from about 21% in the 1970s to 7% in the 1990s. Moreover, proportionality remained fairly steady during the 2000s and 2010s despite some of the reversals in compactness seen in Figures 3a. Thus, in at least some contexts, it does *not* appear that compromising compactness is strictly necessary to achieve reasonably proportionate Black representation.

[Figure 4 about here]

To see if there is any direct relationship between district compactness and racial representation, we can plot both of these measures for various states simultaneously. Figure 4 shows this plot for seven selected states and the average movement across these states, revealing several important patterns. First, we see dramatic moves from the NW toward the SE quadrant in almost every state in

[11] Compactness also falls in the 2010s. Anecdotally, this is likely the result of increasingly extreme *partisan* gerrymanders during this decade, a factor not explored in this article.
[12] For example, Mississippi was 36.8% Black in the 1980 census. For six years during the 1980s, one of the state's representatives was Black, with no Black representatives the other four years. Thus, the delegation averaged 12% Black, and the racial proportionality of the Mississippi delegation in the 1980s is -24.8%. Among all Southern states' congressional maps from the 1970s through 2010s, the overall correlation between racial proportionality and Polsby-Popper compactness is -.28.

the 1990s. And more generally, we see the strength of the trade-off historically in the frequency of lines running along the NW-to-SE diagonal (and almost never along the SW-to-NE diagonal), suggesting that when changes to a map are made, states are usually sacrificing compactness to increase racial proportionality (or vice versa); they almost never simultaneously increase or decrease both measures. The moves in the 2000s and 2010s tend to be much more vertical than horizontal, with states maintaining similar racial representation while changing more on compactness. While compactness and Black representation were indeed heavily negatively correlated prior to the 2000s, this correlation has slightly abated in the last two decades, as several maps have alternately improved and worsened compactness without significant impact on racial composition. Also note that only three points on this graph lie to the right of zero on the x-axis, reflecting that the *only* state during this period in which Blacks were *not* underrepresented was Georgia in the 1990s through 2010s.

Of course, just because previous maps have shown an adverse relationship between compactness and Black representation does not mean that this trade-off is necessary for present or future districting. Plaintiffs in cases like *Milligan* assert that Black representation can be improved by reducing racial packing while maintaining compactness. Indeed, our algorithm finds many circumstances in which less-packed Black opportunity districts improve expected Black representation, but as we explore in the next section, the degree to which compactness must be sacrificed varies.

V. Algorithmically Generated District Results

Algorithm Specifications

We now move from analysis of historical enacted maps to maps generated using the hill-climbing method described in Section III. Using our heuristic based optimization, we generate 100 maps for each of six objective function combinations on each of nine states (a total of 600 maps per state). Each map is the result of running the algorithm for 25,000 steps, with each step defined as involving a strict improvement to the incumbent map.[13] Each objective function is a mix of maximizing Polsby-Popper compactness and expected Black representation in the following combinations:

- Objective 1.) 100% compactness
- Objective 2.) 50% compactness, 50% Black representation
- Objective 3.) 25% compactness, 75% Black representation
- Objective 4.) 10% compactness, 90% Black representation
- Objective 5.) 5% compactness, 95% Black representation
- Objective 6.) 100% Black representation

The weights on the components of the objective function do not correspond directly to prioritization of the component objectives since the different components of the objective function are in different units. Although both objectives are expressed as proportions between 0 and 1, the range of plausible outcomes tends to be much larger in the case of compactness than Black representation. As a result, compactness needs to be assigned a much smaller weight to be given equal practical influence on the resulting maps. For this reason, four of the six objective functions chosen overweight the Black representative objective, but this is just a consequence of the measures' scaling and not intended as a judgment on the normative value of each goal.

[13] Through a series of benchmarking trials, we determined a run length of 25,000 steps provided the optimal balance for the efficient generation of maps; see Appendix B for discussion.

Alabama Example

We begin by using the motivating example of Alabama. Figure 5 shows the 600 maps generated for Alabama along these dimensions, along with averages for all 100 runs of each objective functions, averages for the top 20% of runs of each objective, and the actual map enacted for the 2022 election (and subject of the *Milligan* litigation). The blue lines connecting the average values form a kind of approximate "Pareto curve" or "Pareto frontier" indicating the effective maximum value for compactness that one can achieve in a state at any level of Black representation, and vice versa.[14] The steepness of these curves suggests the extent to which one value must serve as a trade-off for the other, or whether it is possible to maximize both simultaneously.

When the algorithm is set to consider *only* compactness excluding any other objectives (beyond continuity and equal population constraints that are always enforced), the resulting maps tend to produce districts with an average compactness of .43, and elect around .55 Black representatives (Objective 1; the red squares). When the algorithm is set to *only* consider Black representation, it produces maps electing around 1.22 Black representatives, but with compactness of only around .10 (Objective 6; the green circles), though a subset of these maps produce close to 1.5 Black members.

[Figure 5 about here]

[Figure 6 about here]

[14] Formally, a Pareto frontier is the subset of all solutions (measured on two or more measurements) in which it is not possible to find an alternate solution which improves on one measurement without simultaneously decreasing another measurement. A Pareto curve is a visual depiction of the Pareto frontier set along two or more dimensions. However, we use the term more casually to describe the curves showing the approximate 50^{th} and 90^{th} percentile results of each objective function in Figures 5 and 7. As described in Section VI, the shape of these curves can be used to evaluate the extent to which compactness and Black representation are mutually exclusive in a given state. See Altman and McDonald (2018) for another use of Pareto curves to evaluate trade-offs in redistricting norms.

There is clearly a large trade-off between these two norms when only one of the objectives is valued. However, it is striking that once the objective function is set to incorporate just a small consideration for compactness, we see significant compactness gains while sacrificing very little Black representation. Note the average values under Objective 4 (90% Black representation, 10% compactness; the yellow triangles). Compared with Objective 6, average compactness increases from .10 to .31, while average Black representation declines only from 1.22 to 1.11. In other words, it is possible to achieve a map that is almost as compact as the ideally compact map while only sacrificing a 10% chance of electing one more Black representative.

The trade-off between these two dimensions (or lack thereof) is visually striking just by looking at examples of some maps. Figure 6 shows the best map produced under each of Objectives 1, 3, 4, and 6, alongside the 2022 enacted map. The Alabama maps drawn solely to maximize one objective perform largely as expected. The Objective 1 compactness-maximizing map indeed draws very regular, aesthetically pleasing districts, but represents minority voters poorly, yielding no districts with Black population greater than 35%, and usually electing zero Black representatives. The Objective 6 Black representation-maximizing map is likely to elect two Black representatives, but does so in districts that are extremely irregular. But the Objective 3 and 4 maps, which mix both objectives, perform better than the enacted map on *both* measures of compactness and Black representation. Objective 3 heavily packs a single Black majority district similarly to the enacted map, but with significantly more compact districts overall. Objective 4 instead draws two opportunity districts with ~50% Black population, creating Black representation close to the Objective 6 map, but still more compact than the enacted map. Thus, it does appear that while there is a slight trade-off necessary between these two dimensions, they are very far from perfect substitutes. Instead, we can come close to maximizing both goals simultaneously if we are willing to make minimal sacrifices on both ends.

Hill-Climbing Algorithm Results in Nine States

Generalizing from the single state example, Figure 7 below replicates the Figure 5 graph of Alabama for nine states, showing both the values of every map generated by the algorithm and the associated approximate Pareto curves. Figure 8 depicts the best and median maps for each of the six objective functions for each state, alongside the enacted map used in the 2022 election.[15] Table 1 shows the average and top quintile number of expected Black representatives produced for each state under each objective function, along with a few other reference points: the expected number of Black representatives elected under the 2022 enact maps, the number required for proportional representation, and a "non-spatial optimum", the maximum expected Black representatives possible without any spatial or contiguity constraints.[16] This table also summarizes how well each objective function compares to proportional representation and the pure Black representation objective (Objective 6) in expected Black members, and sums up the results across all nine states.[17]

---

[15] "Best" and "median" are defined by ranking the 100 maps generated for each objective function by their objective function value.

[16] The "Non-Spatial Optimum" is meant to provide a theoretically (but usually wildly unrealistic) absolute ceiling on Black representation in the state. Its bounds were computed using the following optimization model:

$$\begin{aligned}
\text{Maximize} \quad & \textstyle\sum_{j\in J} \phi\left(\frac{y_j}{z_j}\right) & \\
\text{s.t.} \quad & \textstyle\sum_{j\in J} y_j = B & \\
\textstyle\sum_{j\in J} z_j \quad & = V & \\
\textstyle\sum_{j\in J} \zeta_j \quad & = N & \\
z_j + \zeta_j \quad & = \tilde{P} & \forall j \in J \\
y_j, z_j, \zeta_j \geq 0 \quad & & \forall j \in J
\end{aligned}$$

where V, B, and N are the total number of voters, Black-voters, and non-voters in the region respectively and $\tilde{P} = \frac{V+N}{|J|}$ is the number of people that each district is meant to contain. Then y_j and z_j are the number of black voters and total voters in district j, respectively. ζ_j is the number of non-voters in district j. See [redacted] for details. These solutions were obtained by employing an optimization model that was implemented in the integer programming solver GUROBI.

[17] Across all nine states, 20.45 Black members would be elected out of 80 total seats if Black voters were represented proportionally. Our formula estimates than an average of 14.37 Black members would be elected under the maps enacted following the 2020 census; this compares closely to the 14.2 Black members actually elected on average during the 2010s (across 79 districts).

As with the historical data, Figure 7 reveals several similarities across the nine states. First, every state shows some trade-off between compactness and Black representation, while simultaneously showing some potential to significantly improve upon one dimension while sacrificing little along the second dimension. The approximate Pareto curve in each state follows roughly the same shape, and all the maps drawn with a predominantly compactness objective significantly under-represent the Black population. However, the actual enacted map in seven of the nine states falls well below the curves, suggesting that almost all existing maps could be substantially and simultaneously improved along both dimensions. But Figure 7 also reveals several differences across states. Most significantly, the steepness of the approximate Pareto curve does vary, indicating there is more of a trade-off between these dimensions in some states than others, with a steeper curve (more convex) indicating greater potential to simultaneously maximize both objectives. It appears steepest in Virginia and Georgia, where the Black population is primarily urban, suggesting this trait might make it easier to draw compact districts that promote Black representation, a factor further analyzed in the next section.

[Table 1 about here]

[Figure 7 about here]

[Figure 8(a-c) about here]

We can also generalize about the effects of each objective function from Table 1 and Figures 7 and 8.

1.) *Objective 1 (100% compactness):* This objective draws very regular shaped districts with overall average compactness of .41, much higher than any enacted Southern map in the recent past, but at a significant cost to Black representation. Together, these maps elect an average of 11 Black representatives across all nine states, well below the expected number of Black members elected in the enacted maps (14.4), and 46% below proportional representation. The typical Objective 1 map draws no Black opportunity districts in at least five of the nine states.

2.) *Objective 2 (50% compactness, 50% Black representation)*: Objective 2 also draws very compact districts (average Polsby-Popper=.40), in many cases visually indistinguishable from Objective 1, while noticeably improving upon Black representation, electing at least two more Black members across the nine states. But in many states, this is done by drawing multiple districts with Black populations in the 35-40% range rather than creating genuine Black-majority or Black opportunity districts, and these maps still average 36% below Black proportionality.

3.) *Objective 3 (25% compactness, 75% Black representation):* Objective 3 maps closely resemble the enacted maps with respect to Black representation, usually drawing a single Black-majority district in Alabama, South Carolina, and Tennessee, and multiple such districts in Maryland and Georgia. They elect almost 15 expected Black members across all states, just over the 14.4 expected from the enacted maps. Yet they are still dramatically more compact than any real southern congressional map (average Polsby-Popper=.37). Expected Black representation is 27% below proportionality and 11% below the performance of Objective 6.

4.) *Objective 4 (10% compactness, 90% Black representation)*: Here we observe some irregularly shaped districts, especially where the algorithm draws a second Black opportunity district in Alabama, Louisiana, and South Carolina. But most districts remain fairly compact, with an average Polsby-Popper of .29, still superior to most real southern congressional maps averaging .244. And this objective gets close to optimal realistic Black representation (4% below Objective 6), though still 23% short of proportionality.

5.) *Objective 5 (5% compactness, 95% Black representation):* We see several highly irregular districts in most Objective 5 maps, especially in central Alabama, coastal Louisiana, and southern Maryland. Average compactness is .23, slightly below the current average. However, Black representation under this objective falls just 1% below Objective 6. The *best quartile* of these maps elect 22% more Black members than the enacted maps, to within 14% of proportionality.

6.) *Objective 6 (100% Black representation):* Maps drawn with Black representation as their sole objective perform only slightly better than Objective 5 on this measure, but produce districts that are dramatically worse on compactness. Average Polsby-Popper scores hover barely above .10, substantially worse than almost all real enacted maps, even recent partisan gerrymanders, with visually bizarre districts in every state. There appears to be little benefit to this objective over the previous one, but a very substantial cost. Especially comparing this last objective function (6) to (4) and (5), it is obvious that a small amount of consideration for compactness can go a long way without significantly sacrificing other important values.

[Figure 9 about here]

Figure 9 summarizes the average results across all nine states for each objective function. Black representation (the solid line), increases substantially as the objective mix shifts from (1) to (4). But the gains as the mix shifts from (4) to (6) are extremely small. Conversely, the gains in compactness (the dashed line) are similarly dramatic as the objective shifts from (6) to (3), but almost imperceptible as the objective shifts from (3) to (1). Overall, there is indeed a trade-off between these two norms. But the trade-off seen in the middle two objectives involves making only a modest sacrifice in one norm while being almost costless for the second norm. Objective (3) sacrifices 12% of Black representation compared to the pure objective (6), but only .04 in average compactness compared to pure objective (1), while objective (4) sacrifices only 4% of Black representation but .12 in compactness compared to (6) and (1) respectively. And almost all of the maps created in most states from Objectives 3 through 5 simultaneously perform better on both metrics than the actual enacted maps following the 2020 census, and they are almost all substantially more compact than even the pre-*Gingles* maps from the 1970s and 1980s.

Comparison to Specific Enacted Maps

If we analyze the differences between the algorithmic maps that heavily weigh Black representation (Objectives 4, 5, and 6) and the enacted maps, a general pattern emerges with a few notable exceptions. *In seven of the nine states, at least one objective function draws maps that are both more compact and elect more Black representatives than the enacted map*, often by drawing more districts that are less racially concentrated than existing Black opportunity districts. Most commonly, in the smaller Deep South states where the enacted map packs a single Black-majority district, the algorithm often cracks this district to create two Black opportunity districts. For example, in Mississippi, the enacted map includes a single 62% Black district that we estimate is 97% likely to elect a Black representative, but no other district more than 10% likely to do so. By contrast, the Objectives 4-6 algorithm tends to draw two districts with 45-50% BVAP, each electing a Black representative 50-70% of the time. Alabama, South Carolina, and Louisiana see similar results, where again the best algorithmic maps draw two Black opportunity districts even if one or both such districts are slightly below a BVAP majority. Mississippi's results are also exceptional in that the enacted map produces fewer expected Black representatives than even the compactness-maximizing maps, which due to the dispersed Black population often include two weak Black-opportunity districts. This notably suggests that the actual map of Mississippi, with one super-majority Black district, possibly underrepresents Black voters compared even to a completely race-blind map.

The enacted maps in Maryland and Georgia already contain multiple Black majority districts. But the algorithm follows a similar strategy in cracking these districts to create more Black opportunities. In Maryland, the enacted map includes two Black majority districts around Baltimore and Prince George's County. But the best algorithmic maps split these to draw four districts with 42-50% BVAP stretching from Baltimore across the southern coast of the state. Georgia's enacted map includes three Black opportunity districts in metro Atlanta and a fourth around Savannah. The best maps drawn by the algorithm split these compact Atlanta districts to extend across pockets of Black population in the central and eastern parts of the state, with BVAPs tightly clustered in the 46-53% range. For example, the "best" Objective 5 and 6 Georgia maps shown in Figure 7 draw *six* Black opportunity districts, including one on the eastern coast completely absent in the enacted map.

There are notably two states where the algorithm does *not* perform substantially better than the enacted map on Black representation, though for very different reasons. In Tennessee, both the enacted maps and the simulated maps under all objectives almost always draw a single strong Black-majority district around Memphis. This one district is drawn even if compactness is the sole objective. And even if Black representation is the sole objective, a second Black opportunity district appears impossible; the difference between the average Objective 1 and Objective 6 map is only 0.15 Black members.

By contrast, the Objective 4 through 6 maps in Virginia usually draw two opportunity districts, both in the southeast portion the state, most often respectively centered around Richmond and Norfolk. These are mostly "weaker" Black opportunities with 40-45% BVAP, yielding expected Black probabilities of 50-70%. Yet uniquely, this split of the Black population is exactly what the enacted map does, drawing two districts with BVAPs of 45% and 42% in the same regions. In Figure 8, Virginia is the sole state where the enacted map does not fall well below the "Pareto curve", instead sitting highly among the best algorithmic maps on both dimensions. Virginia was forced create a second Black opportunity district in 2016 (in *Whittman v. Personhuballah)*, which the Virginia Supreme Court largely retained in drawing the 2022 map, while substantially improving overall compactness. Despite the fact that these two districts do not have Black majorities, in practice they are heavily Democratic and seem likely to elect Black representatives for the foreseeable future. Due to a combination of federal court intervention and a new bipartisan redistricting process, it appears that Virginia alone has drawn a close to optimal map on both dimensions.

VI. Defining the Trade-Off in Each State

Although the sample of nine states explored in this paper is quite small, we can still attempt to quantify the degree to which a trade-off between our two normative goals is necessary in each state, and suggest common factors that may be related to this trade-off. The varied shapes of the distributions of the maps generated across the nine states suggest significant variation in the necessity to compromise compactness to gain additional Black representation (and vice versa). We can define a "trade-off ratio" to measure the severity of this sacrifice by contrasting the shapes of these distributions and the approximate Pareto curves generated by their average points.

This trade-off ratio is defined through the following steps:

1.) Define the approximate Pareto curve by the means for each set of 100 points generated from each objective, with the ends drawn from Objectives 1 and 6.[18]

2.) Draw a right triangle with the hypotenuse defined by the ends of the approximate Pareto curve.

3.) Take the ratio of the area above the approximate Pareto curve contained within the triangle to the area of the whole triangle. This number will be between 0 and 1, with larger number indicating a more severe trade-off. A ratio of 1 indicates you can only increase one objective by proportionally decreasing the other objective. A ratio of 0 means you can jointly maximize both objectives simultaneously without any cost to the other.

Figure 10 gives an example of this method applied to the Alabama graph from Figure 5.

[Figure 10 about here]

[Figure 11 about here]

Figure 11 shows the trade-off ratio for each state shown in Figure 7 plotted against both statewide BVAP, and the percent of the state's Black population living in rural counties, with data

[18] The curve can also be defined using the means of the top quintile of runs, or some other subset, with substantively very similar results.

drawn both from the set of all runs and from the top quintile. Every state has a ratio below 0.5, indicating that compactness and Black representation are far from perfect substitutes, and that it is possible to make large improvements along one dimension with relatively modest sacrifices along the other. But a clear pattern emerges, with the trade-off proving much stronger in states with a more rural Black population (Alabama and South Carolina) than in states where the Black population is more urban (Georgia and Tennessee), and thus Black influence districts can be drawn more compactly centered around large cities. This is most easily seen in the single Black-majority district comprising Memphis that emerges even in the purely compactness-maximizing maps of Tennessee. But beyond the obvious fact that urban population centers are likely themselves more compact than rural areas, Blacks in urban areas may also be more likely to live in proximity to white progressives, Latinos, and other non-Black groups more amenable to voting for a Black Democrat than rural Southern white voters. Previous research across several decades (e.g. Hood and McKee 2023; Voss 1996) has shown urban Southern white voters to be less racially polarized than rural whites. This factor likely contributes especially to the great number of Black opportunity districts our simulations find feasible in the Atlanta and DC metro areas.

The effects of both Black population and urbanization on this trade-off are significant at $p<.05$ in simple bivariate regressions, though their effects are not consistently independently significant in models incorporating both variables (full results in Appendix E). It appears that the one largely non-Southern state in the data set, Maryland, is an outlier in the bottom two graphs, and the correlation between trade-off ratio and rural Black population increases substantially if Maryland is excluded. Despite having a very urban Black population, Maryland's Black-representation maximizing maps are very noncompact, largely owing to the unique geography of the state, as the algorithm often draws such districts to awkwardly snake across the Chesapeake Bay. Nevertheless, in determining the necessity for non-compact districts to represent minority voters in a state, practitioners and courts may be well-advised to consider the urbanization of that population as a crucial factor, with more rural populations permitting less rigorous adherence to compactness goals.

VII. Discussion and Conclusion

The analysis of both real and approximately optimized districts in this article is contained to a single narrow yet substantively important question in legislative districting: the role of the compactness norm in potentially suppressing the representation of Black voters in southern states where they have been historically and systematically underrepresented. We find that although the trade-off is real at the extremes, it is likely mostly unnecessary to drastically sacrifice compactness to achieve close to optimal or proportional Black representation. With respect to the specific controversies in the current redistricting cycle, we find that it is indeed possible to draw many configurations of maps with *two* districts likely to elect Black representatives in Alabama, Louisiana, and South Carolina while simultaneously *improving on the compactness* of the existing maps.

Nevertheless, there are certainly instances where the simplifications required to model district outcomes may lead to imperfect conclusions. Although our model does account for differences across states with a subregional interaction, many other across- or within-state factors may contribute to differences in Black representation. For example, variation in the proportion of White liberals between suburban Atlanta and rural northern Georgia may lead to significant differences in the likelihood of neighboring districts to elect a minority even if the minority population were constant. Even with the subregional interaction, factors like this are likely responsible for our algorithm seeming to under-predict Black representation in Virginia and over-predict it in Mississippi.

The narrow specifications of this paper also do not attempt to model the representation of other minority groups, especially Latinos. In many states, other minorities both merit large consideration for their own representation and how they influence the representation of Black voters. While this paper deliberately excludes states that could feasibly draw Latino influence districts, both the inclusion of a

Latino representation objective and the impact of other minority populations are important avenues for future research. Finally, the paper models only one measure of compactness, to the exclusion of other "traditional districting principles" such as respect for existing political boundaries and core retention.

This method also allows us to explore the potential trade-off between *descriptive* racial representation and *substantive* partisan representation that has been at the heart of much of the literature on racial gerrymandering since the 1990s, beginning with works like Lublin (1997) and Cameron et al. (1996). While detailed analysis of this subject is beyond the scope of this article,[19] Appendix D shows the correlation between expected Black seats and expected Democratic seats (estimated as a probit function of Cook's PVI) for all the simulations runs used here. Generally, Black representation and Democratic representation is found to be highly *positively* correlated in all Deep South states, but negatively correlated in only Maryland and Virginia, where coalition districts between Black and White liberal voters are much more feasible.

This article represents just one application of the general algorithmic approach to approximately optimized districts, albeit an application of immediate importance to scholars and practitioners of constitutional law, legislative studies, and race and ethnic politics. But the methodology advanced here can also easily extend to research exploring such dimensions as partisan goals, communities of interest, and existing political boundaries in the future, as well as opening up this exploration for the public in forthcoming redistricting cycles.

[19] We explore trade-offs in partisan gerrymandering using this method in other work (citation redacted).

References

Altman, Micah, and Michael McDonald. 2011. "BARD: Better automated redistricting." *Journal of Statistical Software* 42.1: 1-28.

Altman, Micah, and Michael McDonald. 2018. "Redistricting by formula: An Ohio reform experiment." *American Politics Research* 46.: 103-131.

Becker, Amariah, et al. 2021. "Computational redistricting and the voting rights act." *Election Law Journal: Rules, Politics, and Policy*.

Becker, Amariah, and Justin Solomon. 2021. "Redistricting Algorithms." In *Political Geometry*, eds. M. Duchin and O. Walsh. Basel: Birkhauser. Under contract.

Borodin, Allan, et al. "Little House (Seat) on the Prairie: Compactness, Gerrymandering, and Population Distribution." *AAMAS*. 2022.

Bowen, Daniel. 2014. "Boundaries, redistricting criteria, and representation in the US house of representatives." *American Politics Research* 42.5 (2014): 856-895.

Buchanan, Austin. "Political Districting." 2023. Optimization Online. URL: https://optimization-online.org/wp-content/uploads/2023/06/political_districting.pdf.

Butler, David, and Bruce Cain. 1992. *Congressional Redistricting: Comparative and Theoretical Perspectives*. New York: Macmillan.

Cameron, Charles, David Epstein and Sharyn O'Halloran. 1996. “Do Majority-Minority Districts Maximize Substantive Black Representation in Congress?”. *American Political Science Review*, 90(4):794-812.

Canon, David. 2022. "Race and Redistricting." *Annual Review of Political Science* 25: 24.1-24.20.

Casellas, Jason. 2011. *Latino Representation in State Houses and Congress*. Cambridge: Cambridge University.

Chen, Jowei. 2017. "The impact of political geography on Wisconsin redistricting: An analysis of Wisconsin's Act 43 Assembly Districting Plan." *Election Law Journal* 16(4): 443-452.

Chen, Jowei and Jonathan Rodden. 2015. ‘‘Cutting through the Thicket: Redistricting Simulations and the Detection of Partisan Gerrymanders.’’ *Election Law Journal* 14(4): 331– 345.

Cirincione, Carmen, Thomas Darling, and Timothy O'Rourke. 2000. "Assessing South Carolina's 1990s congressional districting." *Political Geography* 19.2: 189-211.

Duchin, Moon. Metric Geometry and Gerrymandering Group. 2018. “Comparison of Districting Plans for the Virginia House of Delegates”. Available at: https://mggg.org/VA-report.pdf.

Epstein, David, and Sharyn O’Halloran. 1999b. "Measuring the Electoral and Policy Impact of Majority-Minority Voting Districts.” *American Journal of Political Science* 43(2): 367-95.

Fabio, Anthony et al. 2009. “The Association Between County-Level Injury Rates and Racial Segregation Revisited: A Multilevel Analysis.” *Amer. J. of Public Health*. 99(4): 748-753.

Fravel, Jamie, Robert Hildebrand, Nicholas Goedert, Laurel Travis, and Matthew Pierson. “Dual Bounds for Redistricting Problems with Non-Linear Objective Functions.” Manuscript. 2022.

Garfinkel, Robert, and George Nemhauser. 1970. "Optimal political districting by implicit enumeration techniques." *Management Science* 16.8: B-495.

Herschlag, Gregroy, Robert Ravier, and Jonathan Mattingly, 2017. “Evaluating Partisan Gerrymandering in Wisconsin.” Available at: https://arxiv.org/pdf/1709.01596.pdf

Hayes, Danny, and Seth McKee. 2009. "The participatory effects of redistricting." *American Journal of Political Science* 53.4 (2009): 1006-1023.

Hicks, William, et al. 2017. "Revisiting Majority-Minority Districts and Black Representation." *Political Research Quarterly*: 1065912917738574.

Hood III, M. V., and Seth C. McKee. 2022. *Rural Republican Realignment in the Modern South: The Untold Story*. Univ of South Carolina Press.

Kaufman, Aaron, Gary King, and Mayya Komisarchik. 2021. "How to measure legislative district compactness if you only know it when you see it." *American Journal of Political Science* 65.3: 533-550.

Ladewig, Jeffrey W. 2018. "“Appearances do matter”: Congressional district compactness and electoral turnout." *Election Law Journal* 17.2: 137-150.

Lublin, David, Handley, Lisa, Brunell, Thomas, and Grofman, Bernard. 2020. “Minority success in non-majority minority districts: Finding the “sweet spot”.” *Journal of Race, Ethnicity and Politics,* 5(2): 275-298.

Lublin, David. 1997. *The Paradox of Representation: Racial Gerrymandering and Minority Interests in Congress*. Princeton University Press.

Massey, Douglas and Nancy Denton. 2008. “The Dimensions of Racial Segregation.” *Social Forces* 67(2): 281-315.

McDonald, Michael, and Robin E. Best. 2015. "Unfair partisan gerrymanders in politics and law: A diagnostic applied to six cases." *Election Law Journal* 14.4: 312-330.

McDonald, Michael, and Daniel Magleby, Jonathan Krasno, Shawn Donahue, and Robin Best. 2018. “A Proposal to Give Seat Effect Second Billing when Investigating Gerrymandering”. Presented at the 2018 Annual Conference of the MPSA, Chicago, IL.

McKee, Seth, and Melanie Springer. 2015. "A tale of “Two Souths”: White voting behavior in contemporary Southern elections." *Social Science Quarterly* 96.2: 588-607.

Ruszczyński, Andrzej. 2006. *Nonlinear Optimization.* Princeton, NJ: Princeton University.

Saxon, James. 2020. Reviving Legislative Avenues for Gerrymandering Reform with a Flexible, Automated Tool. *Political Analysis*. 1-23.

Stephanopoulos, Nicholas, and Eric McGhee. 2015. “Partisan Gerrymandering and the Efficiency Gap.” *University of Chicago Law Review* 82: 831-900.

Swamy, Rahul, Douglas King, and Sheldon Jacobson. 2022. "Multi-Objective Optimization for Political Districting: A Scalable Multilevel Approach." *Operations Research* (preprint).

Tam Cho, Wendy, and Yan Liu. 2016. "Toward a talismanic redistricting tool: A computational method for identifying extreme redistricting plans." *Election Law Journal* 15.4: 351-366.

Validi, Hamidreza and Buchanan, Austin. 2022. “Political districting to minimize cut edges.” Mathematical Programming Computation. 14: 623 – 672.

Voss, Stephen. 1996. "Beyond racial threat: Failure of an old hypothesis in the new South." *The Journal of Politics* 58.4: 1156-1170.

Cases Cited

Abrams v. Johnson, 521 U.S. 74 (1997)

Allen v. Milligan, 599 U. S. 1 (2023)

Bartlett v. Strickland, 556 U.S. 1 (2009)

Bethune Hill v. Virginia State Board of Elections, 326 F. Supp. 3d 128 (E.D. Va. 2018)

Bethune Hill v. Virginia State Board of Elections, 137 S. Ct. 788 (2017)

Cooper v. Harris, 136 S. Ct. 2512 (2017)

Common Cause v. Lewis, 834 S.E.2d 425 (N.C. 2019)

Hunt v. Cromartie, 526 U.S. 541 (1999)

Karcher v. Daggett, 462 U.S. 725 (1983)

League of Women Voters vs. Commonwealth of Pennsylvania, 178 A.3d 737 (Pa. 2018)

Merrill v. Milligan, 595 U.S. ____, Nos. 21A375 (21–1086) (2022)

Miller v. Johnson, 515 U.S. 900 (1995)

Shaw v. Reno, 509 U.S. 630 (1993)

Shelby County v. Holder, 570 U.S. 529 (2013)

Singleton v. Merrill, 2:21-cv-01291 (N. D. Ala. 2022)

Thornburg v. Gingles, 478 U.S. 30 (1986)

Vieth v. Pennsylvania, 195 F.Supp.2d 672 (M. D. Pa. 2002)

Whittman v. Personhuballah, 136 S. Ct. 1732 (2016)

Table 1. Black Representation Results for Algorithmic Districting in Nine States

	AL	GA	LA	MD	MS	NC	SC	TN	VA	9-State Total
Average										
Obj. 1 (100% Compactness)	0.55	3.34	0.75	2.29	1.15	0.87	0.42	0.91	0.74	11.01
Obj. 2 (50% Black Rep)	0.81	3.88	0.99	2.61	1.19	1.15	0.59	0.96	0.96	13.14
Obj. 3 (75% Black Rep.)	0.99	4.21	1.21	2.81	1.23	1.51	0.78	0.98	1.14	14.86
Obj. 4 (90% Black Rep.)	1.11	4.47	1.38	2.96	1.30	1.84	0.90	1.02	1.27	16.24
Obj. 5 (95% Black Rep.)	1.15	4.55	1.43	3.01	1.34	1.95	0.96	1.05	1.30	16.75
Obj. 6 (100% Black Rep.)	1.22	4.57	1.44	3.03	1.32	1.95	1.00	1.06	1.32	16.91
Top 20%										
Obj. 1 (100% Compactness)	0.54	3.33	0.75	2.42	1.15	0.86	0.39	0.91	0.76	11.11
Obj. 2 (50% Black Rep)	0.89	3.92	1.11	2.67	1.20	1.19	0.62	0.96	1.05	13.59
Obj. 3 (75% Black Rep.)	1.05	4.32	1.31	2.91	1.24	1.62	0.82	0.98	1.20	15.45
Obj. 4 (90% Black Rep.)	1.26	4.62	1.45	3.08	1.36	1.96	0.96	1.04	1.35	17.08
Obj. 5 (95% Black Rep.)	1.31	4.69	1.52	3.11	1.39	2.05	1.04	1.06	1.38	17.54
Obj. 6 (100% Black Rep.)	1.42	4.71	1.55	3.14	1.38	2.03	1.08	1.07	1.40	17.78
2020s Enacted Expected	1.00	3.92	1.25	2.58	1.12	1.39	0.72	1.00	1.39	14.37
Proportional Rep.	1.81	4.44	1.88	2.48	1.45	2.99	1.77	1.45	2.19	20.45
Proportionality of Enacted	55%	88%	67%	104%	77%	46%	41%	69%	64%	70%
Non-Spatial Optimum	2.98	8.19	2.99	4.47	2.00	5.51	2.94	2.79	4.79	36.66
Avg. % of Proportionality										
Obj. 1 (100% Compactness)	30%	75%	40%	92%	79%	29%	24%	63%	34%	53.9%
Obj. 2 (50% Black Rep)	45%	87%	53%	105%	82%	38%	34%	66%	44%	64.2%
Obj. 3 (75% Black Rep.)	54%	95%	64%	113%	85%	50%	44%	68%	52%	72.6%
Obj. 4 (90% Black Rep.)	61%	101%	73%	119%	90%	61%	51%	71%	58%	79.4%
Obj. 5 (95% Black Rep.)	64%	103%	76%	121%	93%	65%	54%	72%	60%	81.9%
Obj. 6 (100% Black Rep.)	67%	103%	77%	122%	92%	65%	57%	73%	60%	82.7%
Avg. % of Objective 6										
Obj. 1 (100% Compactness)	45%	73%	52%	76%	87%	45%	41%	86%	56%	65.1%
Obj. 2 (50% Black Rep)	66%	85%	69%	86%	90%	59%	59%	91%	73%	77.7%
Obj. 3 (75% Black Rep.)	81%	92%	84%	93%	93%	78%	77%	93%	87%	87.8%
Obj. 4 (90% Black Rep.)	91%	98%	96%	98%	98%	94%	90%	96%	96%	96.0%
Obj. 5 (95% Black Rep.)	95%	100%	99%	100%	101%	100%	96%	99%	99%	99.1%

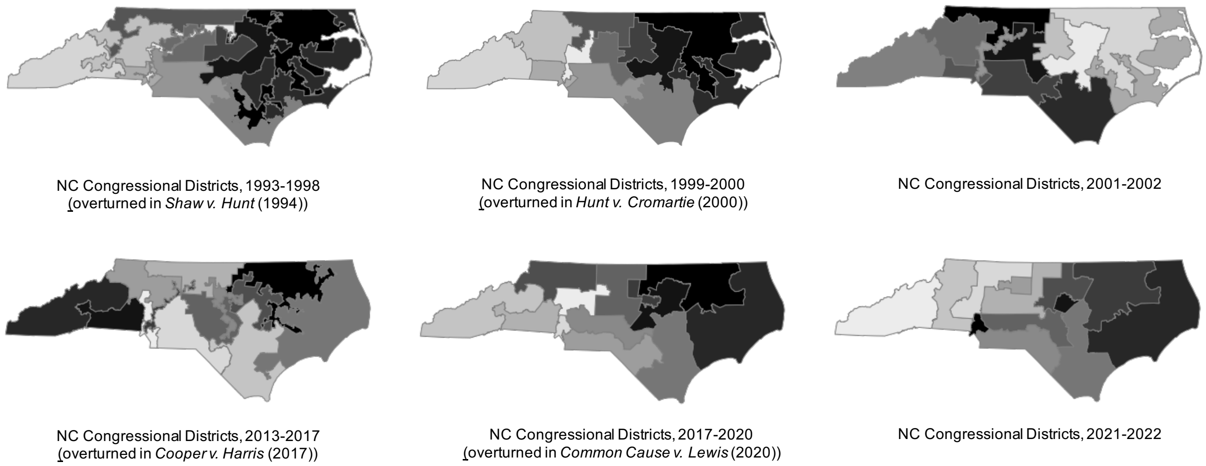

Figure 1. North Carolina Congressional Districts in the 1990s and 2010s

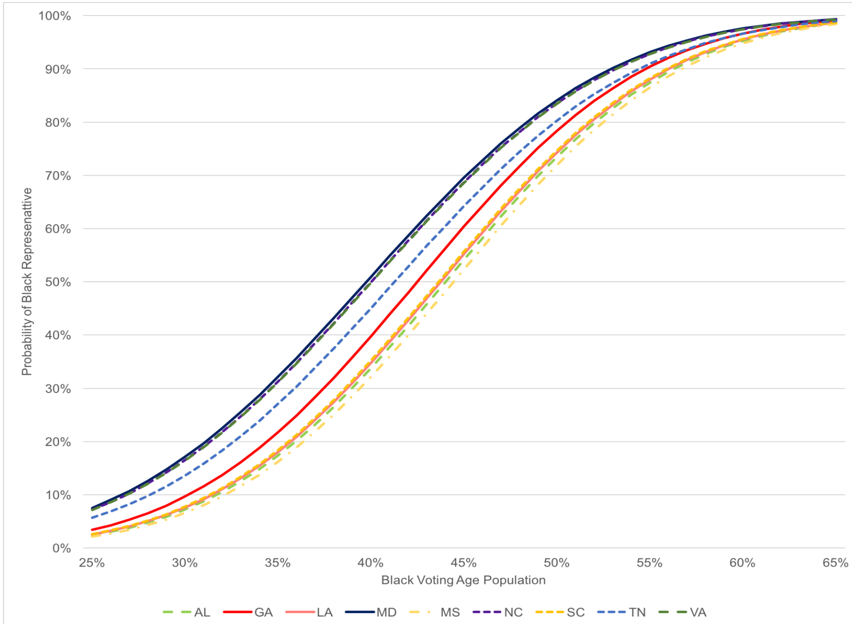

Figure 2. Probability of Election of Black Representative as a Function of Black Population (Probit Estimate, statewide average Latino population imputed)

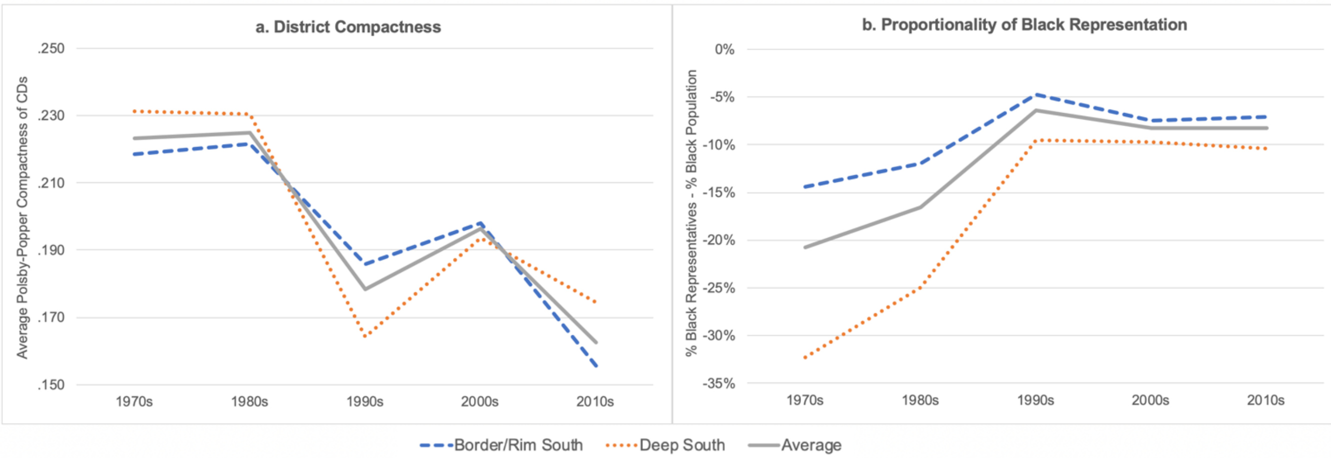

Figure 3. Compactness and Proportionality of Black Representation within Congressional Districts by Decade

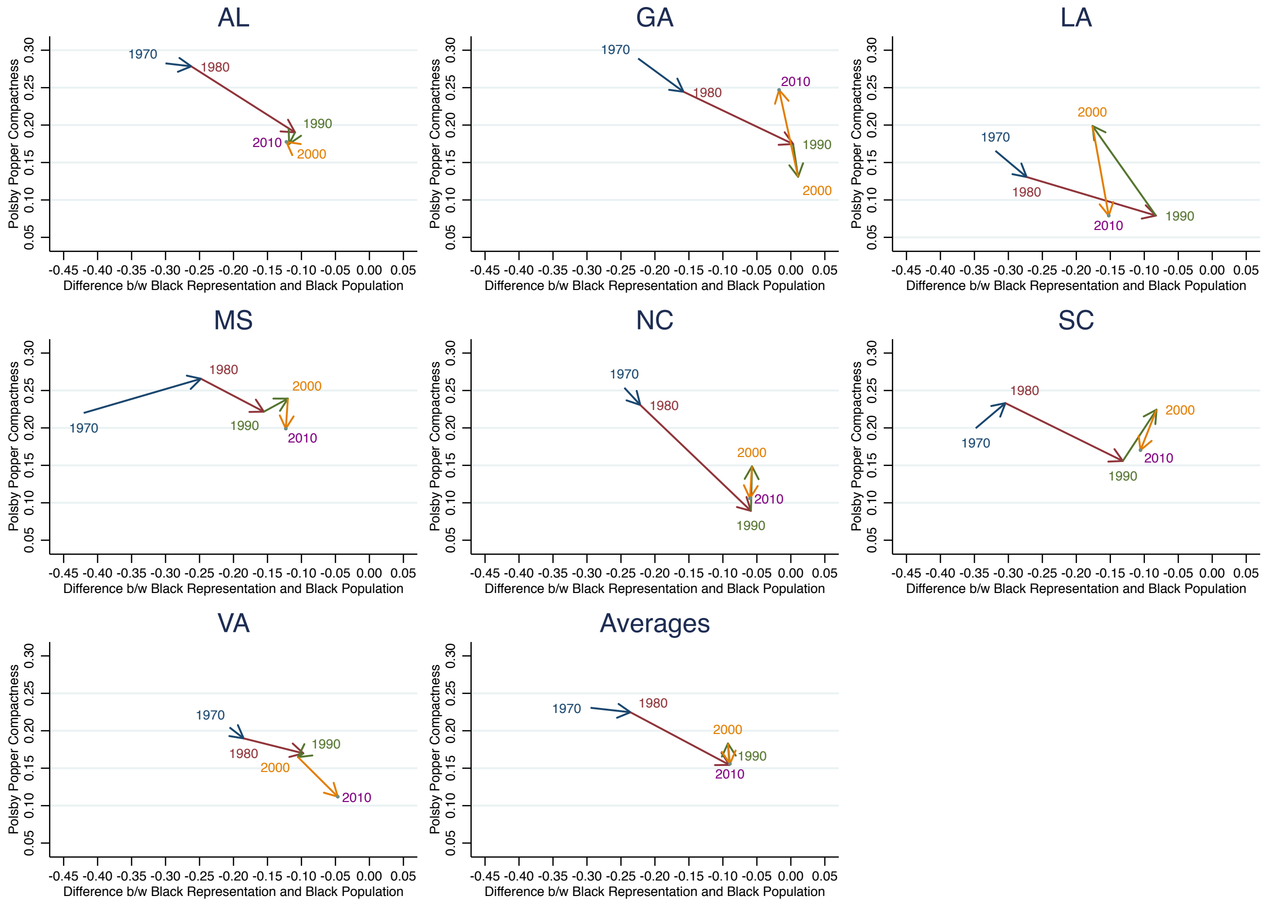

Figure 4. Black Representation Proportionality vs. Compactness in Select Southern States

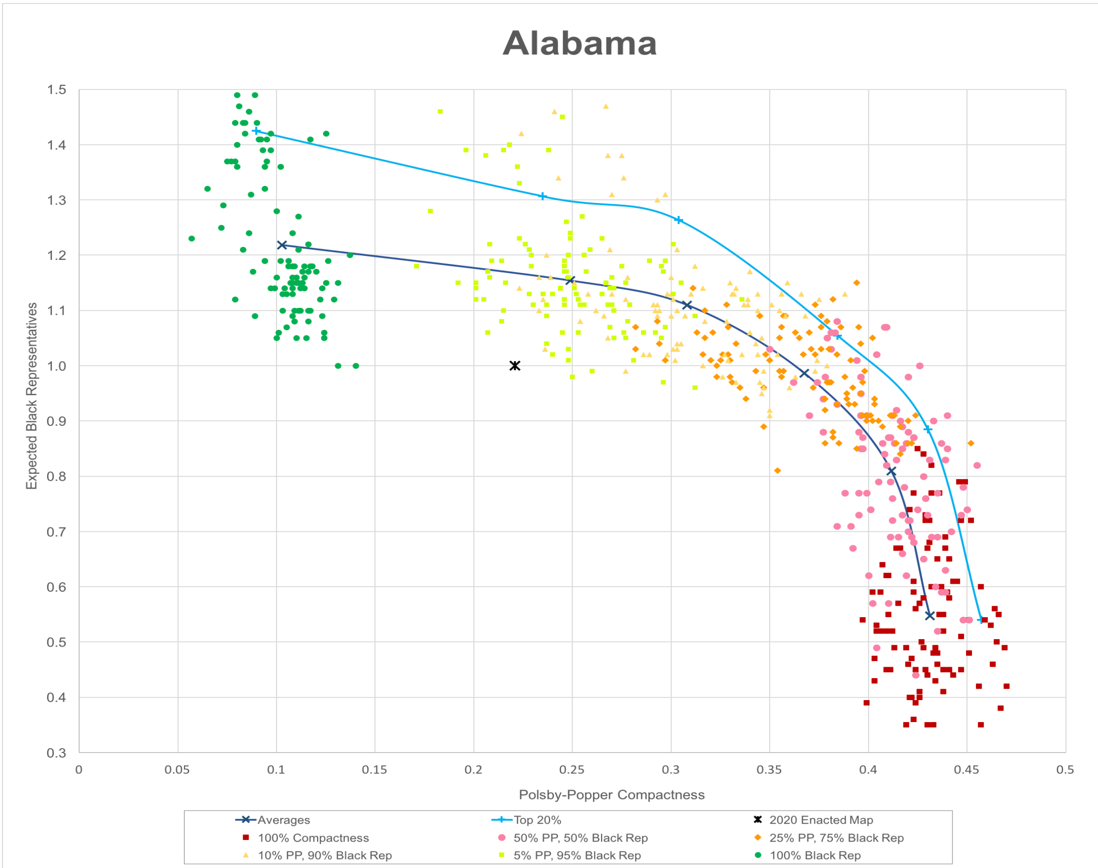

Figure 5. Hill-Climbing Algorithm Results in Alabama

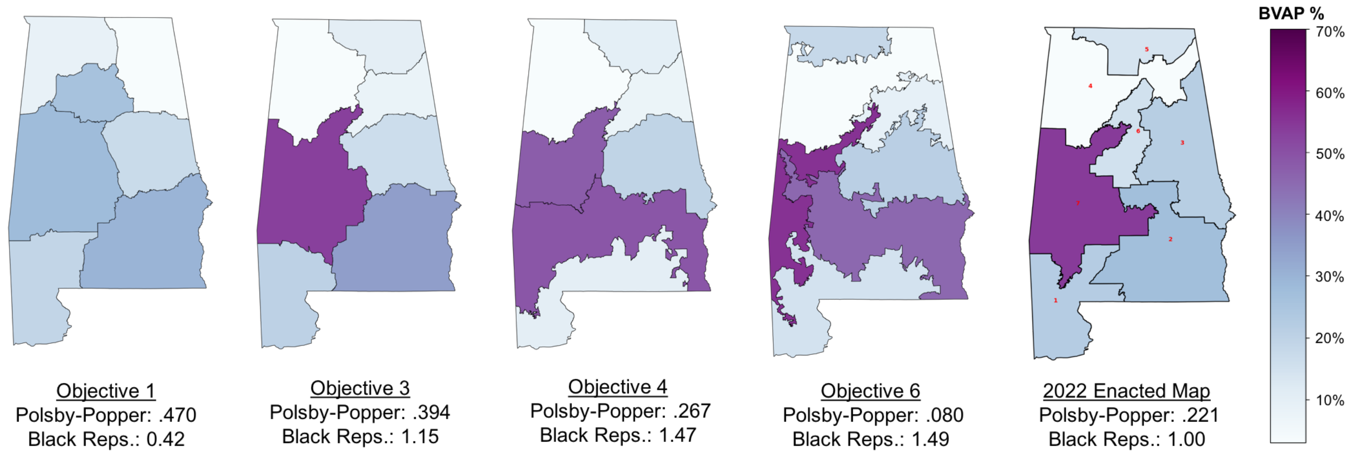

Figure 6. Four Algorithmic Maps of Alabama Compared with Enacted Map

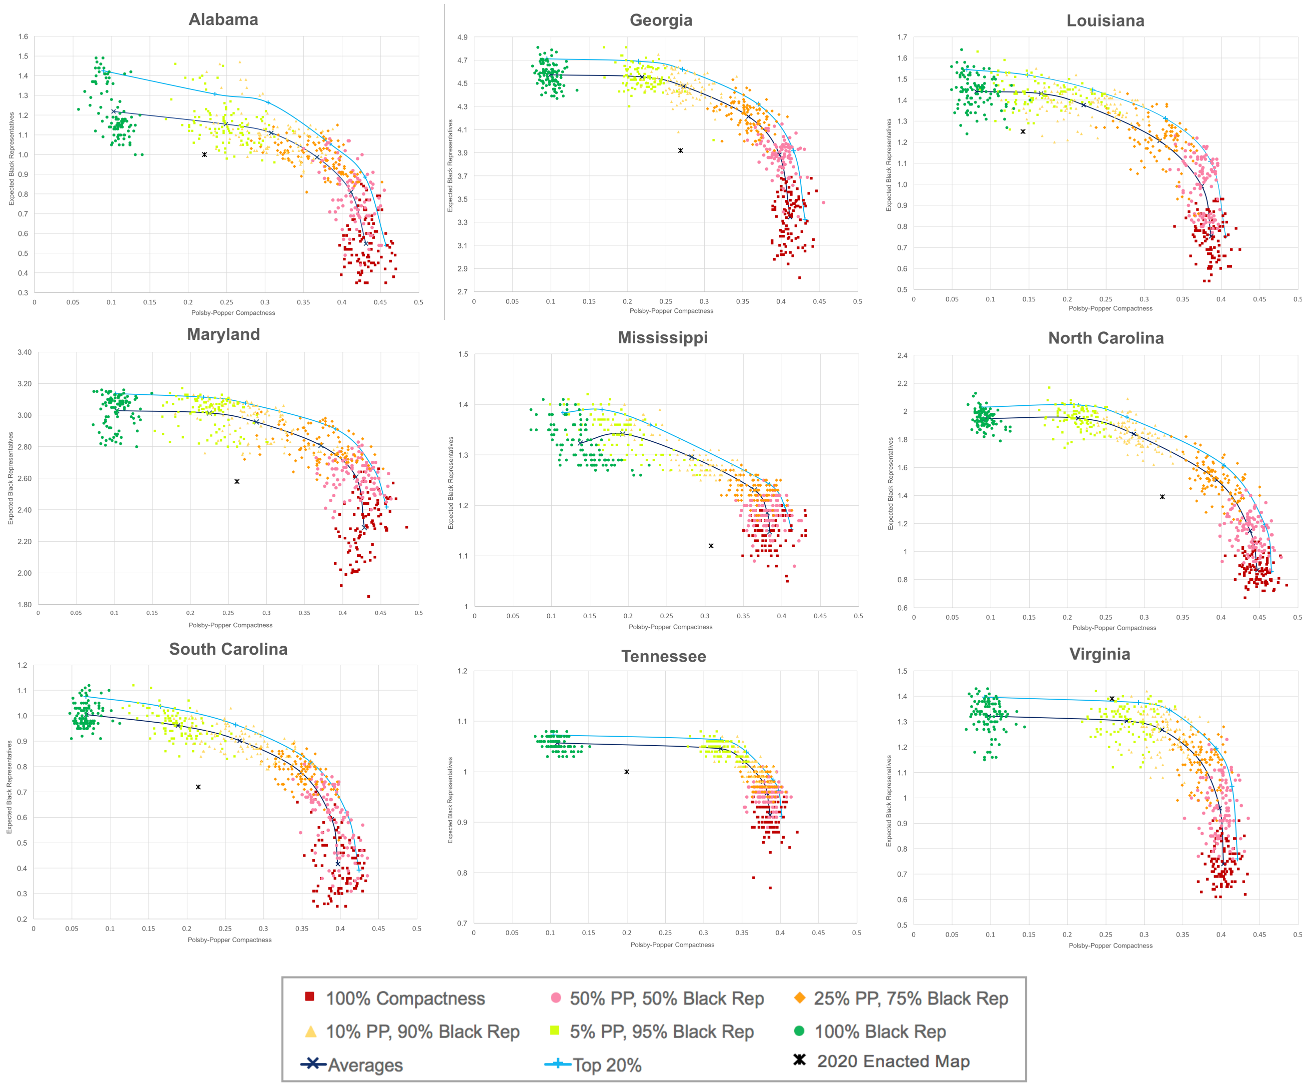

Figure 7. Hill-Climbing Algorithm Results in Nine States

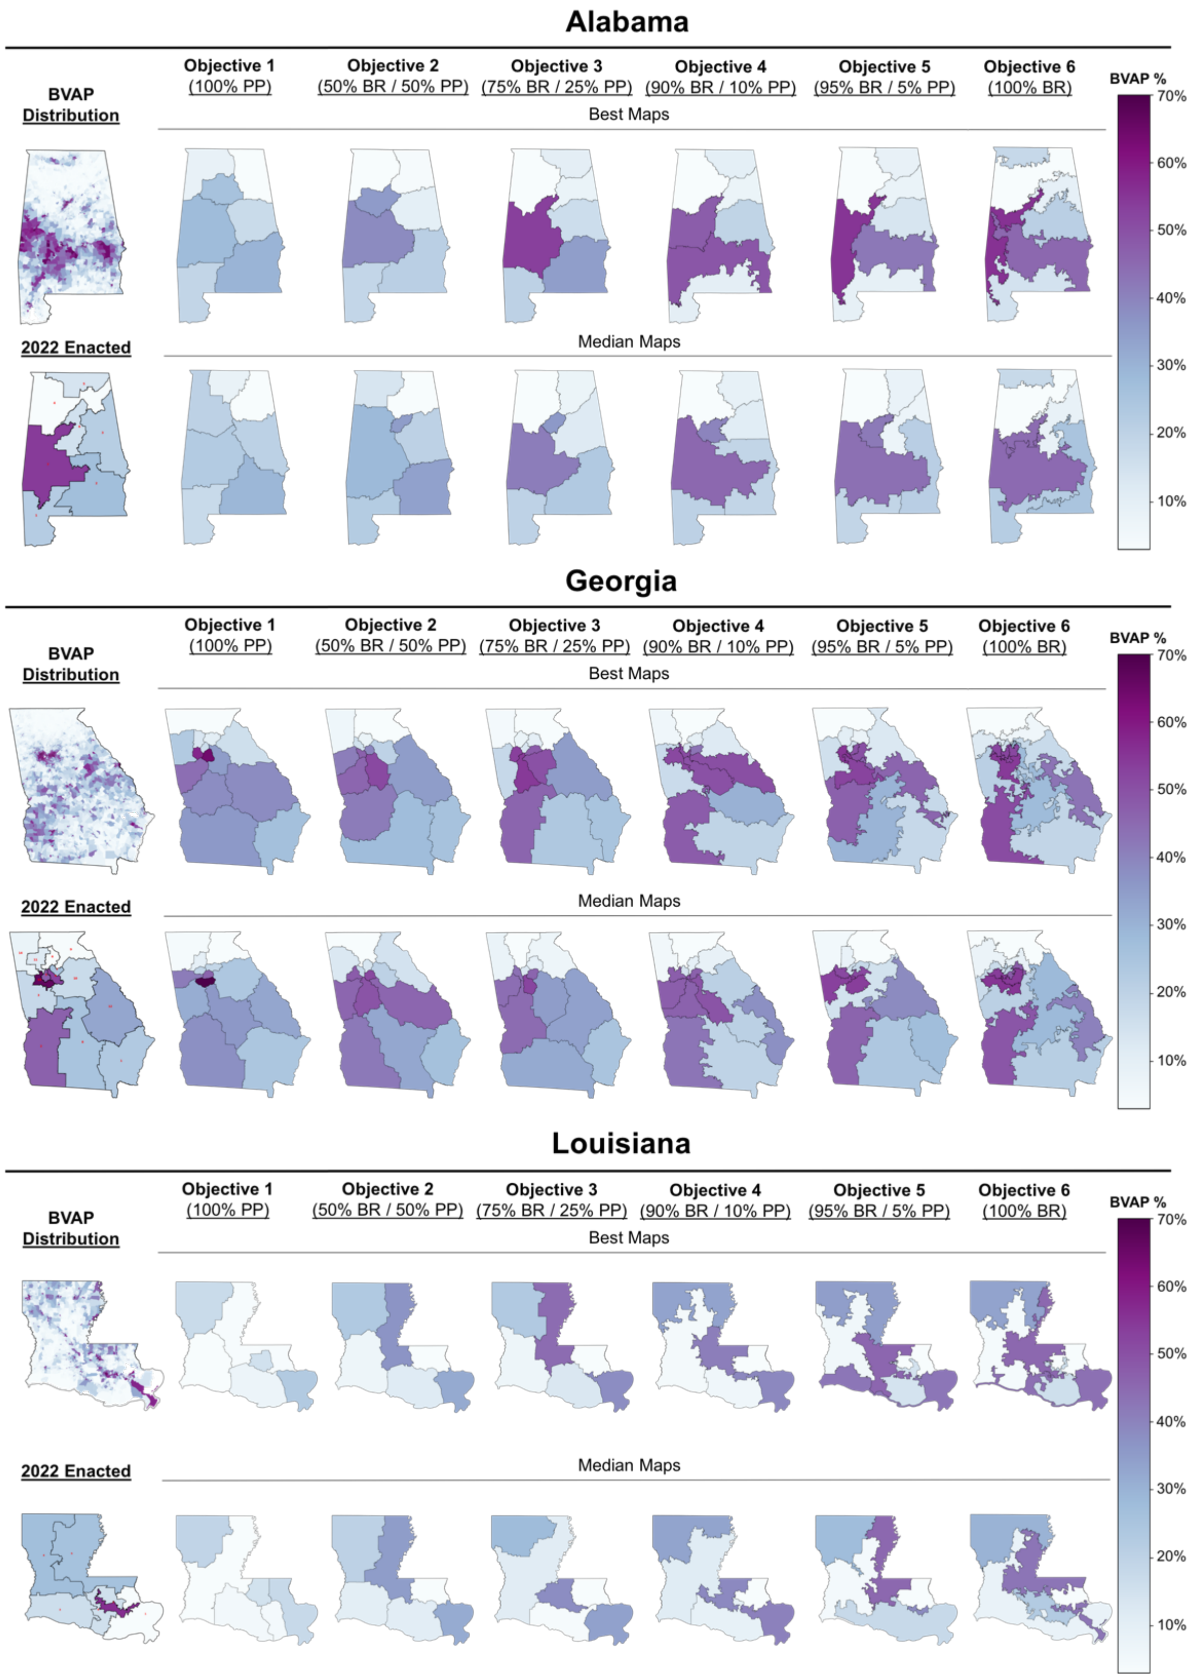

Figure 8a. Maps Generated in Nine States from Six Objective Functions

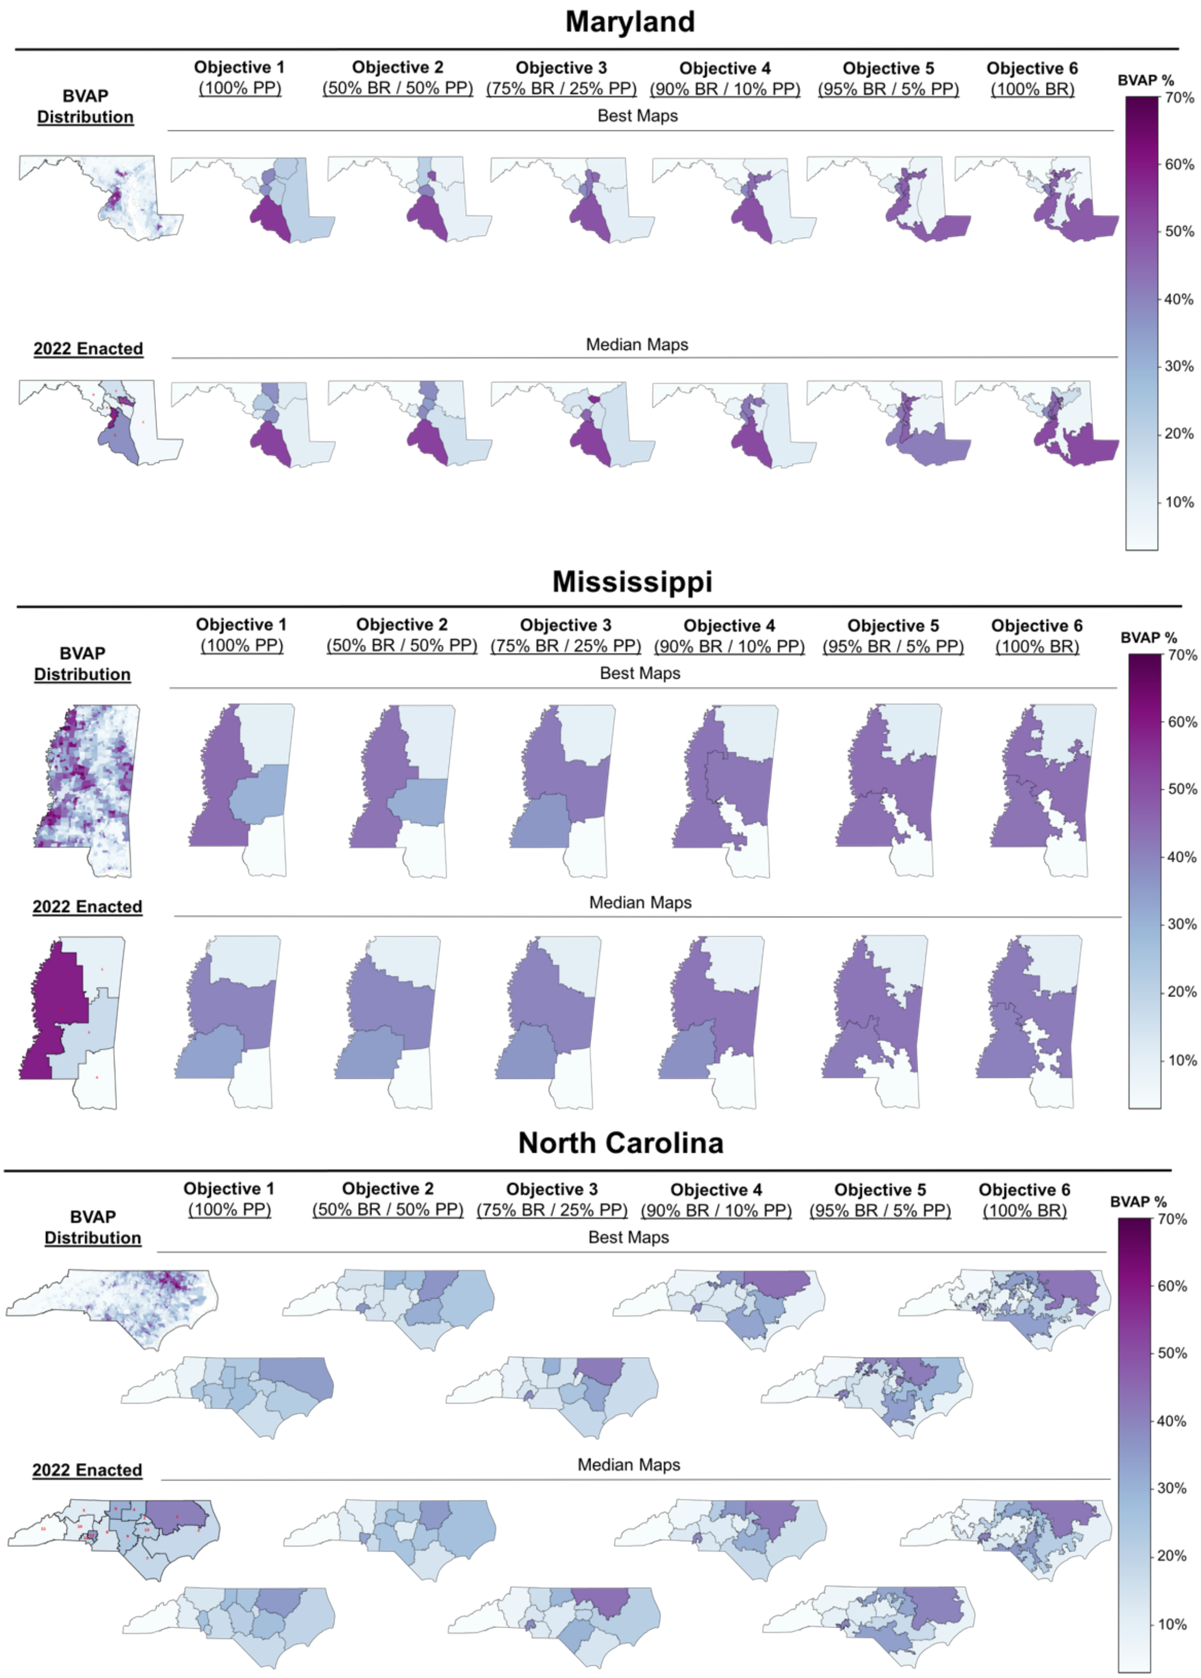

Figure 8b. Maps Generated in Nine States from Six Objective Functions (cont.)

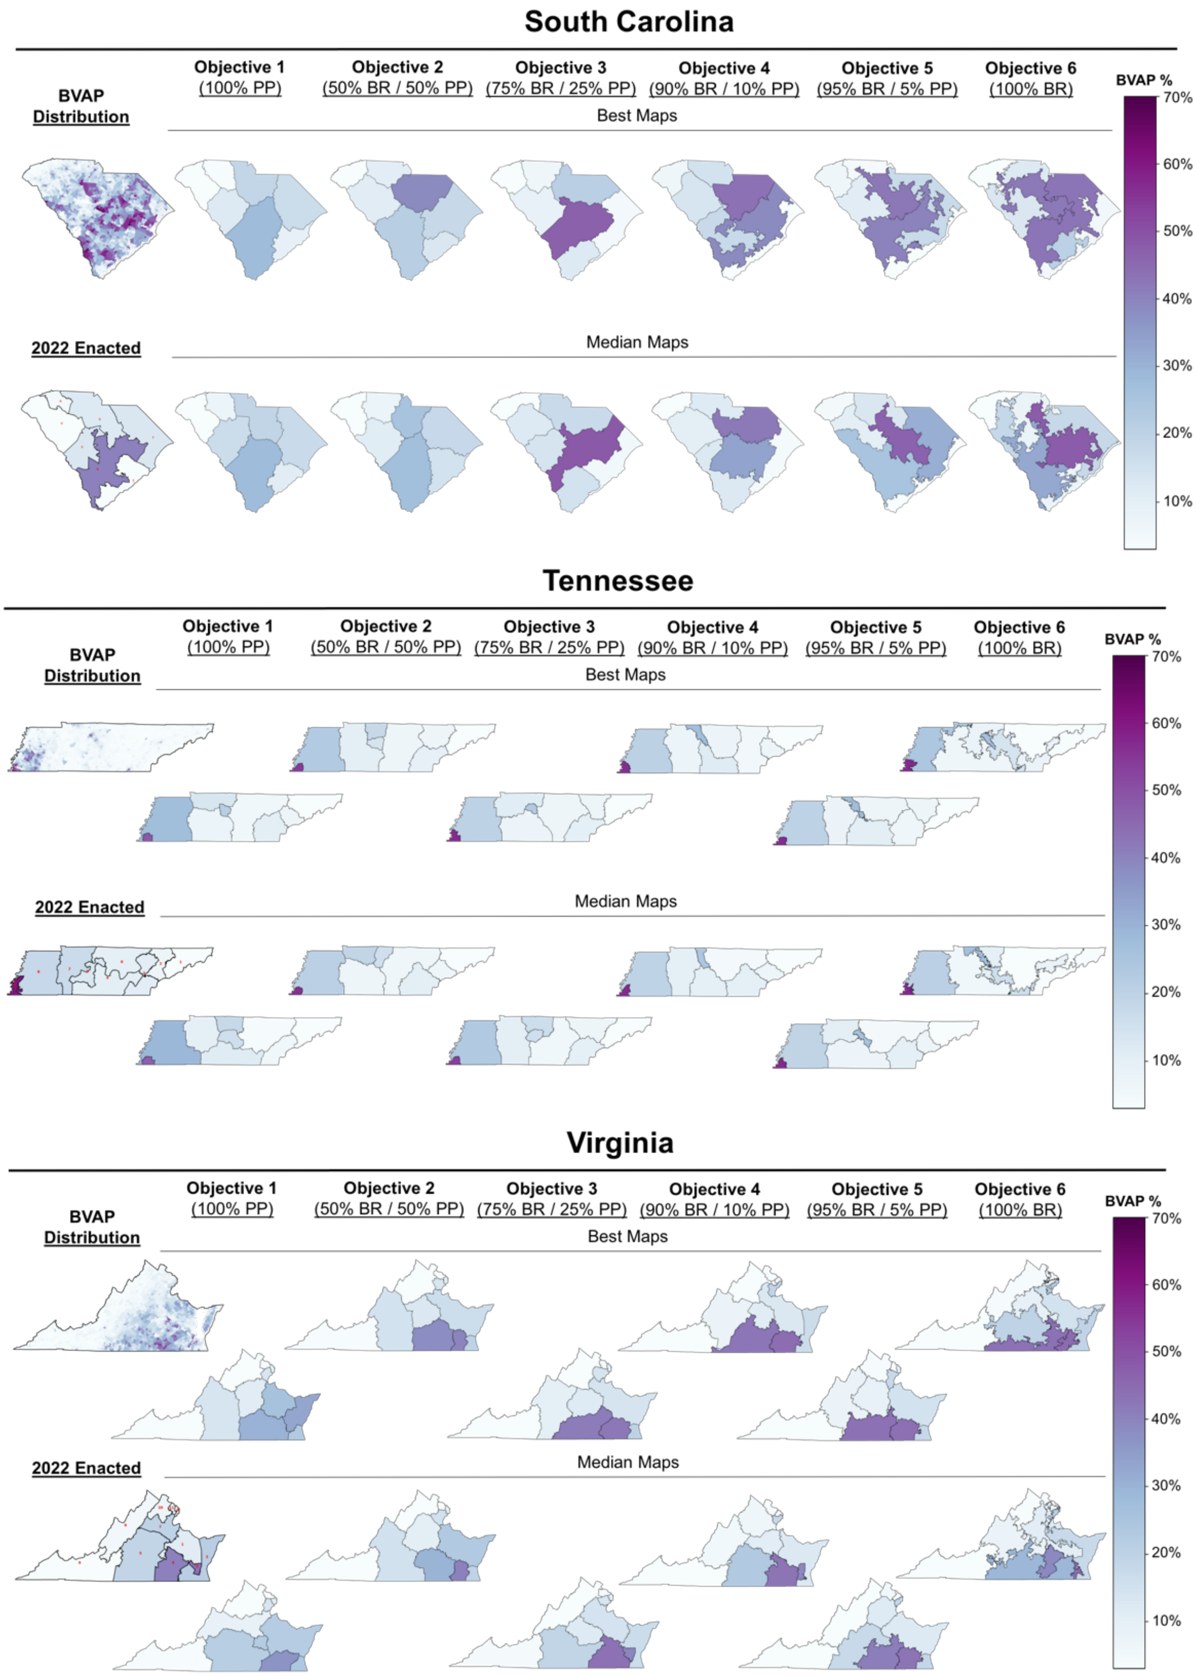

Figure 8c. Maps Generated in Nine States from Six Objective Functions (cont.)

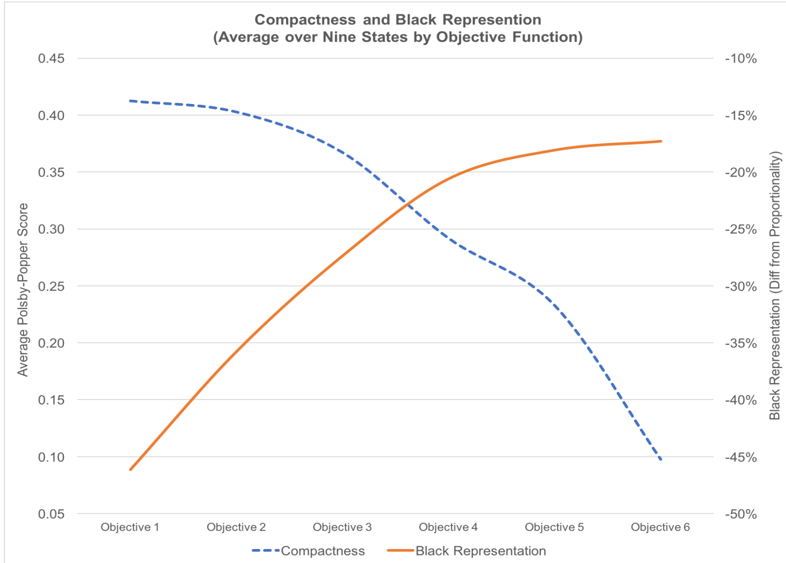

Figure 9. Trade-Off Between Compactness and Black Representation (Average over Nine States by Objective Function)

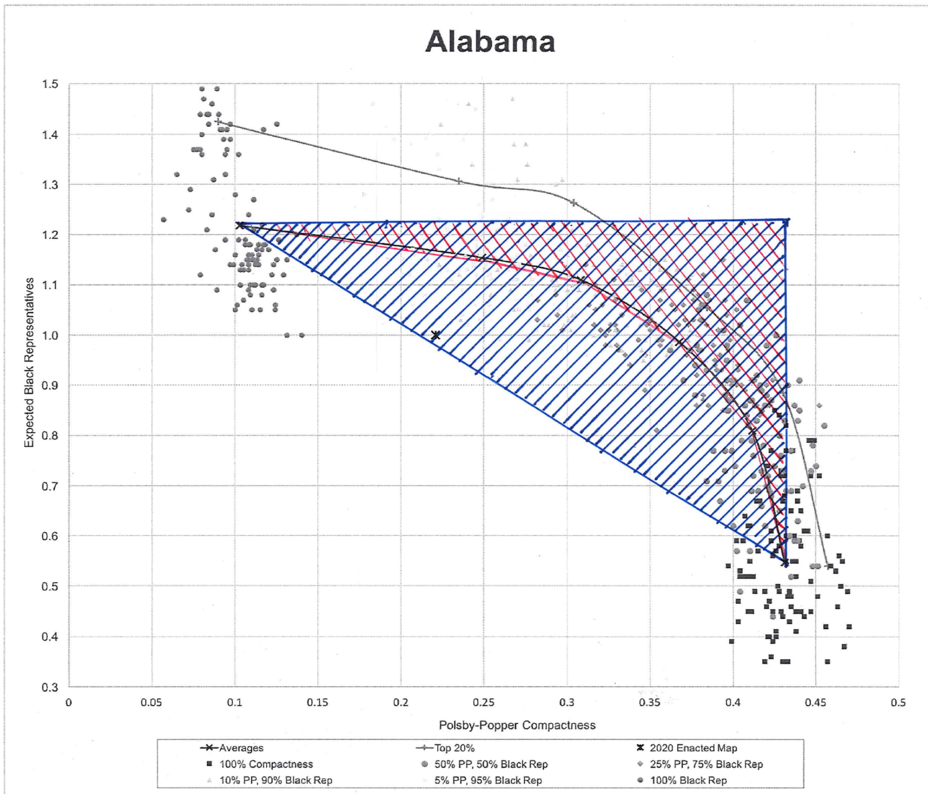

Figure 10. Example of Trade-Off Ratio Calculation

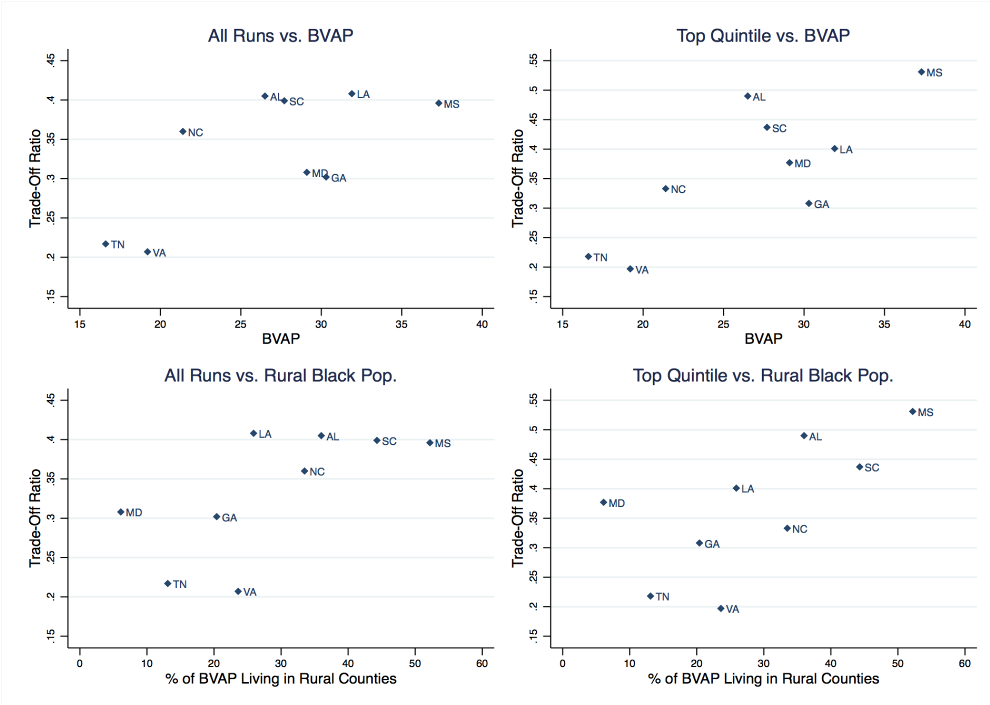

Figure 11. Black Population and Urbanization vs. Trade-Off Ratio

Figure Caption List

- Figure 1. North Carolina Congressional Districts in 1990s and 2010s

Figure 1 Caption: This figure shows the evolution of North Carolina's congressional districts across two decades due to Court decisions involving racial and partisan gerrymandering. The evolution of maps in each decade from less compact to more compact is evident by contrasting the left-hand side maps from those on the right-hand side.

- Figure 2. Probability of Election of Black Representative as a Function of Black Population (Probit Estimate, statewide average Latino population imputed)

Figure 2 Caption: This figure shows the probability assigned by the algorithm that a district in each state will elect a Black representative as a function of its Black population, estimated from election results 1992-2020. For the purpose of this figure, the statewide average Latino population is imputed into the representation formula. Border/Rim states have a higher probability of electing a Black representative than Deep South states at each value for Black population, with the difference being largest in the 35%-50% Black population range.

- Figure 3. Compactness and Proportionality of Black Representation within Congressional Districts by Decade

Figure 3 Caption: This figure shows aggregate Southern trends in compactness and Black representation since the 1970s. In all regions of the South, compactness dropped in the 1990s while Black representation increased substantially. Compactness rose slightly in the 2000s and dropped again in the 2010s without large aggregate change to Black representation.

- Figure 4. Black Representation Proportionality vs. Compactness in Select Southern States

Figure 4 Caption: This figure shows the progression of district compactness and Black representation in seven Southern states since the 1970s. In most states, the largest movement comes in the 1990s with reduced compactness and increased Black representation, with mostly idiosyncratic movement in other decades. Note this figure shows only the first map used in each decade.

- Figure 5. Hill-Climbing Algorithm Results in Alabama

Figure 5 Caption: The figure shows the average district compactness and mean Black representatives elected in 600 simulated maps of Alabama, under six different mixes of compactness and Black representation objectives, along with the average of each objective and the average of the top 20 maps from each objective. Most maps under objectives 4 and 5 improve upon the map used in the 2022 elections in both compactness and Black representation.

- Figure 6. Four Algorithmic Maps of Alabama Compared with Enacted Map

Figure 6 Caption: This figure shows the maps of Alabama congressional districts produced by our algorithm that performs best on each of four objectives. Note the Objective 3 and 4 maps produce both more compact districts and more Black representatives than the map enacted for the 2022 elections.

- Figure 7. Hill-Climbing Algorithm Results in Nine States

Figure 7 Caption: This figure shows the district compactness and Black representation of 600 simulations in each of nine states, with 100 maps for each of six different objective mixes. The enacted maps in each state is also shown (the black cross). In all states except Virginia, many simulated maps perform better on both objectives than the real enacted map.

- Figure 8. Maps Generated in Nine States from Six Objective Functions

Figure 8 Caption: This figure shows examples of maps produced by our algorithms for each of nine states under the six objective function mixes. In each case, both the map that performs best on its designated objective function and the median map (i.e. 50^{th} best map) are shown.

- Figure 9. Trade-Off Between Compactness and Black Representation (Average over Nine States by Objective Function)

Figure 9 Caption: This figure shows average compactness and proportionality of Black representation produced by our simulated maps for each of six objectives aggregated across all nine states. The trade-off between these two goals is severe at the extreme pure objectives (Objectives 1 and 4), but much less pronounced among the compromise objectives.

- Figure 10. Example of Trade-Off Ratio Calculation

Figure 10 Caption: This figure depicts the "trade-off" ratio calculation quantifying how severe the necessary trade-off between compactness and Black representation is in each state. The trade-off ratio is the ratio of the red area to the blue area. A smaller ratio indicates that each objective may be more easily maximized without significant sacrifice to the other.

- Figure 11. Black Population and Urbanization vs. Trade-Off Ratio

Figure 11 Caption: This figure plots the "trade-off ratio" between compactness and Black representation in each state against statewide BVAP (top two panels) and the fraction of statewide Black population living in rural counties (bottom two panels). The right panels show the ratio calculated from only the top 20% of simulations. There appears to be negative correlation between Black urbanization and trade-off ratio (with Maryland as an outlier), suggesting it is easier to draw compact Black opportunity districts where the Black population is less rural.

Appendix Tables for
"Black Representation and District Compactness in Southern Congressional Districts"

Appendix A. Historical Compactness in Southern Congressional Districts, 1970s-2010s

Table A1. Historical Compactness and Proportionality of Black Representation by Decade

		1970s				
State	D/B	CDs	Statewide Black Pop %	Black Rep. %	Black Rep % - Black Pop %	Polsby-Popper Compactness
AL	Deep	8	30.0%	0.0%	-30.0%	0.282
AR	Border	4	21.8%	0.0%	-21.8%	0.271
FL	Border	12	17.8%	0.0%	-17.8%	0.111
GA	Deep	10	28.5%	6.0%	-22.5%	0.289
KY	Border	7	7.1%	0.0%	-7.1%	0.217
LA	Deep	8	31.9%	0.0%	-31.9%	0.166
MD	Border	8	16.7%	12.5%	-4.2%	0.186
MS	Deep	5	42.0%	0.0%	-42.0%	0.220
NC	Border	11	24.5%	0.0%	-24.5%	0.254
OK	Border	6	6.6%	0.0%	-6.6%	0.223
SC	Deep	6	34.8%	0.0%	-34.8%	0.200
TN	Border	9	16.5%	0.0%	-16.5%	0.267
TX	Border	23	12.4%	1.7%	-10.7%	0.234
VA	Border	10	20.6%	0.0%	-20.6%	0.204
Deep South Avg.			33.4%	1.2%	-32.2%	0.231
Border South Avg.			16.0%	1.6%	-14.4%	0.219
All South Avg.			22.2%	1.4%	-20.8%	0.223

		1980s				
State	D/B	CDs	Statewide Black Pop %	Black Rep. %	Black Rep % - Black Pop %	Polsby-Popper Compactness
AL	Deep	7	26.2%	0.0%	-26.2%	0.278
AR	Border	4	18.3%	0.0%	-18.3%	0.292

FL	Border	15	15.3%	0.0%	-15.3%	0.128
GA	Deep	10	25.8%	10.0%	-15.8%	0.245
KY	Border	7	7.2%	0.0%	-7.2%	0.293
LA	Deep	8	29.8%	2.5%	-27.3%	0.131
MD	Border	8	17.8%	12.5%	-5.3%	0.165
MS	Deep	5	36.8%	12.0%	-24.8%	0.266
NC	Border	11	22.2%	0.0%	-22.2%	0.231
OK	Border	6	6.7%	0.0%	-6.7%	0.223
SC	Deep	6	30.4%	0.0%	-30.4%	0.233
TN	Border	8	15.8%	10.0%	-5.8%	0.266
TX	Border	24	12.5%	4.2%	-8.3%	0.207
VA	Border	10	18.5%	0.0%	-18.5%	0.190
Deep South Avg.			29.8%	4.9%	-24.9%	0.230
Border South Avg.			14.9%	3.0%	-12.0%	0.222
All South Avg.			20.2%	3.7%	-16.6%	0.225

		1990s				
State	D/B	CDs	Statewide Black Pop %	Black Rep. %	Black Rep % - Black Pop %	Polsby-Popper Compactness
AL	Deep	7	25.2%	14.3%	-11.0%	0.190
AR	Border	4	15.9%	0.0%	-15.9%	0.273
FL	Border	23	13.6%	13.0%	-0.5%	0.194
GA	Deep	11	26.9%	27.3%	0.3%	0.175
KY	Border	6	7.1%	0.0%	-7.1%	0.234
LA	Deep	7	28.3%	20.0%	-8.3%	0.079
MD	Border	8	24.9%	25.0%	0.1%	0.182
MS	Deep	5	35.5%	20.0%	-15.5%	0.222
NC	Border	12	22.5%	16.7%	-5.9%	0.089
OK	Border	6	7.0%	13.3%	6.4%	0.226
SC	Deep	6	29.8%	16.7%	-13.2%	0.156

TN	Border	9	15.9%	11.1%	-4.8%	0.185
TX	Border	30	11.9%	6.7%	-5.2%	0.121
VA	Border	11	18.8%	9.1%	-9.7%	0.170
Deep South Avg.			29.2%	19.6%	-9.5%	0.164
Border South Avg.			15.3%	10.5%	-4.7%	0.186
All South Avg.			20.2%	13.8%	-6.4%	0.178

		2000s				
State	D/B	CDs	Statewide Black Pop %	Black Rep. %	Black Rep % - Black Pop %	Polsby-Popper Compactness
AL	Deep	7	26.1%	14.3%	-11.8%	0.175
AR	Border	4	15.5%	0.0%	-15.5%	0.262
FL	Border	25	15.4%	12.8%	-2.6%	0.168
GA	Deep	13	29.7%	30.8%	1.1%	0.131
KY	Border	6	7.5%	0.0%	-7.5%	0.278
LA	Deep	7	31.9%	14.3%	-17.6%	0.199
MD	Border	8	28.8%	25.0%	-3.8%	0.082
MS	Deep	4	37.0%	25.0%	-12.0%	0.239
NC	Border	13	21.1%	15.4%	-5.7%	0.149
OK	Border	5	7.3%	0.0%	-7.3%	0.252
SC	Deep	6	28.2%	20.0%	-8.2%	0.224
TN	Border	9	16.5%	4.4%	-12.1%	0.177
TX	Border	32	11.5%	8.8%	-2.8%	0.249
VA	Border	11	19.6%	9.1%	-10.5%	0.165
Deep South Avg.			30.6%	20.9%	-9.7%	0.194
Border South Avg.			15.9%	8.4%	-7.5%	0.198
All South Avg.			21.2%	12.8%	-8.3%	0.196

		2010s				
State	D/B	CDs	Statewide Black Pop %	Black Rep. %	Black Rep % - Black Pop %	Polsby-Popper Compactness

AL	Deep	7	26.5%	14.3%	-12.2%	0.178
AR	Border	4	15.7%	0.0%	-15.7%	0.199
FL	Border	27	15.3%	11.1%	-4.2%	0.109
GA	Deep	14	30.3%	28.6%	-1.7%	0.247
KY	Border	6	7.9%	0.0%	-7.9%	0.191
LA	Deep	6	31.9%	16.7%	-15.2%	0.079
MD	Border	8	29.1%	25.0%	-4.1%	0.046
MS	Deep	4	37.3%	25.0%	-12.3%	0.199
NC	Border	13	21.4%	15.4%	-6.0%	0.106
OK	Border	5	7.2%	0.0%	-7.2%	0.251
SC	Deep	7	27.7%	17.1%	-10.6%	0.170
TN	Border	9	16.6%	0.0%	-16.6%	0.202
TX	Border	36	11.5%	13.9%	2.4%	0.184
VA	Border	11	19.2%	14.5%	-4.7%	0.112
Deep South Avg.			30.7%	20.3%	-10.4%	0.175
Border South Avg.			16.0%	8.9%	-7.1%	0.156
All South Avg.			21.3%	13.0%	-8.3%	0.162

Appendix B: Benchmarking for Choice of Run Length

To determine an optimal run length for generating each single map under our hill-climbing algorithm, we conducted 50 benchmarking runs in one Deep South state (AL) and one Border South state (NC) under three objectives using seven different run lengths (1000, 2500, 5,000, 10,000, 25,000, 50,000, and 100,000 steps). Figure B1 plots the mean objective function value for each of the three objective in Alabama and North Carolina at each run length, while Table B1 shows the average values for compactness, Black representation, and time per run.

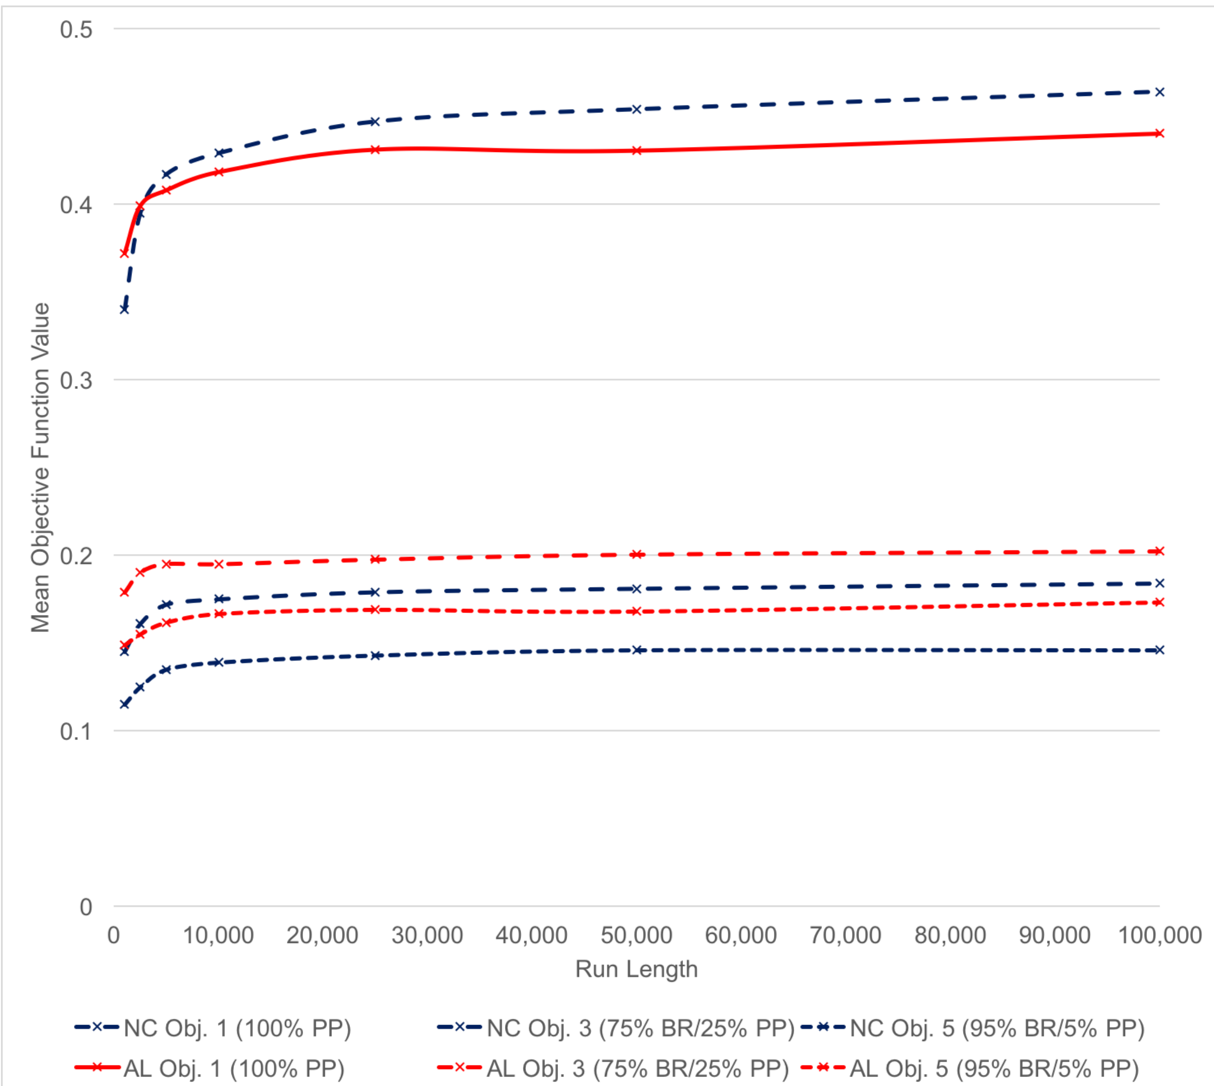

Figure B1. Objective Function Value at Various Run Lengths, NC and AL

In almost all cases, performance increased significantly up to a run length of 10,000 steps, but improvements beyond that were minimal. In particular, extending the run length beyond 25,000 saw very modest continued improvements to compactness, but almost no improvement to Black representation. Across states and objective functions, increasing the run length from 25,000 to 100,000 steps saw an average increase in compactness of 3%, but less than a 1% average increase to Black representation. Meanwhile, increasing the run length from 25,000 to 100,000 increased the average time of each run from just over two hours to more than eight hours. Based on this, we chose a universal run length of 25,000 steps.

Table B1. Performance of Algorithm on NC & AL CDs Under Different Run Lengths

North Carolina

Run Length	Objective	Runs	Polsby-Pop.	Black Rep.	Time
1,000	100% Compactness	48	0.340	0.89	0:04:02
2,500	100% Compactness	50	0.395	0.87	0:09:14
5,000	100% Compactness	50	0.417	0.85	0:18:55
10,000	100% Compactness	50	0.429	0.89	0:39:25
25,000	100% Compactness	100	0.447	0.87	1:36:50
50,000	100% Compactness	50	0.454	0.89	3:12:41
100,000	100% Compactness	50	0.464	0.87	6:34:09
1,000	75% Black Rep; 25% Compact.	50	0.299	1.32	0:04:19
2,500	75% Black Rep; 25% Compact.	50	0.345	1.41	0:10:31
5,000	75% Black Rep; 25% Compact.	50	0.376	1.45	0:21:39
10,000	75% Black Rep; 25% Compact.	50	0.381	1.48	0:43:05
25,000	75% Black Rep; 25% Compact.	100	0.394	1.51	1:50:53

50,000	75% Black Rep; 25% Compact.	50	0.407	1.47	3:43:03
100,000	75% Black Rep; 25% Compact.	50	0.412	1.52	7:42:13
1,000	95% Black Rep.; 5% Compact.	49	0.204	1.54	0:04:30
2,500	95% Black Rep.; 5% Compact.	50	0.219	1.68	0:11:38
5,000	95% Black Rep.; 5% Compact.	50	0.212	1.83	0:25:30
10,000	95% Black Rep.; 5% Compact.	50	0.217	1.89	0:52:10
25,000	95% Black Rep.; 5% Compact.	101	0.214	1.95	2:12:45
50,000	95% Black Rep.; 5% Compact.	50	0.219	1.98	4:15:15
100,000	95% Black Rep.; 5% Compact.	54	0.219	1.99	9:27:58

Alabama

	Run Length / Objective	Runs	PP	Black Rep	Time
1,000	100% Compactness	50	0.372	0.79	0:04:33
2,500	100% Compactness	50	0.399	0.81	0:11:11
5,000	100% Compactness	50	0.408	0.79	0:22:52
10,000	100% Compactness	50	0.418	0.80	0:46:29
25,000	100% Compactness	100	0.431	0.82	1:52:49
50,000	100% Compactness	50	0.430	0.82	3:39:07
100,000	100% Compactness	50	0.440	0.78	6:58:41
1,000	75% Black Rep; 25% Compact.	50	0.330	1.14	0:05:00
2,500	75% Black Rep; 25% Compact.	50	0.353	1.18	0:12:45
5,000	75% Black Rep; 25% Compact.	50	0.359	1.21	0:27:54
10,000	75% Black Rep; 25% Compact.	50	0.361	1.21	0:54:55
25,000	75% Black Rep; 25% Compact.	100	0.367	1.21	2:28:06

50,000	75% Black Rep; 25% Compact.	50	0.371	1.22	4:30:37
100,000	75% Black Rep; 25% Compact.	50	0.376	1.23	8:35:51
1,000	95% Black Rep.; 5% Compact.	50	0.248	1.23	0:05:28
2,500	95% Black Rep.; 5% Compact.	50	0.249	1.26	0:14:22
5,000	95% Black Rep.; 5% Compact.	50	0.249	1.32	0:30:40
10,000	95% Black Rep.; 5% Compact.	49	0.252	1.35	1:04:06
25,000	95% Black Rep.; 5% Compact.	105	0.249	1.36	2:49:37
50,000	95% Black Rep.; 5% Compact.	50	0.258	1.35	5:13:55
100,000	95% Black Rep.; 5% Compact.	49	0.250	1.40	9:49:12

Appendix C. Complete Compactness and Black Representation Results from Hill-Climbing Algorithm in Nine States

	Average of 100 Runs				Average of Top 20% (20 Runs)			
	Compactness		Black Reps.		Compactness		Black Reps.	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Alabama								
100 % PP	0.43	0.02	0.55	0.12	0.46	0.01	0.54	0.13
50% PP, 50% Black Rep	0.41	0.02	0.81	0.15	0.43	0.01	0.89	0.10
25% PP, 75% Black Rep	0.37	0.04	0.99	0.08	0.38	0.02	1.05	0.06
10% PP, 90% Black Rep	0.31	0.04	1.11	0.11	0.30	0.05	1.26	0.12
5% PP, 95% Black Rep	0.25	0.03	1.15	0.10	0.23	0.03	1.31	0.09
100% Black Rep	0.10	0.02	1.22	0.13	0.09	0.01	1.42	0.03
Georgia								
100 % PP	0.41	0.01	3.34	0.21	0.43	0.01	3.33	0.22
50% PP, 50% Black Rep	0.40	0.02	3.88	0.15	0.42	0.01	3.92	0.15
25% PP, 75% Black Rep	0.36	0.02	4.21	0.12	0.37	0.02	4.32	0.10
10% PP, 90% Black Rep	0.27	0.03	4.47	0.11	0.27	0.03	4.62	0.05
5% PP, 95% Black Rep	0.22	0.02	4.55	0.11	0.21	0.02	4.69	0.05
100% Black Rep	0.10	0.01	4.57	0.10	0.09	0.01	4.71	0.04
Louisiana								
100 % PP	0.39	0.01	0.75	0.10	0.41	0.01	0.75	0.10
50% PP, 50% Black Rep	0.37	0.01	0.99	0.13	0.39	0.01	1.11	0.05
25% PP, 75% Black Rep	0.32	0.03	1.21	0.11	0.33	0.02	1.31	0.06
10% PP, 90% Black Rep	0.22	0.05	1.38	0.06	0.23	0.05	1.45	0.05
5% PP, 95% Black Rep	0.17	0.04	1.43	0.06	0.15	0.04	1.52	0.04
100% Black Rep	0.08	0.02	1.44	0.08	0.07	0.01	1.55	0.03
Maryland								
100 % PP	0.43	0.02	2.29	0.20	0.46	0.01	2.42	0.11
50% PP, 50% Black Rep	0.42	0.02	2.61	0.11	0.44	0.01	2.67	0.10
25% PP, 75% Black Rep	0.37	0.04	2.81	0.10	0.39	0.03	2.91	0.06
10% PP, 90% Black Rep	0.29	0.04	2.96	0.11	0.27	0.03	3.08	0.03
5% PP, 95% Black Rep	0.22	0.03	3.01	0.10	0.22	0.03	3.11	0.02
100% Black Rep	0.10	0.01	3.03	0.10	0.10	0.02	3.14	0.01

	Average of 100 Runs				Average of Top 20% (20 Runs)			
	Compactness		Black Reps.		Compactness		Black Reps.	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Mississippi								
100 % PP	0.38	0.02	1.15	0.03	0.41	0.01	1.15	0.04
50% PP, 50% Black Rep	0.38	0.02	1.19	0.03	0.40	0.01	1.20	0.03
25% PP, 75% Black Rep	0.36	0.02	1.23	0.02	0.38	0.01	1.24	0.02
10% PP, 90% Black Rep	0.28	0.05	1.30	0.04	0.23	0.04	1.36	0.02
5% PP, 95% Black Rep	0.19	0.04	1.34	0.04	0.17	0.03	1.39	0.02
100% Black Rep	0.14	0.03	1.32	0.04	0.11	0.02	1.38	0.02
North Carolina								
100 % PP	0.45	0.01	0.87	0.09	0.47	0.01	0.86	0.09
50% PP, 50% Black Rep	0.44	0.02	1.15	0.12	0.46	0.01	1.19	0.12
25% PP, 75% Black Rep	0.39	0.02	1.51	0.12	0.40	0.03	1.62	0.11
10% PP, 90% Black Rep	0.29	0.02	1.84	0.09	0.28	0.02	1.96	0.05
5% PP, 95% Black Rep	0.21	0.02	1.95	0.07	0.21	0.02	2.05	0.04
100% Black Rep	0.09	0.01	1.95	0.06	0.09	0.01	2.03	0.04
South Carolina								
100 % PP	0.40	0.02	0.42	0.11	0.42	0.00	0.39	0.06
50% PP, 50% Black Rep	0.39	0.02	0.59	0.13	0.41	0.02	0.62	0.11
25% PP, 75% Black Rep	0.35	0.02	0.78	0.06	0.36	0.01	0.82	0.04
10% PP, 90% Black Rep	0.27	0.04	0.90	0.05	0.26	0.04	0.96	0.03
5% PP, 95% Black Rep	0.19	0.03	0.96	0.06	0.17	0.03	1.04	0.03
100% Black Rep	0.07	0.01	1.00	0.05	0.07	0.01	1.08	0.02
Tennessee								
100 % PP	0.39	0.01	0.91	0.04	0.40	0.01	0.91	0.03
50% PP, 50% Black Rep	0.38	0.01	0.96	0.02	0.40	0.01	0.96	0.01
25% PP, 75% Black Rep	0.38	0.01	0.98	0.02	0.39	0.01	0.98	0.02
10% PP, 90% Black Rep	0.35	0.02	1.02	0.02	0.36	0.02	1.04	0.02
5% PP, 95% Black Rep	0.32	0.02	1.05	0.01	0.32	0.01	1.06	0.01
100% Black Rep	0.11	0.01	1.06	0.01	0.10	0.01	1.07	0.00
Virginia								
100 % PP	0.40	0.01	0.74	0.07	0.42	0.01	0.76	0.07
50% PP, 50% Black Rep	0.40	0.02	0.96	0.11	0.41	0.01	1.05	0.10
25% PP, 75% Black Rep	0.37	0.02	1.14	0.08	0.39	0.01	1.20	0.05
10% PP, 90% Black Rep	0.32	0.03	1.27	0.07	0.33	0.03	1.35	0.03
5% PP, 95% Black Rep	0.28	0.03	1.30	0.06	0.29	0.03	1.38	0.02
100% Black Rep	0.10	0.01	1.32	0.06	0.09	0.01	1.40	0.02

Appendix D. Relationship between Racial Representation and Partisan Representation

Figure D1. Expected Democrats vs. Expected Black Representative in Maps Algorithmically Generated to Maximize Compactness and Black Representation

Notes: Expected Democrats estimated from a probit function of Cook's PVI, estimating district partisanship based on the district's average deviation from the national vote in the 2016 and 2020 presidential elections. The slope coefficient of this function is derived from estimates of the partisan variance in U.S. congressional elections over the past 50 years. See [redacted] for details. This data is drawn from the same set of algorithmic runs used in the analysis in Table 1 and Figure 7.

Table D1. Correlation Coefficient between Expected Black Representative and Expected Algorithmically Generated to Maximize Compactness and Black Representation

State	ρ	State	ρ
Louisiana	.917	North Carolina	.472
Alabama	.903	South Carolina	.447
Georgia	.834	Maryland	-.500
Mississippi	.827	Virginia	-.575
Tennessee	.677		

Notes: Represents correlation coefficient between x- and y-axes in Figure D1

Appendix E. Effect of Black Population and Urbanization of Black Population on Trade-Off Ratio between Black Representation and Compactness

	All Runs					
	Include Maryland			Exclude Maryland		
Trade-Off Ratio	(1)	(2)	(3)	(4)	(5)	(6)
BVAP	.0084**	-	.0061	.0088**	-	.0053
	(.0033)		(.0033)	(.0034)		(.0041)
% BVAP Rural	-	.0036*	.0024	-	.0048**	.0030
		(.0015)	(.0015)		(.0018)	(.0022)
Constant	.11	.23	.10	.10	.19	.10
	(.090)	(.048)	(.081)	(.093)	(.061)	(.087)
Observations	9	9	9	8	8	8
R-squared	.483	.443	.646	.524	.534	.650

	Top Quintile					
	Include Maryland			Exclude Maryland		
Trade-Off Ratio	(7)	(8)	(9)	(10)	(11)	(12)
BVAP	.014**	-	.011**	.014**	-	.0076
	(.0040)		(.0039)	(.0044)		(.0042)
% BVAP Rural	-	.0052**	.0032	-	.0080***	.0054*
		(.0022)	(.0017)		(.0021)	(.0022)
Constant	.00055	.22	-.0074	-.0014	.12	-.0032
	(.11)	(.068)	(.095)	(.12)	(.069)	(.088)
Observations	9	9	9	8	8	8
R-squared	.624	.455	.760	.629	.714	.828

Notes: Standard errors in parentheses; *** p<.01, ** p<.05, * p<.10