

A Partisan Solution to Partisan Gerrymandering: The Define-Combine Procedure

Maxwell Palmer* Benjamin Schneer[†] Kevin DeLuca[‡]

August 2022

Abstract

Redistricting reformers have proposed many solutions to the problem of partisan gerrymandering but they all require either bipartisan consensus or the agreement of both parties on the legitimacy of a neutral third party to resolve disputes. In this paper we propose a new method for drawing district maps, the Define-Combine Procedure, that substantially reduces partisan gerrymandering without requiring a neutral third party or bipartisan agreement. One party defines a map of $2N$ equal-population contiguous districts. Then the second party combines pairs of contiguous districts to create the final map of N districts. Using real-world geographic and electoral data, we use simulations and map-drawing algorithms to show that this procedure dramatically reduces the advantage conferred to the party controlling the redistricting process and leads to less biased maps without requiring cooperation or non-partisan actors.

*Associate Professor, Department of Political Science, Boston University. mbpalmer@bu.edu

[†]Assistant Professor, John F. Kennedy School of Government. benjamin_schneer@hks.harvard.edu

[‡]Ph.D Candidate, John F. Kennedy School of Government. kevinluca@g.harvard.edu

Following the 2020 Census, every state in the United States redrew their Congressional and legislative district boundaries (for an overview see Warshaw, McGhee and Migurski (Forthcoming)). Many city councils, school boards, county commissions, and other representative bodies are also in the process of drawing new maps. Redistricting creates opportunities for political actors to benefit certain groups over others. In particular, partisan gerrymandering is used to advantage one political party, even enabling a party to win a majority of the seats in a legislative chamber without winning a majority of votes in the election. For example, in 2018, Republican candidates for the Wisconsin State Assembly won only 45% of the state-wide vote but, due to partisan gerrymandering, won 63 of 99 seats (64%).

Gerrymandering hinders democratic representation. When districts are drawn to amplify disproportionately some voices over others, public policy is less likely to reflect constituent preferences. Gerrymandering also allows parties to undermine electoral competition by drawing districts that create, enhance, or lock in a partisan advantage and shield representatives from accountability.

Current reforms for partisan gerrymandering, including bipartisan or non-partisan commissions, all require either bipartisan cooperation or a neutral third-party actor, such as a judge, special master, or independent tie-breaker, trusted by both political parties, to select a fair map. However, in today's hyper-partisan environment, there are few such actors considered able to fulfill this role fairly by both sides, and attempts at cooperation have generated acrimony. Intense controversies continue to surround redistricting, even in states that have enacted anti-gerrymandering reforms.

We propose a new method for drawing maps, the *Define-Combine Procedure* (DCP), which reduces partisan gerrymandering without a neutral third party or bipartisan cooperation. We develop a simple framework that allows each party to act in their own partisan self-interest but achieves a significantly fairer map than would be drawn by either party on its own. We divide the districting process into two stages. Suppose a state must be divided into N equal-population districts. One party—the “definer”—draws $2N$ contiguous, equal-

population districts. Then, the second party—the “combiner”—selects contiguous pairs of districts from the set defined by the first party to create the final districts. This method produces N equally populated, contiguous districts. By dividing the responsibility of drawing districts into two separate stages, in which each party retains complete autonomy in their own stage, the parties counteract each other’s partisan ambitions while maintaining considerable flexibility to achieve other objectives, including maintaining compactness and communities of interest. We show that DCP produces maps with large reductions in bias advantaging either political party. Using simulations based on real-world geographic and electoral data, we assess DCP’s performance in all states where congressional redistricting occurs. To our knowledge, this represents the first effort to date demonstrating how a new map-drawing procedure could lead to tangible improvements across such a wide variety of geographic and electoral contexts.

Limitations of Current Partisan Gerrymandering Fixes

Reformers have sought to reduce partisan gerrymandering by shifting redistricting authority from partisan legislatures to nonpartisan independent redistricting commissions. State legislatures draw districts in over half of the U.S. states, with the remainder relying on commissions (see Appendix Table A.1). However, a majority of commission-based states are not actually independent from interference by a state legislature, and those that are rely on a member or members to act neutrally, sometimes touching off controversy—as when the Governor of Arizona sought to impeach the independent chair of her state’s commission (Druke, 2017).

The 2020 cycle revealed some of the challenges of independent redistricting commissions. In New York and Ohio, the commissions failed to produce maps, and the state legislatures passed maps that were found to be partisan gerrymanders by each state’s Supreme Court. A popular initiative passed in 2018 in Utah establishing an independent commission was

altered by Republican state legislators in 2020, allowing them to bypass the commission’s proposed districts. In Virginia, the commission had an equal number of Democratic and Republican members and failed to agree on a map, forcing the state Supreme Court to appoint a special master. Even in states where commissions did pass maps—such as Michigan and Colorado—the process engendered heated accusations of partisan and racial gerrymandering.¹ Appendix A.1.1 further details challenges presented by redistricting in legislatures, commissions, and through litigation.

Academic researchers have proposed a variety of approaches to measuring and identifying partisan gerrymandering, including methods to quantify partisan gerrymandering and using simulations to identify when a map is a partisan gerrymander (e.g. Stephanopoulos and McGhee, 2018; Chen and Rodden, 2015). Others have drawn inspiration from the cake-cutting problem to propose processes by which actual maps could be drawn fairly. While this logic has inspired several redistricting proposals (Landau, Reid and Yershov, 2009; Pegden, Procaccia and Yu, 2017; Ely, 2019; Alexeev and Mixon, 2019; Brams, 2020), none of the existing work on map-drawing processes has addressed implementation across different states or electoral contexts using real-world data. The complexity of most previous proposals also precludes real-world application or even simulation using modern computing techniques. Appendix A.1.2 describes these proposals.

DCP contrasts with the existing proposals in several ways. First, *it does not require either forging bipartisan support between parties or appointing a third-party arbiter to resolve points of difference*. A procedure that sidesteps these common stumbling blocks could reduce the contentious political disputes that accompany decennial redistricting, leading to fairer maps and less partisan acrimony. Second, *DCP is a process-based solution; it does not require courts to define and agree on an empirical standard for partisan gerrymandering*. State

¹In Michigan, Democratic state legislators argued that the commission’s districts discriminated against Black voters: <https://www.bridgemi.com/michigan-government/michigan-supreme-court-dismisses-redistricting-suit-over-black-representation>. In Colorado, Democrats and Hispanic advocacy groups argued the map discriminated against Hispanic voters: <https://coloradosun.com/2021/11/01/colorado-congressional-map-approved-supreme-court/>.

courts could resolve redistricting disputes by ordering the parties to use DCP to draw a remedial map. Third, *DCP is simple and could be implemented efficiently*. Several other proposed solutions have appealing game-theoretic properties but in practice require multiple rounds of bargaining or map-drawing and are difficult or impossible to solve computationally; in contrast, DCP’s two-stage process efficiently produces a complete and valid districting plan.

The Define-Combine Procedure

Suppose a state needs to be divided into N contiguous, equally-populated, single-member districts. Elections in the state are contested by two parties, A and B . We assume that both parties are risk-neutral seat-maximizers; their goal is to win as many of the N seats as possible in the next election.²

We consider two methods of drawing district maps. First, one party has unilateral control of the process. Given a set of potential valid maps, when party i draws the maps, it will select a map that maximizes the number of seats it wins. Under this method of redistricting, we should expect that, where possible, party i will maximize partisan advantage relative to party j by strategically cracking and packing party j ’s voters to minimize the seats won by party j . In many cases, it will be possible for party i to win a substantially larger share of the seats than its statewide vote share.

This method, which we will call the “unilateral redistricting process” (URP), approximates the redistricting process in states where one party controls redistricting for a given map. While other factors, such as incumbency or vote margins in close seats, factor into districting decisions, gerrymandering for partisan gain is a primary objective for the party in power. For example, when independent experts have used map-drawing software to simulate thousands of possible maps in a state, the observed maps in states where redistricting is

²Parties may also have other goals. For example, parties might seek to maximize bias, or they might have preferences over the level of responsiveness in the plan (Katz, King and Rosenblatt, 2020).

controlled by a single party often appear to be among the most partisan possible.³

The second method we consider is our own innovation, the two-stage Define-Combine Procedure. In this model, map-drawing is divided between the two parties but in each stage one party acts unilaterally.

Suppose Party A acts in the first stage as the “Definer,” and Party B acts in the second stage as the “Combiner.”

1. Party A defines a set of $2N$ contiguous, equally populated districts. To avoid confusion with the following stage, we refer to these units as *subdistricts*.
2. Party B creates the final map of N districts by combining pairs of 2 contiguous sub-districts.

Appendix A.2 describes the properties of the model. The optimal solution can be found with backward induction. Given the subdistricts defined by A , B will select the combination that maximizes the number of seats B will win. Knowing this, A selects the set of subdistricts that minimizes the number of seats B can win in the second stage.

Appendix A.3 provides an in-depth application of URP and DCP to the state of Iowa to illustrate the intuition from the perspectives of the two parties.

Evaluating the Define-Combine Procedure Using Simulations

To evaluate DCP, we use map-drawing algorithms to simulate and compare both redistricting procedures. Unlike other studies of gerrymandering that use simulation methods, we are interested not in the *distribution* of possible outcomes in each state but rather in the limit—the most extreme possible map drawn by each party using each process. To identify these

³Research on redistricting in Florida (Chen and Rodden, 2015), Maryland (Cho and Liu, 2016), Wisconsin (Chen, 2017), Virginia (Magleby and Mosesson, 2018), and Pennsylvania (Duchin, 2018) find that chosen plans in states with URPs are extreme outliers as compared to the set of simulated possible maps.

maps, we employ the “shortburst” algorithm (Cannon et al., 2020), which begins with a starting map and assigns it a numerical score, such as the number of seats won by a given party.⁴ The algorithm then generates a set of N variations of the starting map in a “burst,” calculates a score for each using the same scoring function, and selects the map with the highest score (or the starting map if there is no improvement). We repeat this burst process until the score remains stable. As Cannon et al. (2020) shows, this approach finds more extreme plans than those identified by random walks and other distributional approaches. Appendix A.4 provides our detailed methodology.

To simulate maps drawn under unilateral redistricting, we apply the shortburst algorithm to find the maps that maximize seats won for each party. For each state and party, we run 10 separate sets of simulations. We (1) generate a random starting map, (2) run the shortburst algorithm 2000 times, generating 10 maps per burst, and (3) save the final map. The scoring function maximizes first the number of seats won by the party and second the party’s vote share in the next closest district.⁵

Second, we use a “nested shortburst” algorithm to simulate DCP. For each state and party we run 10 separate sets of simulations. We (1) generate a random starting map of $2N$ districts and (2) run the shortburst burst algorithm 250 times, with 10 maps generated per burst. However, we score each using a second iteration of the shortburst algorithm, which finds the pairing of districts that maximizes seats for the second party. This approach gives both parties the benefit of actively trying to maximize their own seat advantage, rather than looking at a random set of proposals and choosing the best option.

Using election data from the Voting and Election Science Team (2022) merged with 2020 census data from the ALARM Project (Kenny and McCartan, 2021), we ran simulations for every state with at least two congressional districts. We measure partisanship using the 2020 presidential election results. For each simulation we use a 1% population deviation

⁴Implemented in the `redist` package at https://alarm-redist.github.io/redist/reference/redist_shortburst.html.

⁵This scoring function finds extreme gerrymanders more efficiently than scoring based on seats won alone.

constraint (using 2020 total population), a simple compactness constraint, and require all districts to be contiguous.⁶

Results

We compare the performance of DCP to two key benchmarks—simulated unilateral redistricting and actual, adopted plans—along a variety of different metrics. The metrics we examine include: (1) the advantage, which we will term *definer’s advantage*, conferred to the map-drawing party by a redistricting procedure; (2) deviation from proportionality; (3) deviation from a fairness target (e.g., that accounts for geographic biases in a state that might advantage one party); and, (4) partisan bias induced by a redistricting procedure.⁷ Appendix A.5 describes these measures in detail. We are agnostic as to the metric used to evaluate a given redistricting plan; our goal is to show the performance of DCP, and its benefits, across a variety of measures.

The simulation results reveal that DCP dramatically reduces the partisan advantage conferred to a map maker as compared to URP. Figure 1 presents four maps that compare seats won by each party under four redistricting procedures. The top two maps show the results of the unilateral redistricting simulations, with Democrats redistricting on the left and Republicans redistricting on the right; the bottom two maps show the results of DCP.

Of the 429 seats redistricted in the simulations, Democrats win 335 (78.1%) when they draw the maps unilaterally in every state. When Republicans draw all maps unilaterally, Democrats win 132 seats (30.8%). The simulations reveal a possible swing of 203 seats between parties through unilateral redistricting. These results represent a theoretical maximum for the shift in seats that could occur due to partisan gerrymandering. Our unilateral

⁶Weakening the compactness constraint does not substantially affect our results. We do not employ other constraints to comply with the Voting Rights Act or state redistricting laws. Appendix A.6.6 presents a test of DCP with constraints to comply with the VRA.

⁷Partisan bias is the difference between seat share and vote share when votes are split evenly between the parties. For proportionality, we examine the difference between the two-party vote share and seat share. For a fairness target, we take the midpoint of the most extreme maps possible under each party’s simulated unilateral redistricting procedure (Benadè, Procaccia and Tucker-Foltz, 2021).

simulations are more extreme than those actually drawn in many states considered to be partisan gerrymanders because we impose limited constraints and maximize seats won with a strict 50% of the vote cutoff (rather than creating a “toss-up” category).

As illustrated in the bottom two maps, DCP dramatically reduces the swing of seats between parties; in many states the difference in seats won based on who defines and who combines falls to either zero or one. When Democrats draw the Define-stage map, they win 244 (56.9%) seats; when Republicans draw the Define-stage map, Democrats win 211 (50.8%) seats. When calculating the difference state by state (since some parties have a second- rather than first-mover advantage), the swing between parties based on first-mover versus second-mover status amounts to 38 seats, eliminating over 80% of the swing in seats theoretically possible under unilateral partisan control of redistricting. Tables reporting the full underlying results are available in Appendix A.6.1. State by state results are available in Appendix A.6.4.

Table 1 summarizes the results for Definer’s Advantage for URP and DCP, along with the other metrics we consider (we do not report a quantity for Definer’s Advantage for Adopted plans because we do not observe seat share under the counterfactual that the out-party controlled the redistricting process—though we report an imputed value in Appendix A.6.2). The next two metrics evaluate the redistricting procedures based on deviations from target outcomes, specifically proportionality and a geometric fairness target (defined in Benadè, Procaccia and Tucker-Foltz (2021) as the average of the worst- and best-case map for each party). We evaluate each of these metrics based on current party control of the redistricting process in each state. For instance, for a simulated unilateral map for the party in power we calculate the deviation in seat share from the (1) seat share proportional to a state’s vote share, and (2) seat share based on the geometric fairness target. The first of these metrics illustrates how much proportional representation would improve based upon a state’s switch to DCP; the second illustrates how much closer to a geometric target a switch to DCP would achieve.

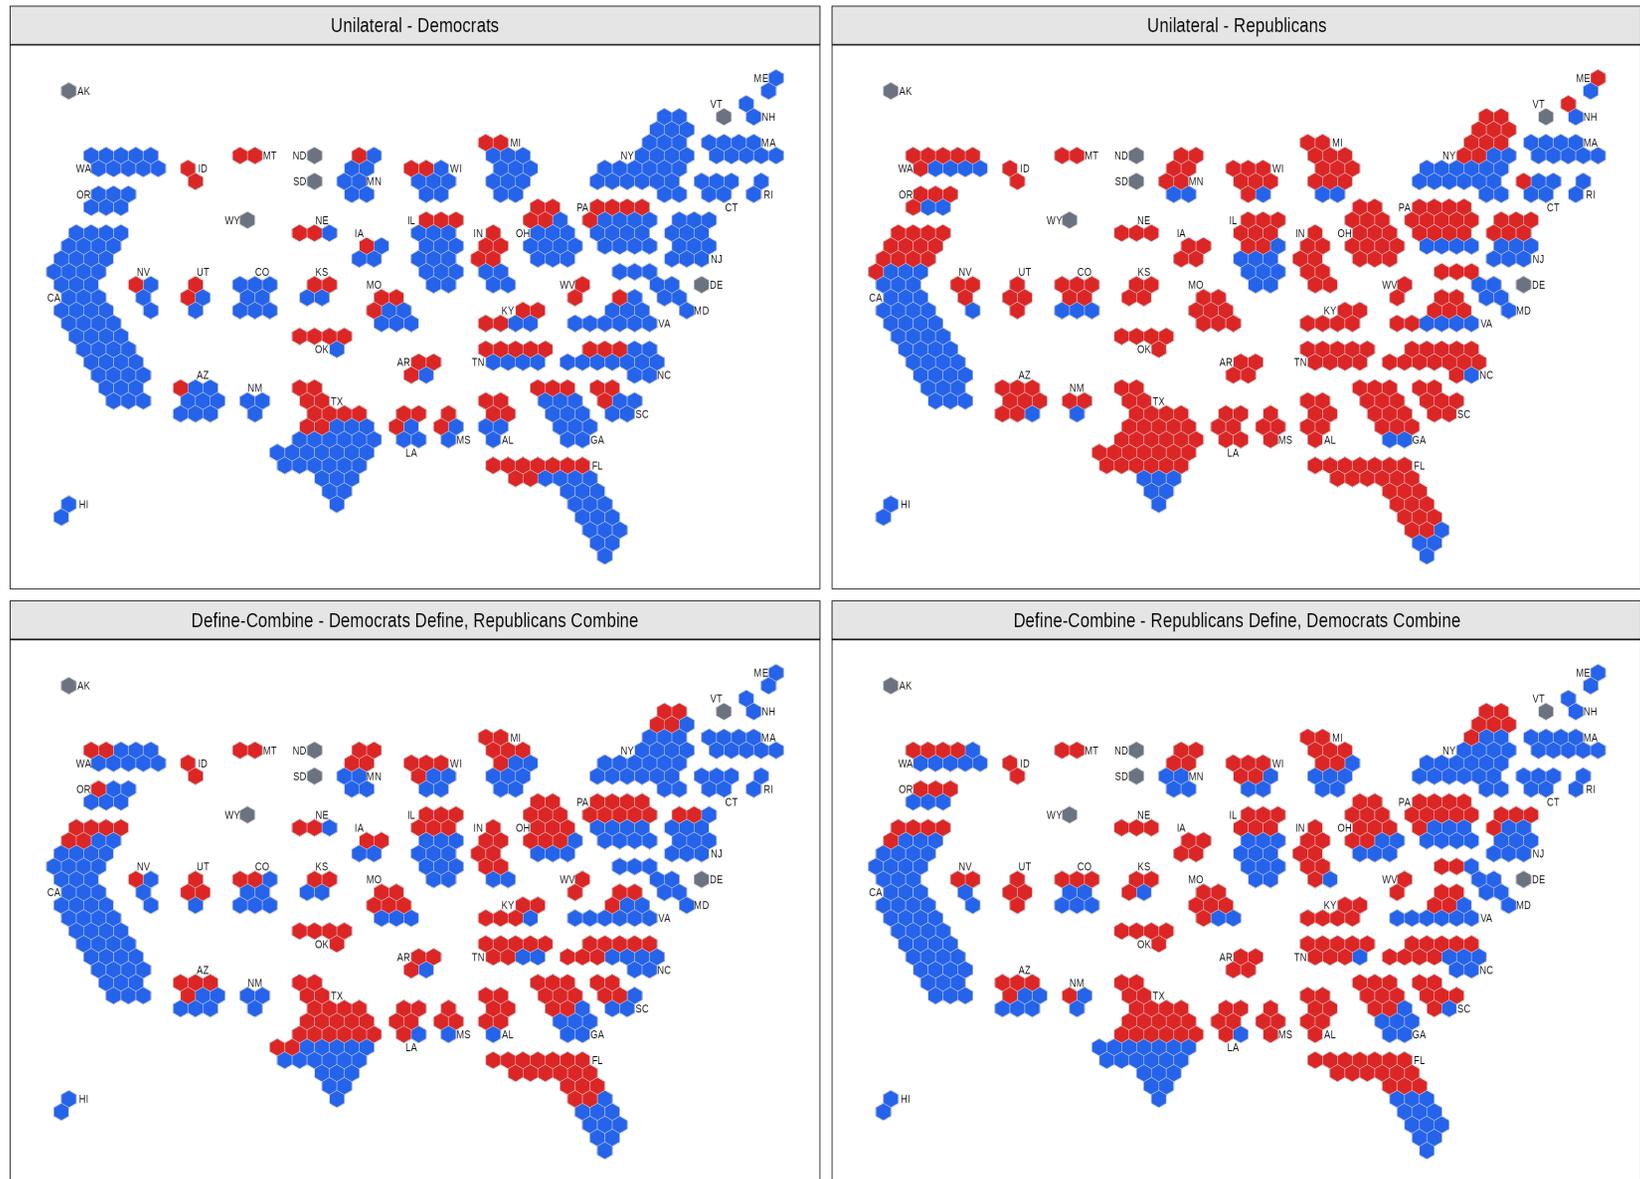


Figure 1: **Maps of Simulation Results by Method and Party:** These figures report the full set results for URP and DCP simulations, exhibiting which party wins each congressional district. Blue hexagons tally Democratic wins; red hexagons tally Republican wins.

Metric	URP	DCP	Adopted Plan
Definer’s Advantage	0.473	0.089	
Deviation from Proportionality	0.270	0.159	0.187
Deviation from Fairness Target	0.237	0.058	0.090
Partisan Bias	0.303	0.060	0.073

Table 1: **DCP Performance Versus Alternatives:** This table reports the performance of the DCP as compared to unilateral redistricting and adopted plans along several different metrics. Definer’s Advantage for adopted plans is omitted since it would involve interpolating seat share under the scenario where the opposing party held control over the redistricting process. Partisan Bias is calculated only for states with 2020 Democratic Presidential Vote Share between 45% and 55%.

For both metrics, DCP improves upon both the simulated outcomes from URP as well as adopted plans (rows 2 and 3 in Table 1). We estimate, for example, that if every state drew their map according to a unilateral process controlled by the majority party of the lower house of the state legislature, and then switched to DCP with the same party serving as the definer, deviation from proportionality would be cut by over one third and deviation from the fairness target would be cut by three quarters. Appendix A.6.3 reports full results broken out by state redistricting procedure. DCP achieves reduced deviations from proportionality compared to enacted maps in states with Legislature or Political redistricting procedures and has comparable or slightly greater deviations from proportionality compared to enacted maps in states with Independent Commissions and Courts or Special Masters. DCP improves upon adopted plans regardless of redistricting procedure in terms of deviations from the geometric fairness target.

Finally, to determine partisan bias for adopted plans, we gathered data on partisan biases from PlanScore.⁸ We replicated the partisan bias calculation for our simulated URP and DCP maps in the same states that PlanScore evaluated partisan bias. For this set of states, we find DCP dramatically reduces partisan bias compared to URP and modestly reduces partisan bias compared to adopted plans in the states (full results in Appendix A.6.5).

⁸Available at <https://planscore.campaignlegal.org/>. PlanScore calculates this metric for states with Democratic vote share between 45% and 55% (AZ, CO, FL, GA, MI, MN, NC, NM, NV, PA, TX, VA, WI).

Most of the improvements compared to enacted maps come from states where legislatures controlled redistricting (see Appendix A.6.3).

Overall DCP exhibits significant improvements compared to our benchmark procedures across a range of different metrics. Importantly, the comparison to enacted maps is not one to one, since our DCP simulations necessarily assume fewer constraints than real-world redistricters face. Thus, the fact that we still observe similar or less extreme maps in our simulations illustrates the powerful tendency of DCP to temper extreme redistricting outcomes.

Finally, we note that real-world examples have just one geographic distribution of voters per state. To test the robustness of DCP and to explore its properties, we also use simulated grid maps that allow for (1) different voter distributions with varying degrees of geographic clustering and (2) different levels of state-wide partisanship, (3) varying objectives beyond seat maximization for the redistricting parties, and (4) varying numbers of districts for a fixed total population. In Appendix A.7, we explore how each of these extensions influence the properties of URP and DCP. Overall, DCP continues to reduce gerrymandering dramatically when varying the degree of state-wide partisanship, the geographic clustering, the parties' objective functions, and the number of districts.

Conclusion

DCP features simple rules, clear strategies for each party, and an efficient framework that can be implemented in every state. One obstacle to implementation for past theoretical approaches has been the difficulty for decision makers in either party to predict outcomes in the real world. DCP, as we have demonstrated, can be applied to real-world geographic data, which allows analysts to predict outcomes and reduce uncertainty surrounding the redistricting process. Furthermore, DCP significantly reduces the advantages conferred to the redistricting party and results in maps more likely to reflect the will of voters. These

advantages hold up across a variety of different contexts reflecting the political and geographic heterogeneity of the states, as well as across a number of different measures of a map’s fairness and level of bias.

There are many challenges in using automated algorithms to aid in the redistricting process (Cho and Cain, 2020). In some cases, advances in computing power and the ability of politicians to consider a large range of maps exacerbates partisan gerrymandering, rather than alleviating it, as partisan mapmakers use map-drawing algorithms to devise increasingly gerrymandered maps. DCP provides an approach that utilizes advances in computing to produce less biased maps—ones where the process-based algorithm itself constrains partisan motives.

Political parties will almost always oppose ceding power, but this is doubly so when the choices they face require embracing significant uncertainty about future political outcomes. Because DCP is a two-stage game, solvable with existing computing resources, it represents a step towards providing an alternative mechanism to court, legislature, or commission-based redistricting that is feasible to implement. This framework gives parties the autonomy to respect communities of interest, geographic boundaries, and other political concerns—that is, to internalize the wide range of factors that play important roles in decisions about redistricting—while nonetheless tempering the partisan biases that emerge during redistricting. By involving both parties but setting them in opposition to each other, rather than requiring bipartisan cooperation or independent third-party mediators, DCP offers a partisan solution to the extraordinarily partisan process that is redistricting.

References

- Alexeev, Boris and Dustin Mixon. 2019. “Partisan Gerrymandering with Geographically Compact Districts.” *Journal of Applied Probability* 55(4):1046–1059.
- Benadè, Gerdus, Ariel Procaccia and Jamie Tucker-Foltz. 2021. “You can have your cake and redistrict it too.” *ACM Transactions on Economics and Computation* .
- Brams, Steven. 2020. “Making Partisan Gerrymandering Fair: One Old and Two New Methods.” *Social Science Quarterly* 101(1):68–72.
- Cannon, Sarah, Ari Goldbloom-Helzner, Varun Gupta, JN Matthews and Bhushan Suwal. 2020. “Voting Rights, Markov Chains, and Optimization by Short Bursts.” *arXiv* .
- 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.
- Cho, Wendy Tam and Bruce Cain. 2020. “Human-centered redistricting automation in the age of AI.” *Science* 369(6508):1179–1181.
- Cho, Wendy Tam and Yan Liu. 2016. “Toward a Talismanic Redistricting Tool: A Computational Method for Identifying Extreme Redistricting Plans.” *Election Law Journal* 15(4):351–366.
- Druke, Galen. 2017. “Want Competitive Elections? So Did Arizona. Then The Screaming Started.” *FiveThirtyEight* .
- Duchin, Moon. 2018. “Outlier Analysis for Pennsylvania Congressional Redistricting.”. <https://maggg.org/uploads/md-report.pdf>.

- Ely, Jeffrey. 2019. “A Cake-Cutting Solution to Gerrymandering.” p. 27. <https://www.jeffely.com/post/a-cake-cutting-solution-to-gerrymandering>.
- Katz, Jonathan, Gary King and Elizabeth Rosenblatt. 2020. “Theoretical Foundations and Empirical Evaluations of Partisan Fairness in District-Based Democracies.” *American Political Science Review* 114(1):1–15.
- Kenny, Christopher and Cory McCartan. 2021. “ALARM Project: 2020 Redistricting Data Files.”. <https://github.com/alarm-redist/census-2020/>.
- Landau, Z., O. Reid and I. Yershov. 2009. “A fair division solution to the problem of redistricting.” *Social Choice and Welfare* 32(3):479–492.
- Magleby, Daniel and Daniel Mosesson. 2018. “A New Approach for Developing Neutral Redistricting Plans.” *Political Analysis* 26(2):147–167.
- Pegden, Wesley, Ariel Procaccia and Dingli Yu. 2017. “A Partisan Districting Protocol with Provably Nonpartisan Outcomes.”. <https://arxiv.org/abs/1710.08781>.
- Stephanopoulos, Nicholas and Eric McGhee. 2018. “The Measure of a Metric: The Debate over Quantifying Partisan Gerrymandering.” *Stanford Law Review* 70(5):1503–1566.
- Voting and Election Science Team. 2022. “Precinct Level Election Results, 2016–2020.” *Harvard Dataverse* . <https://dataverse.harvard.edu/dataverse/electionscience>.
- Warshaw, Christopher, Eric McGhee and Michal Migurski. Forthcoming. “Districts for a New Decade—Partisan Outcomes and Racial Representation in the 2021 Redistricting Cycle.” *Publius: Journal of Federalism* .

Appendix: Supporting Information for
A Partisan Solution to Partisan Gerrymandering

Contents

A.1	Limitations of Current Partisan Gerrymandering Fixes	1
A.1.1	Already-Implemented Solutions	1
A.1.2	Other Proposed Solutions	5
A.2	The Define-Combine Procedure	8
A.3	A Simple Example: Iowa	11
A.4	Simulation Details	14
A.5	Measurement of Redistricting Plans	15
A.6	Additional Simulation Results	17
A.6.1	Simulation Tables	17
A.6.2	Assessing Definer’s Advantage for DCP versus Enacted Plans	19
A.6.3	Assessing Performance of DCP versus Enacted Plans by State Redistricting Procedure	19
A.6.4	State-by-State Shortburst Comparisons, 2020 Presidential Election Democratic Vote Share	21
A.6.5	Simulation Results After Uniform Swing	27
A.6.6	Compliance with the Voting Rights Act	30
A.7	Grid Maps	32
A.7.1	Grid Map Simulations	32
A.7.2	Varying Partisan Composition of Voters	33
A.7.3	Alternative Objectives to Maximizing Seats Won in Next Election	37
A.7.4	Geographic Partisan Clustering	37
A.7.5	Number of Districts and Definer’s Advantage	41
A.7.6	Partisan Clustering Calculations	43

A.1 Limitations of Current Partisan Gerrymandering Fixes

We can divide current solutions to partisan gerrymandering into two classes. First, there are solutions that are currently used in at least one of the fifty states. These include, for example, legislature-involved redistricting commissions, independent redistricting commissions, and judicial intervention to reduce partisan gerrymandering. Second, there are proposals—generally put forth by researchers—that move beyond currently implemented solutions and lay out some other mechanism by which maps are drawn. These include methods where the parties draw districts by alternating back and forth; many of these approaches are inspired by the cake-cutting problem and principles of fair division (i.e., how to divide a good, such as a cake, fairly between two parties) (Brams and Taylor, 1996).

A.1.1 Already-Implemented Solutions

Citizens have expressed deep dissatisfaction with redistricting procedures currently adopted in the states. For example, fewer than 25% of respondents in the Cooperative Congressional Election Survey answered affirmatively when asked whether redistricting in their state was fair (Schaffner and Ansolabehere, 2015).⁹ In a number of states, voters or legislators have responded by establishing redistricting commissions that are meant to de-politicize the map-making process and produce districts that are more fair.

Table A.1 reports the specific type of commission used in each commission-based state.¹⁰ In total, 29 states¹¹ draw their Congressional district maps through the legislature exclusively, while the rest use some sort of redistricting commission. The redistricting process in the majority of commission-based states, however, is not in fact independent from the state legislature. Of the 21 states that use some form of redistricting commissions, only nine states have truly independent commissions that are able to create maps without input or approval of the state legislature. Advisory commissions assist the legislature as it draws district boundaries, but the legislature approves the maps. Political or politician commissions are mainly comprised of elected officials. In backup-commission states, a commission, sometimes comprised of politicians or politician-appointed members, only plays a role if the legislature fails to pass a districting plan within a certain time period. Finally, independent commissions are distinct from the others as they do not include public officials or legislators.¹²

⁹Though commission-based states did register significantly higher approval rates than legislature-based states.

¹⁰We use classifications from Justin Levitt’s website, *All About Redistricting: Who Draws the Lines?*, with some additional classification (Levitt, 2020).

¹¹Four of these 29 states have only have one Congressional districts and do not actually engage in Congressional redistricting.

¹²Selection methods for independent commissions vary significantly—in Arizona, Idaho, Montana, and

All told, 12 of the 21 commission-based states do *not* have a redistricting process for their Congressional districts that is meaningfully independent from the state legislature; as a result, the map-drawing process remains subject to the same partisan pressures as in states with legislature-drawn maps. Two of the nine states with truly independent commissions, Alaska and Montana, use their commissions to draw state legislative districts but since they only have one Congressional district their commissions do not engage in Congressional redistricting. Of the remaining states that have established independent redistricting commissions for redrawing their Congressional districts, nearly all rely on a member or members to act neutrally (often these members are not affiliated with either of the two major political parties). The logic behind this design is that the two parties will have to appeal to a neutral arbiter—the independent member(s) of the commission—in order to achieve a majority and pass a map. In theory, this could cause both parties to curb their partisan gerrymandering efforts in order to create a fairer map appealing to a neutral (and presumably more moderate) commission member.¹³

Scholars have not reached a consensus on the benefits of independent commissions. One study examines the efficacy of redistricting commissions in seven Western states and compares them to five non-commission states in the West, and finds that redistricting commissions do not out-perform legislatures when judged by the metric of drawing compact, competitive districts that preserve preexisting political boundaries (Miller and Grofman, 2013). (On the other hand, the same authors find that commissions seem to excel at producing maps “on time” that avoid litigation.) Others have found that there are more competitive districts in commission-drawn maps in the 1990s and 2000s redistricting cycles (Carson and Crespino, 2004), and that independent commissions are more likely to follow traditional redistricting principles, including compactness, splitting fewer political subdivisions, and preserving the cores of existing districts (Edwards et al., 2017). Recent research using simulations to consider a set of alternative maps that could have been enacted by independent commissions

Washington majority and minority party leaders appoint commissioners, while judges make appointment decisions in Colorado. Alaska has two members chosen by the governor, two by party leaders in the state legislature, and the last by the state supreme court chief justice. California has a process that involves narrowing a pool of applicants down, randomly selecting some members, and then having those members choose the remaining members. Eight of New York’s commissioners are appointed by politicians from each party, and the last two members are chosen by the appointees and must not have been a registered Democrat or Republican in the past five years. Utah’s independent commission has members chosen by state legislators, with the governor choosing the commission chair, though all commission members must not be affiliated with any political party nor have voted in any political party’s primary elections in the past five years. Michigan’s independent commission members were selected randomly from a pool of qualified applicants.

¹³Only Idaho (6 members) and Washington (4 members and 1 non-voting member) have perfectly balanced (by partisanship) independent commissions, and in these cases some bipartisan cooperation is needed for them to successfully create district maps. Researchers have illustrated that, in practice, balanced commissions may produce incumbent-protecting gerrymanders (McDonald, 2004).

finds that independent commissions insulate incumbent legislators to the same degree that party-controlled legislative redistricting does, suggesting that independent commissions may not be as neutral as many suppose (Henderson, Hamel and Goldzimer, 2018).

The effectiveness of independent commissions also hinges crucially on who staffs them. A Brennan Center report notes that “the strength and independence of the [commissioner] selection process was, by far, the most important determinant of a commission’s success” (*Redistricting Commissions: What Works*, 2018). Even with an independent staffing process, however, independent commissions do not quell the partisan anger over redistricting controversies. Those who have studied independent commissions note that “the decisions of such commissions may generate partisan rancor comparable to what we see from states where one party entirely controls the redistricting process and engages in a partisan gerrymander” (Miller and Grofman, 2013, p. 648), and that “[o]ften, commissioners have strong common prior beliefs about the likely partisanship of the tiebreaker, and therefore balk at compromise during initial negotiations. Once chosen, the tiebreaker then sides with one of the parties and a partisan plan is adopted” (McDonald, 2004, p. 383). Similarly, the Brennan Center report notes that “states that used a tiebreaker model popular in earlier reforms experienced much lower levels of satisfaction, mainly because the tiebreaker tended to end up siding with one party or the other, resulting in a winner-take-all effect” (*Redistricting Commissions: What Works*, 2018).

Last of all, the establishment of an independent redistricting commission is not a realistic options for many citizens. According to the National Conference of State Legislatures, slightly more than half of U.S. states do not have a legislative process allowing statutes or state constitutional amendments by initiative.¹⁴ Of the 24 states that do, nine have independent redistricting commissions already. The states with the most intense partisan gerrymandering do not have an initiative process, and legislatures in those states are also very unlikely to voluntarily relinquish authority over redistricting to an independent commission. For example in Maryland, North Carolina, Pennsylvania, Texas, Virginia, and Wisconsin, voters cannot feasibly establish non-partisan independent redistricting commissions because these states do not have an initiative process.

These problems with the creation and effectiveness of commissions show that independent redistricting commissions do not offer a silver-bullet solution to partisan gerrymandering in most states. Regardless of who draws the lines, many states have instead looked to the courts for relief from partisan gerrymandering. Existing legal remedies, however, have met with several obstacles.

¹⁴See <https://www.ncsl.org/research/elections-and-campaigns/initiative-and-referendum-processes.aspx>.

One of the largest obstacles to effective judicial intervention is that courts lack effective guidelines and standards to adjudicate partisan gerrymandering litigation. At a minimum, courts must decide (1) how to measure and evaluate partisan gerrymandering,¹⁵ (2) how to compare multiple maps,¹⁶ and (3) at what threshold there is too much partisan gerrymandering. But none of these three issues have been settled. Any solution needs to cut through the “sociological gobbledygook” in a way perceived as non-partisan and legally sound (quoting Chief Justice Roberts during Oral Arguments for *Gill v. Whitford*, October 3, 2017). Additionally, the Supreme Court’s decision *Rucho v. Common Cause* (2019) effectively barred the federal judicial from future intervention in partisan gerrymandering litigation. This has left state courts to adjudicate partisan gerrymandering claims, and relegates the judicial intervention option to a much less effective state-by-state approach.

State Supreme Courts have recently struck down redistricting plans for being unconstitutional partisan gerrymanders (according to state law). Most recently, both Ohio’s and New York’s legislature-passed redistricting plans were struck down as impermissible partisan gerrymanders (in Ohio, *League of Women Voters of Ohio v. Ohio Redistricting Comm.*, 2022; in New York, *Harkenrider v. Hochul*, 2022). In Florida, the courts based their decision in *League of Women Voters v. Detzner* (2015) on a “Fair Districts” amendment prohibiting partisan gerrymandering, which voters had previously added to the state constitution through a popular initiative.¹⁷ In both Pennsylvania (*League of Women Voters v. Commonwealth of Pennsylvania*, 2018) and North Carolina (*League of Women Voters v. Rucho*, 2018; reconsidered and reaffirmed 2019), the courts relied on more generic language in the state constitutions ensuring “free elections.”¹⁸ Only some states, however, have existing state laws

¹⁵There exists no legal consensus on how to best identify instances of partisan gerrymandering, despite a plethora of new partisan gerrymandering metrics developed in the past few decades. Since the Supreme Court’s decision in *Vieth v. Jubelirer* (2004), finding a standard to judge partisan gerrymandering has remained a challenge. Measures like the Efficiency Gap (Stephanopoulos and McGhee, 2015), the Mean-Median Difference (McDonald and Best, 2015), and Partisan Fairness (King and Browning, 1987; Grofman and King, 2007) have grown increasingly common, but courts have not settled on one. Each approach has some mix of desirable and undesirable features (Stephanopoulos and McGhee, 2018).

¹⁶Courts sometimes rely on simulated or counterfactual election results in order to create a distribution of possible maps against which the actual or proposed redistricting plans can be compared. Current computational limitations make it impossible to create the full distribution of possible maps, so simulations rely on creating a representative sample of possible maps as a baseline (Cho and Liu, 2016). Experts continue to debate whether particular simulation methods create a “true” distribution of possible maps, and the courts must navigate among competing methods, (Cirincione, Darling and O’Rourke, 2000; Altman and McDonald, 2011; Chen and Rodden, 2013; Chen and Cottrell, 2016; Cho and Liu, 2016; Magleby and Mosesson, 2018; Duchin, 2018a; Fifield et al., 2020).

¹⁷Amendment 5: “Legislative districts or districting plans may not be drawn to favor or disfavor an incumbent or political party.” https://web.archive.org/web/20101208155829/http://projects.palmbeachpost.com/yourvote/ballot_question/florida/2010/amendment-5-and-6-2010/.

¹⁸For North Carolina, the courts concluded that the redistricting process was not consistent with a broad reading of Section 10 of the North Carolina State Constitution, which states that “All elections shall be

or constitutional amendments that provide a legal basis to limit partisan gerrymandering. For example, while 30 states have some version of “free election” clauses in their constitutions, only 18 also require “equal” or “open” elections.¹⁹ Notably, of 41 the states that do *not* have an independent redistricting commission, 18 of them have neither a “free election” clause nor an “open” or “equal” provision in their state constitutions.

Overall, the existing attempts to fix partisan gerrymandering have resulted in a patchwork of solutions with highly limited effectiveness. Many voters live in states that cannot feasibly implement the commission-based solutions that have had success in other states, and judicial intervention is limited to state courts which often do not have constitutional provisions that allow them to reduce partisan gerrymandering. And lack of citizen-led initiative procedures in many states makes it impossible for voters to solve this issue without the help and approval of their partisan state legislators. This insufficient patchwork leaves citizens with little recourse to address the degradation of representation in their states caused by partisan gerrymandering.

A.1.2 Other Proposed Solutions

Some of the most promising alternative solutions to gerrymandering draw inspiration from the cake-cutting problem; how do two people perform the fair division of a piece of cake without the need of third-party intervention? The solution is to arbitrarily choose one as the first mover; she divides the cake and then the second-mover may choose between either of the pieces. This logic, applied to geography, has inspired several redistricting proposals.

One proposal is to have an independent third party divide the state into two and then each party negotiates over who gets to redistrict one section of the state (Landau, Reid and Yershov, 2009). The parties each independently redistrict their agreed-upon parts of the state. Combining the two sections results in a final map. In another proposal, each of two parties alternate back and forth drawing district maps (Pegden, Procaccia and Yu, 2017). Termed “I-cut-you-freeze,” the protocol involves a back and forth where one party draws a map, the other party freezes in place one district from that map, and then redraws a new district map for the remaining area in the state. The players alternate between “cutting” and “freezing” until producing a full map.

Neither of these approaches has seen any take-up in the real world. The difficulties of implementing these solutions in practice are several-fold. In the first proposal, the process requires a neutral third party to take the initial step of dividing the state into two parts,

free.” Similarly in Pennsylvania, the courts found that the challenged map violated the “Free and Equal Elections” Clause (Article 1, Section 5) of the Pennsylvania State Constitution.

¹⁹See <https://www.ncsl.org/research/redistricting/free-equal-election-clauses-in-state-constitutions.aspx>

which has proven to be a stumbling block in the past (Landau, Reid and Yershov, 2009). Both approaches abstract from real-world geographies and do not place constraints on how voters are assigned to districts (Landau, Reid and Yershov, 2009; Pegden, Procaccia and Yu, 2017). Furthermore, because they involve multiple stages of bargaining between the parties, these approaches are impractical to simulate in real-world contexts using actual geographies and voter rolls. Thus, lack of information about implementation and potential results with real electoral geography and population information make it unlikely that decision makers would adopt these protocols.

Other researchers have proposed a protocol with a similar “I-cut-you-freeze” style, but with an explicitly spatial addition to the process (Ely, 2019). The first party draws a full set of districts. Any district that is convex²⁰ is locked into place. However, the second party has the ability to redraw any non-convex districts so that they are convex. This two-stage process assures the creation of a map without misshapen districts. However, this proposal also meets with some practical issues. First, in some states it is likely not possible to meet equal population requirements while also maintaining convex districts. Second, even states with convex districts can be extraordinarily biased in favor of one party, depending on the geographical distribution of voters (Alexeev and Mixon, 2019). A final proposal involves a method that divides the state in two and allows each party to redistrict their half, with the additional constraint that each party draws a share of districts roughly proportional to the party’s statewide vote share in the last Congressional election (Brams, 2020). In essence, this method seeks to let the parties create their fair share of gerrymandered districts.

All of the proposed solutions involve either a third-party neutral arbiter, are difficult to implement in practice, or have uncertain outcomes that are hard if not impossible to predict computationally. Our Define-Combine Procedure is designed to address these difficulties. Unlike the already-implemented fixes or other proposed solutions, DCP does not require an independent third party to ensure that districts are fair, and it is possible to predict the outcomes of DCP using simulations. An additional benefit is that DCP could be combined with many existing solutions or proposals - for example, by having an existing redistricting commission use DCP to create legislative maps for a state, or by first freezing certain districts and then using DCP on the rest to produce a final map. This represents a substantial step towards implementing a process-based solution to the problem of partisan gerrymandering.

²⁰For a district to be convex, a straight line can be drawn between any two points in the district and all of the line remains inside the district.

Table A.1: Redistricting Procedures by State

Legislature Only (31)			Legislature-Involved Commissions (10)			Independent Commissions (9)
			Advisory (4)	Political (3)	Backup (3)	
Alabama	Minnesota	Pennsylvania	Iowa	Hawaii	Connecticut	Alaska*
Arkansas	Mississippi	Rhode Island	Maine	New Jersey	Indiana	Arizona
Delaware*	Missouri	South Carolina	Utah	Virginia	Ohio ¹	California
Florida	Nebraska	South Dakota*	Vermont*			Colorado
Georgia	Nevada	Tennessee				Idaho
Illinois	New Hampshire	Texas				Michigan
Kansas	New Mexico	West Virginia				Montana
Kentucky	North Carolina	Wisconsin				New York
Louisiana	North Dakota*	Wyoming*				Washington
Maryland	Oklahoma					
Massachusetts	Oregon					

Source: Justin Levitt, *All About Redistricting: Who Draws the Lines?*, website: <https://redistricting.lls.edu/redistricting-101/who-draws-the-lines/>, along with authors' classifications.

Advisory Commission: Assists the legislature in drawing the maps, but the legislature has the ultimate power to approve or alter the final district maps; **Political Commission:** Legislature as a whole isn't officially involved, but the members of the commission are politicians or elected officials; **Backup Commission:** Steps in if the legislature does not pass a districting plan by a certain deadline—these backup commissions vary in their composition and procedures as well, but are almost always comprised of politicians (governor, secretary of state, state legislators, or members selected by political leadership); **Independent Commission:** Commissions that have no politicians or elected officials on them, and whose maps are not subject to legislature approval.

¹ Ohio has a seven-person politician commission that draws lines if the legislature does not create a map with three-fifths legislature support. The commission's map must have the support of two minority party legislators, who are required to be on the commission. If both the legislature and politician commission fail to enact a map, the majority party can adopt a map without minority support that would last for four years. <https://redistricting.lls.edu/state/ohio/>

*State only has one U.S. House district; state legislative redistricting authority used for classification.

A.2 The Define-Combine Procedure

Suppose a state (or city, school district, or other entity engaged in redistricting) with population P needs to be divided into N contiguous single-member districts with equal population P/N . Elections in the state are contested by two parties, A and B . We assume for simplicity that all people in the state vote in all elections, and their voting decision is based solely on their personal partisan preference; the makeup of their district and the candidates who run have no impact on their vote choice.²¹ Let v_A be the number of votes in the state for Party A, and $v_B = P - v_A$ be the number of votes for Party B. For each district d , let v_{dA} and v_{dB} be the number of votes in the district for each party. A districting map M is a set of N districts, and each district is itself a set of P/N specific voters. Thus, for this framework there is a *finite* set of possible maps \mathcal{M} (though it grows extraordinarily large as the population P increases). And, for any map, both parties can foresee the number of votes they will receive in each district and the number of districts that they will win.

We first assume that both parties are seat-maximizers; their goal is to win as many of the N seats as possible in the next election²²:

$$U_A(M) = \frac{\sum_{d=1}^N s_d}{N}$$

where

$$s_d = \begin{cases} 1 & \text{if } v_{dA} > v_{dB} \\ 1/2 & \text{if } v_{dA} = v_{dB} \\ 0 & \text{if } v_{dA} < v_{dB} \end{cases}$$

and the utility of Party B is $U_B(M) = 1 - U_A(M)$. Both parties are risk-neutral; they are indifferent between winning one district and tying in two districts (with a 50% chance of winning the election in each). U_i is equivalent to the percentage of seats won by party i .²³

²¹The levels of turnout and vote choice are not essential to this game. The most important things are (1) that districting does not affect vote choice, and (2) each party can anticipate which districts they will win and lose (or the probability of winning and losing). Additionally, our baseline example abstracts away from the incumbency advantage, but it could be incorporated as well into voter decision-making.

²²Later in the paper we consider instances where parties may have other goals. For example, parties might seek to maximize bias, or they might have preferences over the level of responsiveness in the plan (Katz, King and Rosenblatt, 2020).

²³Our characterization here is of a one-shot game. Parties and candidates need not account for uncertainty in future elections or shifting voter preferences over time.

The Unilateral Redistricting Process

We consider two methods of drawing district maps. In the first method, one party has unilateral control of the process. Given a set of potential valid maps \mathcal{M} , when party i draws the maps, it will select a map $\widetilde{M}_i \in \mathcal{M}$ that maximizes U_i .²⁴ Under this method of redistricting, we should expect that, where possible, party i will maximize partisan advantage relative to party j by strategically cracking and packing party j 's voters to minimize the seats won by party j . In many cases, it will be possible for party i to win a substantially larger share of the seats than its statewide vote share.

This method, which we will call the “unilateral redistricting process” (URP), approximates the redistricting process in states where one party controls redistricting for a given map. The party in control seeks to maximize the number of seats it will win in the next election. While other factors, such as incumbency or vote margins in close seats, factor into districting decisions, gerrymandering for partisan gain is a primary objective for the party in power. For example, when independent experts have used map-drawing software to simulate thousands of possible maps in a state, the observed maps in states where redistricting is controlled by a single party often appear to be among the most partisan possible.²⁵

The Define-Combine Procedure

The second method we consider is our own innovation, the two-stage Define-Combine Procedure. In this model, the power to draw the map is divided between the two parties (i.e., the players in the game are Party A and Party B), but in each stage of the process one party acts unilaterally.

Suppose Party A acts in the first stage as the “Definer,” and Party B acts in the second stage as the “Combiner.” The game proceeds as follows:

1. Party A defines a set of $2N$ contiguous, equally populated districts. To avoid confusion with the following stage, we refer to these districts as *subdistricts*.²⁶
2. Party B creates the final map of N districts by combining pairs of 2 contiguous subdistricts.

²⁴In practice, there will be a large number of maps in \mathcal{M} (even in some of the simple examples here, there may be millions or billions of legal maps). Thus, \mathcal{M} does not have to be the complete set of feasible maps, but rather a subset of all feasible maps in which U_i varies.

²⁵Research on redistricting in Florida (Chen and Rodden, 2015), Maryland (Cho and Liu, 2016), Wisconsin (Chen, 2017), Virginia (Magleby and Mosesson, 2018), and Pennsylvania (Duchin, 2018*b*) find that chosen plans in states with URPs are extreme outliers as compared to the set of simulated possible maps.

²⁶A related game could involve defining kN subdistricts where k is a positive integer greater than 2.

Party A moves first and so has a strategy profile consisting of a selection of a map $M_A \in \mathcal{M}$, the set of all maps with $2N$ valid districts.²⁷ Party B combines subdistricts to create a map $M_B \in Q(M_A)$, where $Q(M_A)$ is the set of valid groupings of the subdistricts in M_A .²⁸

Party B, the second mover, will select a best-response to any proposed set of subdistricts. The strategy σ_B is a mapping from the set of valid groupings of sub-districts to a single map, \widehat{M}_B , such that

$$\widehat{M}_B \equiv \sigma_B(M_A) \in \operatorname{argmax} U_B(Q(M_A)).$$

Because voters themselves are indivisible and the districts in this setup consist of sets of voters, any game has a fixed number of possible districts. Also, the second-mover knows what the first mover has chosen to do. In a finite extensive game with perfect information, such as DCP, there exists a subgame perfect equilibrium (Osborne et al., 2004, p. 173). Furthermore, it can be solved using backward induction. For the map M_A selected by Party A, Party B will examine all of the possible maps it could draw, $Q(M_A)$. From these possibilities, it will select the map $\widehat{M}_B \in Q(M_A)$ that maximizes the percentage of seats won by Party B. For every possible map $M_A \in \mathcal{M}$, Party A can anticipate what ultimate map M_B Party B would draw. Therefore, it selects the map \widehat{M}_A that maximizes the percentage of seats won by Party A subject to Party B’s best-response pairings, with payoff $U_A(\sigma_B(\widehat{M}_A))$.²⁹

This procedure reduces partisan gerrymandering by limiting the efficacy of the most important strategy for gerrymandering—packing the opposing party’s voters into as few noncompetitive districts as possible (Friedman and Holden, 2008). Consider a state, to be divided into five districts, with voters evenly divided between the two parties. If Republicans can draw one district with 80% Democrats, they could likely draw the four remaining districts to produce Republican wins by a narrow margin. However, this “packing” strategy fails under DCP—if Republicans draw two 80% Democratic subdistricts, then in the combine stage Democrats would refuse to combine the subdistricts and instead pair them with neighboring Republican subdistricts to win the two resulting districts. While specific political geographies can still enable packing in some places, the ability to do so is significantly constrained.

²⁷Valid districts are contiguous (except where required by geographic features such as islands) and have equal population. We do not impose any compactness, split geography, or other restrictions, but such limitations could be included here. The one exception to this is that valid districts may not include “donuts,” where one district entirely encircles another.

²⁸To rule out a Definer drawing maps with one or fewer valid Combine stage responses, the procedure can also require a Define stage map M_A with more than one valid Combine state map in $Q(M_A)$. For evaluation of real-world state and congressional district maps in this paper, we require $|Q(M_A)| > 10$. We have found that this restriction rarely binds in practice.

²⁹Equivalent to minimizing the percentage of seats won by opposing Party B.

A.3 A Simple Example: Iowa

A simple example illustrates the DCP framework. Consider a (simplified) map of Iowa, as in Figure A.1, with 30 equally-populated precincts and an overall statewide vote share of 50% for the Democrats and 50% for the Republicans.³⁰

Suppose redistricting requires that the state be divided into five equally populated contiguous districts (which could occur if Iowa were to gain an additional Congressional district due to reapportionment). Given that each of the thirty precincts has the same population, the state must be divided into five districts of six precincts each. Under these assumptions, there are 27,250 possible maps.

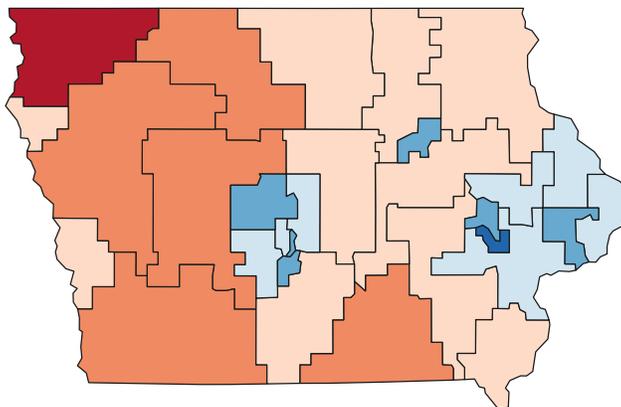
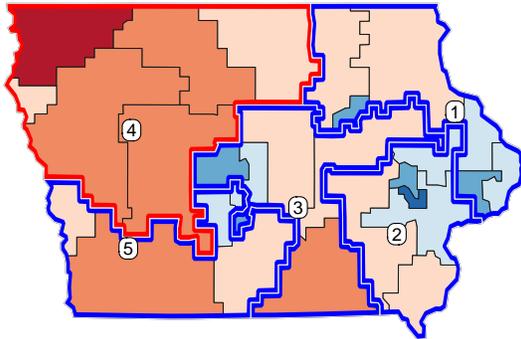


Figure A.1: A simplified map of Iowa with 30 equally populated precincts. Dark red (blue) precincts denote higher Republican (Democratic) vote shares.

If Democrats re-draw the map unilaterally and maximize the number of seats won, they can construct a map where they win four seats and Republicans win one seat. Figure A.2(a) shows one such map (out of many equivalent possibilities) drawn by the Democrats unilaterally. Districts won by Democrats (Republicans) are denoted with a blue (red) outline. In this example, Republicans are packed into District 4, and Democrats win majorities in Districts 1, 2, 3, and 5. Conversely, as illustrated in Figure A.2(b), when Republicans act unilaterally, the opposite result is possible; Republicans win four seats, and Democrats win one.

³⁰We used a 2016 Iowa precinct shape file to generate the map, simplifying by creating thirty precincts with equal population. Vote share in each precinct is based on 2016 Presidential election totals, with a uniform swing applied to each precinct so that the statewide average is 50% for each party. Note that two-party Democratic vote share in Iowa averaged over the 2012 and 2016 presidential elections is 49%, so this scenario does not stray far from reality.

(a) Best map for Democrats
Dems. win 4–1



(b) Best map for Republicans
Reps. win 4–1

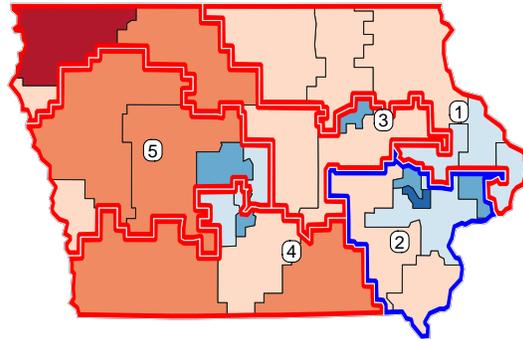
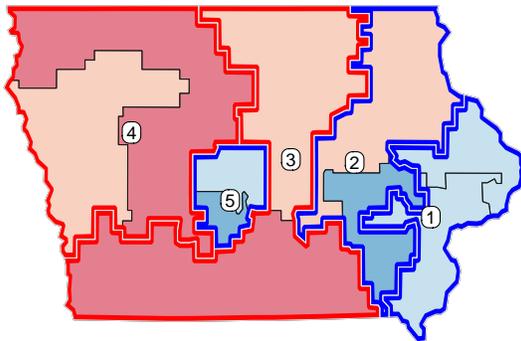


Figure A.2: Examples of Valid Maps for Iowa (simplified).

(a) Dems. Define; Reps. Combine



(b) Reps. Define; Dems. Combine

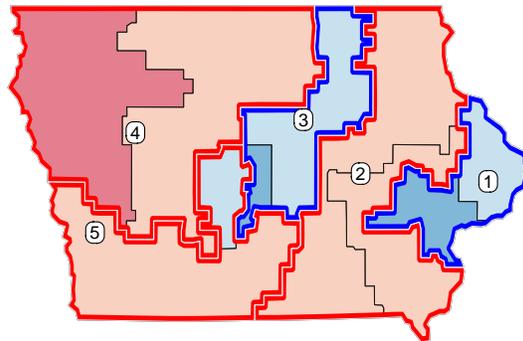


Figure A.3: Define-Combine Procedure Results for Iowa (simplified)

We now apply DCP to this simple example. In the first stage, the defining party will draw a map consisting of ten contiguous subdistricts, each with three precincts. There are 7,713 valid divisions of this map of Iowa. In the second stage, the combining party selects contiguous pairs of subdistricts to create the final district map. The number of possible combinations in the second stage varies based on the subdistricts defined in the first stage. In this example, the number of combinations varies from 2 to 20 possibilities.

For any possible proposed map (i.e., for each sub game), the defining party analyzes the resulting combinations and determines the best-response for the combining party. The defining party chooses the map that minimizes the utility that the combining party gets from making the optimal pairing in the sub game. Given the distribution of voters in our running example, Figure A.3 presents the results of DCP for this simple example if the Democrats go first, on the left, and if the Republicans go first, on the right. In this example,

there are multiple equilibria and we present just one graphically. The defined subdistrict plan selected by the Democrats results in three seats for Democrats and two seats for Republicans. The Republicans cannot choose any other combination of these subdistricts to improve the outcome. Similarly, if the Republicans move first, then in equilibrium the Republicans win three seats and the Democrats win two seats. Thus, DCP reduces the advantage conferred to the map-drawer/definer. Under URP, there is a three-seat (of five total seats) difference in partisan outcomes depending on who controls the process, while under DCP there is only a one-seat difference depending on who draws the define-stage map.³¹

³¹We will denote this difference based on first-mover status as δ . So, for URP $\delta_5^U = 0.8 - 0.2 = 0.6$ and for DCP $\delta_5^D = 0.6 - 0.4 = 0.2$.

A.4 Simulation Details

We use two different applications of shortburst algorithms to generate our results. First, we apply a simple implementation of shortburst to find the map that maximizes seats won by each party under unilateral redistricting. For each state, chamber, and party, we run 10 separate sets of simulations. In each simulation, we (1) generate a random starting map, (2) run the shortburst algorithm 2000 times, with 10 maps generated per burst, and (3) save the final map. We use a scoring function that maximizes first the number of seats won by the party, and second the party’s vote share in the next closest district.³² Across 10 separate chains for each state, we ultimately simulate and score 200,000 plans, selecting the most extreme result for each party.³³

Second, we use a “nested shortburst” algorithm to simulate the Define-Combine Procedure. For each state, chamber, and party we run 10 separate sets of simulations. In each simulation, we (1) generate a random starting map of $2N$ districts, (2) run the shortburst algorithm 250 times, with 10 maps generated per burst. Importantly, our scoring function for these maps involves a second iteration of the scoring algorithm. Instead of calculating the seats won, we use a different shortburst algorithm to generate different pairings of districts and maximize the seats won by the second party. In particular, for every first stage map generated, we next run 50 shortbursts of the pairing algorithm and generate 20 different combinations in each burst—generating different pairings of districts and searching for the instance that maximizes seats won by the second party. Thus, the second implementation of the algorithm maximizes seats for the second party across 1,000 possible maps for each define-stage map generated, and then the first algorithm uses that as the score when maximizing seats for the first party.³⁴ This approach gives both parties the benefit of actively trying to maximize their own advantage, rather than looking at a random set of proposals and choosing the best option.³⁵

The full R code and data to run our simulations is included with the replication data.

³²We find that this scoring function more efficiently finds extreme gerrymanders than scoring based on seats won alone.

³³If a map is generated where the districting party is able to win every district in the state, the shortburst algorithm is terminated and additional maps are not simulated.

³⁴For states with seven or fewer congressional districts, we enumerate *all* possible combinations in the second stage and use the map that yields the most seats for the combining party as the score for the first stage.

³⁵An earlier version of this manuscript used a random sampling approach rather than shortburst. Our results with the shortburst algorithm are more extreme — parties under both processes are able to win more seats, but the overall pattern of the results is similar.

A.5 Measurement of Redistricting Plans

Consider an electoral system with seats-votes function $S_M(\nu_1, \dots, \nu_N)$ for a map M , which takes as an input district-level vote shares ν_1, \dots, ν_N and yields as an output a seat share.³⁶ The state-wide vote share V is the average of district-level vote shares (importantly, elections with identical state-wide average vote share V but different realizations of ν_1, \dots, ν_N could result in different winning candidates). Conditioning on a state-wide vote share V , we can find the average seat share by taking the expected value of the function, over the joint distribution for ν_1, \dots, ν_N —e.g., $E(S_M(\nu_1, \dots, \nu_N) | V) = S_M(V)$. Note that $S_M(0.5) \neq 0.5$ indicates an electoral system with partisan bias, which could be due to inherent geographic bias (Chen and Rodden, 2013), gerrymandering, or both.

The definer’s advantage depends on how the seat share changes when Party A unilaterally redistricts compared to when Party B unilaterally redistricts, for example in a 50-50 state $\delta_{.5}^U = S_{\widetilde{M}_A}(0.5) - S_{\widetilde{M}_B}(0.5)$. Similarly, $\delta_{.5}^D = S_{\widehat{M}_A}(0.5) - S_{\widehat{M}_B}(0.5)$ indicates the definer’s advantage under DCP. A large positive value for $\delta_{.5}^U$ indicates that the party controlling the URP can reap a significant electoral advantage through gerrymandering; a large positive value for $\delta_{.5}^D$ indicates that the definer or first-mover in DCP can reap a significant electoral advantage. A negative value indicates a second-mover advantage or combiner’s advantage. Redistricting procedures that minimize the absolute value of this quantity tend towards providing both parties equal treatment.

Partisan gerrymandering may pose a problem for an electoral system if there exist large differences in seats won *depending on which party controls the redistricting process*. Consider a state with unilateral redistricting and vote share $V = 0.5$; suppose Party A wins 75% of the seats if it draws the map whereas Party B wins 70% of the seats if it draws the map. Such a map appears to confer a large partisan advantage to whichever party controls redistricting; 45% of seats in the legislature change hands depending on the party that draws the map. Alternatively, suppose that Party A wins 52% of the seats if it draws the map, and Party B wins 50% of the seats if it draws the map. In this case, partisan gerrymandering represents a smaller problem, with a swing of 2 percentage points depending on the party controlling the process.

Second, to determine how a redistricting procedure affects partisan bias, we directly compare seat shares for each procedure when the two parties evenly split votes, which we estimate by a uniform swing. If $|S_{\widetilde{M}_A}(0.5)| > |S_{\widehat{M}_A}(0.5)|$ (where \widetilde{M}_A denotes the optimal map for Party A from unilateral redistricting, and \widehat{M}_A the optimal map for Party A from DCP, then DCP reduces bias due to redistricting as compared to URP.

³⁶Notation used here is similar to Katz, King, and Rosenblatt (Katz, King and Rosenblatt, 2020).

Third, to determine a method's deviation from proportionality, we simply compare the Democratic seat share under a plan to the vote share under the plan. Thus, $|S_{\widetilde{M}_A} - V| > |S_{\widetilde{M}_A}(0.5) - V|$ denotes a case where the URP map for Party A deviates from proportionality more than the DCP map.

Fourth, to determine a method's deviation from a fairness target, we first define such a target by following a similar definition to the one in Benadè, Procaccia and Tucker-Foltz (2021), which proposes a geometric target that is the average of the outcomes under the worst and best possible map partitions in a state for a given party. In our framework, this is equal to $\theta = \frac{S_{\widetilde{M}_A}(V) - S_{\widetilde{M}_B}(V)}{2}$. Under this framework, a map M performs better than a map N when $|S_M(V) - \theta| < |S_N(V) - \theta|$.

A.6 Additional Simulation Results

A.6.1 Simulation Tables

State	Seats	R Alone	R Then D	D Then R	D Alone
AL	7	0	0	1	3
AR	4	0	0	1	1
FL	28	4	12	10	19
GA	14	2	6	6	11
IL	17	9	11	11	14
KS	4	0	1	2	2
KY	6	0	0	1	2
LA	6	0	1	1	3
MA	9	9	9	9	9
MO	8	0	2	3	5
MS	4	0	0	1	2
NE	3	0	0	1	1
NM	3	1	2	3	3
NV	4	1	2	3	3
OK	5	0	0	0	1
OR	6	2	3	5	6
RI	2	2	2	2	2
SC	7	0	1	3	4
TN	9	0	1	2	4
TX	38	6	19	17	28
WV	2	0	0	0	0

Table A.2: **Legislatures Redistricting:** Simulation Results

State	Seats	R Alone	R Then D	D Then R	D Alone
CT	5	4	5	5	5
MD	8	5	6	8	8
MN	8	2	4	4	7
NC	14	1	5	6	11
NH	2	1	2	2	2
NY	26	16	20	22	26
PA	17	4	7	8	12
VA	11	4	7	8	10
WI	8	1	3	4	6

Table A.3: **Court or Special Master Redistricting:** Simulation Results

State	Seats	R Alone	R Then D	D Then R	D Alone
AZ	9	1	4	4	8
CO	8	1	4	5	7
MI	13	2	6	5	10

Table A.4: **Independent Commission Redistricting:** Simulation Results

State	Seats	R Alone	R Then D	D Then R	D Alone
HI	2	2	2	2	2
IA	4	0	0	2	3
IN	9	0	1	2	4
ME	2	1	2	2	2
NJ	12	6	8	10	12
OH	15	0	5	4	11
UT	4	0	0	1	2

Table A.5: **Political Redistricting:** Simulation Results

A.6.2 Assessing Definer’s Advantage for DCP versus Enacted Plans

Making a comparison of DCP’s results for Definer’s Advantage to the Definer’s Advantage for actual, enacted plans presents several complications because we do not observe the counterfactual of how the minority party (e.g., the party not controlling redistricting in a state) would behave when enacting their plan. In this section, we impose an additional assumption about the behavior of the party in order to be able to estimate Definer’s Advantage for enacted plans.

To compare in terms of Definer’s Advantage, we first interpolate the Democratic seat share for the counterfactual where the current minority party controlled the redistricting process. We find the difference between the theoretical maximum number of seats the current party in power could win (estimated from URP) and the number of seats the party is projected to win based on the enacted map. We then impute the difference for the counterfactual where the minority party redistricted. The maintained assumption is that each party would deviate from its most extreme map by the same margin.

Under this assumption, controlling the redistricting procedure in each state leads to a swing between the parties of 17.6% of seats on average. Compared to this benchmark, DCP reduces the advantage conferred to the redistricting party by roughly half.

A.6.3 Assessing Performance of DCP versus Enacted Plans by State Redistricting Procedure

This Appendix Section describes simulated DCP results as compared to enacted plans, depending on the redistricting procedure implemented in each state. DCP achieves improvements compared to adopted plans, particularly with regard to reducing the extent of deviation from the geometric fairness target. Table A.6 presents the metrics for DCP versus the other benchmarks broken out by state redistricting procedure. DCP achieves reductions in deviation from the fairness target compared to adopted plans across all of the redistricting procedure types. We view this metric as one of our most informative since the geometric fairness target essentially captures the seat share for a map that grants neither party an advantage due to unilateral redistricting. It also accounts for the natural advantages one party may have in a state due to a state’s geography and how it interacts with the spatial distribution of voters. Along this dimension, DCP performs better better than URP and better than adopted maps regardless of the method used (court or special master, commissions, legislature, etc.). One reason this may depart from the results for partisan bias is that we calculate the fairness target for all states (rather than just the subset used for partisan bias calculations).

Metric	Category	URP	DCP	Adopted Plan
Definer’s Advantage	Court or Special Master	0.495	0.081	<i>0.121</i>
Definer’s Advantage	Independent Commission	0.417	0.052	<i>0.125</i>
Definer’s Advantage	Legislature	0.468	0.097	<i>0.204</i>
Definer’s Advantage	Political	0.562	0.146	<i>0.271</i>
Deviation from Proportionality (Actual State Party Control)	Court or Special Master	0.257	0.144	0.107
Deviation from Proportionality (Actual State Party Control)	Independent Commission	0.279	0.179	0.172
Deviation from Proportionality (Actual State Party Control)	Legislature	0.265	0.142	0.217
Deviation from Proportionality (Actual State Party Control)	Political	0.296	0.216	0.268
Deviation from Fairness Target (Actual State Party Control)	Court or Special Master	0.247	0.056	0.066
Deviation from Fairness Target (Actual State Party Control)	Independent Commission	0.208	0.036	0.062
Deviation from Fairness Target (Actual State Party Control)	Legislature	0.234	0.062	0.105
Deviation from Fairness Target (Actual State Party Control)	Political	0.281	0.094	0.135
Partisan Bias	Court or Special Master	0.302	0.060	0.040
Partisan Bias	Independent Commission	0.350	0.067	0.058
Partisan Bias	Legislature	0.287	0.057	0.100

Table A.6: **DCP Performance Versus Alternatives, by Redistricting Procedure:** Breaking results out by redistricting procedure, this table reports the performance of the Define-Combine Procedure (DCP) as compared to unilateral (URP) and adopted plans along several different metrics. The cells in italics are based on interpolated values for Definer’s Advantage. This quantity is calculated by interpolating the seat share under the scenario where the opposing party held control over the redistricting process. The partisan bias calculations do not include any states with a “Political” redistricting process. Partisan Bias is calculated only for states with 2020 Democratic Presidential Vote Share between 45% and 55%.

For partisan bias, DCP clearly reduces bias compared to the adopted plans in states with legislature drawn maps (from 10% to 5.7%). In states with maps drawn by independent commissions, courts or special masters, DCP achieves comparable or slightly higher levels of bias. Importantly, these partisan bias simulations are for a limited set of states, impose a different set of constraints, and assume parties are seat-maximizers—so this is not a perfect comparison. Nonetheless, given that DCP achieves comparable results, we think it suggests that in practice DCP would perform comparably (or better) in terms of partisan bias to maps produced by commissions, courts or special master—all without requiring agreement on an independent third-party arbiter or cooperation between the parties.

Figure A.4, which plots Appendix seat share for a state and method (y-axis) against the geometric fairness target, summarizes these state-by-state results. We focus here on states with at least four congressional districts and a URP midpoint between 0.3 and 0.7, as redistricting matters primarily in states with some degree of support for more than one political party. Among all states that did not ultimately default to maps drawn by a court or special master, the outcome from DCP is closer to a proportional one than the actual outcome achieved in the state. The figure illustrates how the actual redistricting outcomes, represented by the point where a line with an arrow originates, are less proportional than the DCP outcome in states where redistricting occurred by independent commissions, legislatures, or politician-controlled commissions. The two exceptions we observe are in states with plans drawn by Courts or Special Masters; specifically, we observe that DCP produces a slightly more majoritarian outcome in Virginia and North Carolina than the maps produced for these states, which are both proportional to the geometric fairness target.

A.6.4 State-by-State Shortburst Comparisons, 2020 Presidential Election Democratic Vote Share

Figures A.5, A.6, and A.7 compare DCP outcomes to several other benchmarks for each state. The benchmarks include simulated URP outcomes, actual seat projections based on the post-2020 redistricting cycle, and 2020 Presidential election vote share. We group results by redistricting process: Legislatures, Independent Commissions, Politician-controlled,³⁷ and Court or Special Master. Several key patterns emerge. In some states, the partisan composition tilts so far in the direction of one party that it makes no difference who draws the map or which process is used (see Hawaii, Idaho, Massachusetts, Montana and Rhode Island). However, in states with mixed partisanship, notable differences emerge between the results from DCP and the other benchmarks.

³⁷These include states where state legislatures retain influence over the redistricting process but may have backup commissions, advisory commissions, or commissions composed by or of politicians

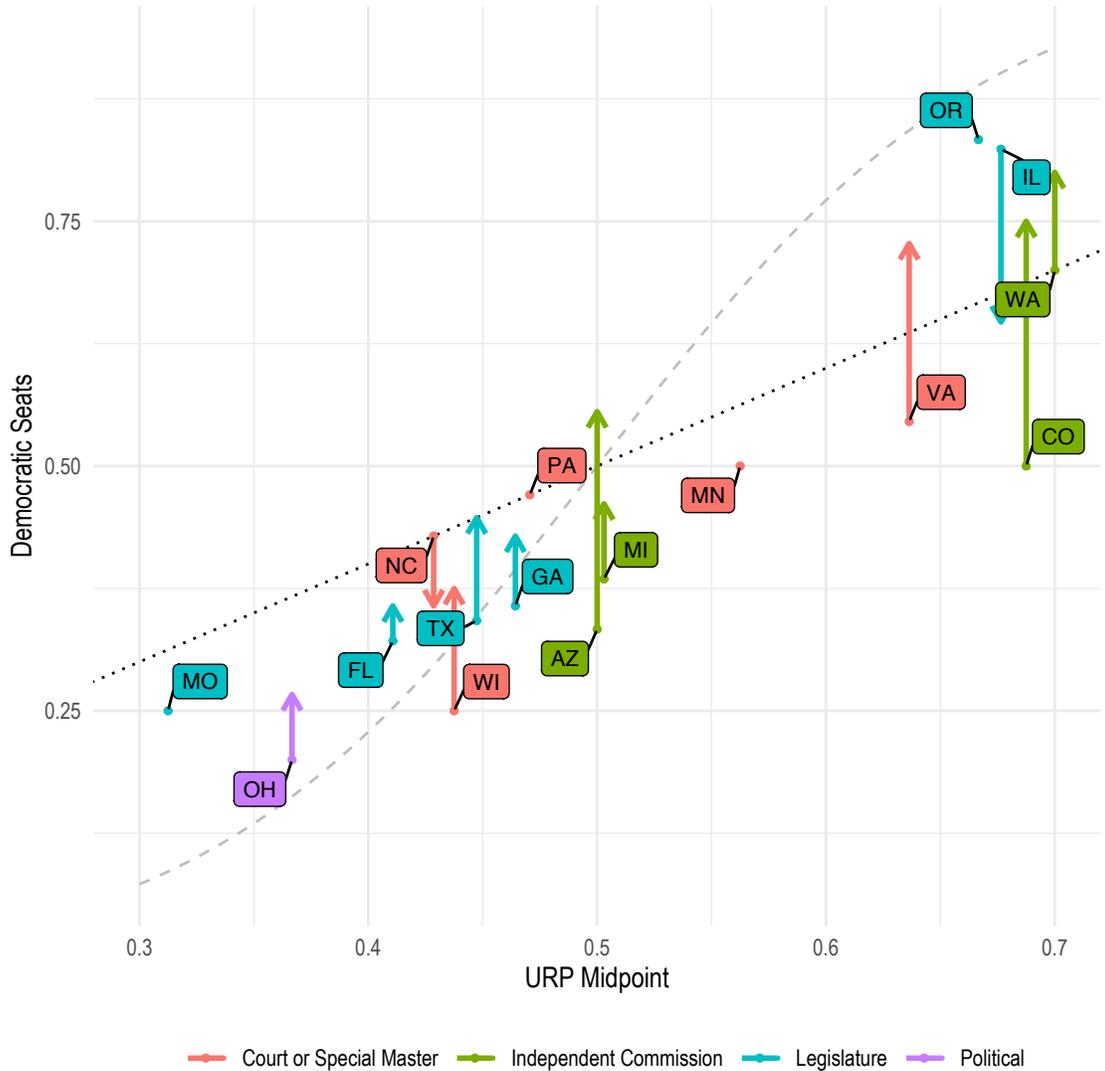


Figure A.4: **Comparing DCP to Actual Outcomes:** This figure displays the comparison between DCP and the actual outcome based on the state’s post-2020 redistricting process. The x-axis reports the midpoint between URP outcomes. The y-axis reports the Appendix seat share for a redistricting procedure. The point where each line originates denotes the projected outcome of the post-2020 redistricting. The point of each arrow denotes the DCP projected outcome. For reference, we also include a 45 degree line (e.g., proportional representation) and a line illustrating seat shares for an unbiased but majoritarian electoral system that follows the cube law. The sample includes states with four or more congressional districts and a URP midpoint between 0.3 and 0.7. First-mover for DCP is determined based upon the party controlling redistricting or (where not applicable) the majority party in each state.

First, as noted in the nationwide results, DCP reduces the sizable gap in seats compared to when one party unilaterally draws district maps. For example, in Virginia, we find that Democrats could draw maps where they win nine of eleven Congressional districts. In contrast, Republicans could draw maps where Democrats win only three Congressional districts—a gap of 6 seats or 55% of seats. Under DCP, the gap shrinks dramatically to 1 seat depending on the party that moves first. Overall, eight of forty-three states exhibit a two seat gap, twenty-one states exhibit a one seat gap, and fourteen states exhibit a zero seat gap under DCP. In every state where URP produces more than a one seat difference in seats, DCP narrows the range of outcomes as compared to URP, regardless of which party defines subdistricts and which party combines them. All states other than California, Florida, Ohio and Texas exhibit a first-mover or definer’s advantage (all states with large numbers of districts; we later explore the relationship between number of districts and first-mover advantage in more detail).

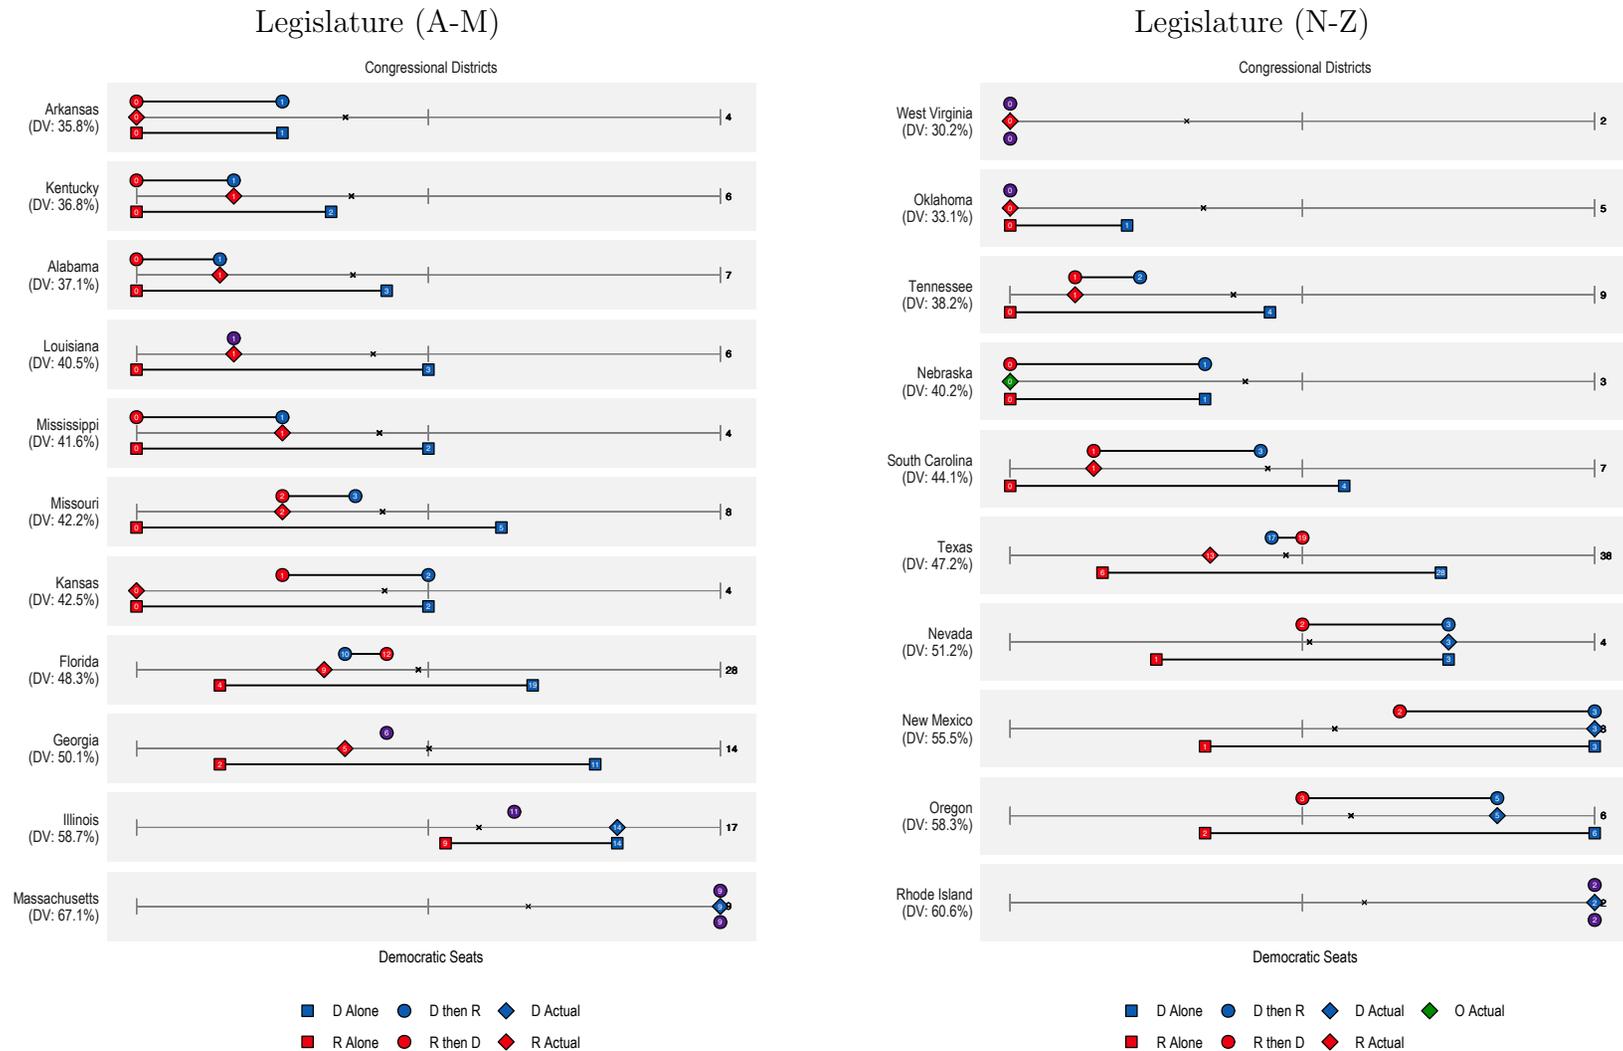


Figure A.5: State by State Simulation Benchmarks: This figure displays DCP (labelled “D/R then R/D”) results to several other benchmarks for state congressional districts where legislatures performed redistricting in the post-2020 redistricting cycle. Points labelled “D/R Alone” denote unilateral redistricting. Points labelled “D/R/O Actual” denote the actual outcome for the state in terms of projected seats after redistricting occurred. The small “X” mark denotes 2020 Presidential Election Democratic vote share, also reported in parentheses under each state name.

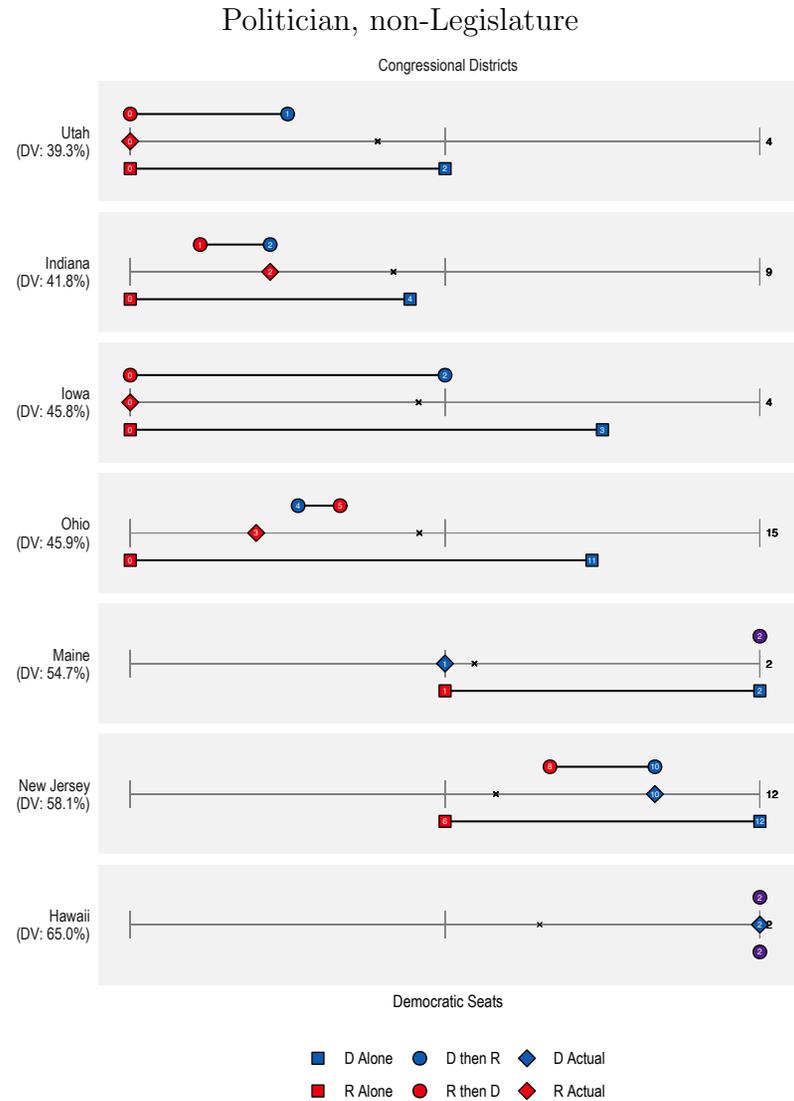
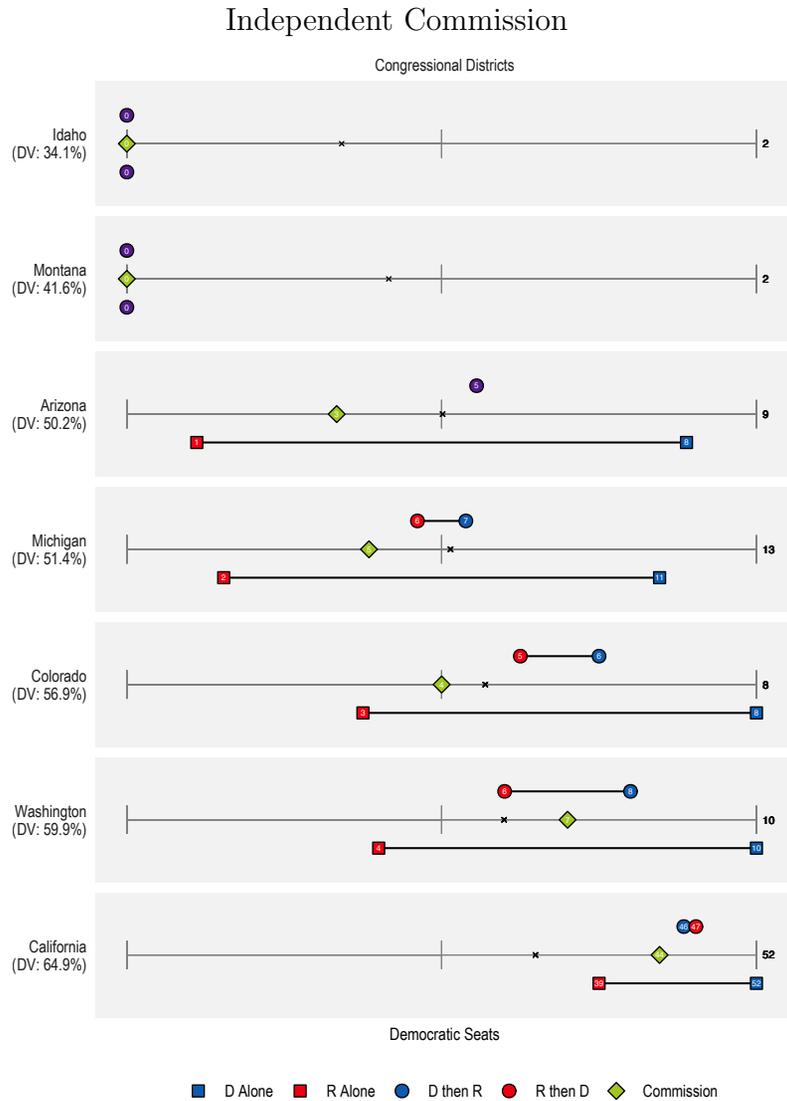


Figure A.6: **State by State Simulation Benchmarks:** This figure compares DCP (labeled “D/R then R/D”) results to several other benchmarks for state congressional districts where commissions or politicians guided redistricting in the post-2020 redistricting cycle. Points labelled “D/R Alone” denote simulated unilateral redistricting. Points labelled “D/R Actual” or “Commission” denote the actual outcome for the state in terms of projected seats after redistricting occurred. The small “X” mark denotes 2020 Presidential Election Democratic vote share, also reported in parentheses under each state name.

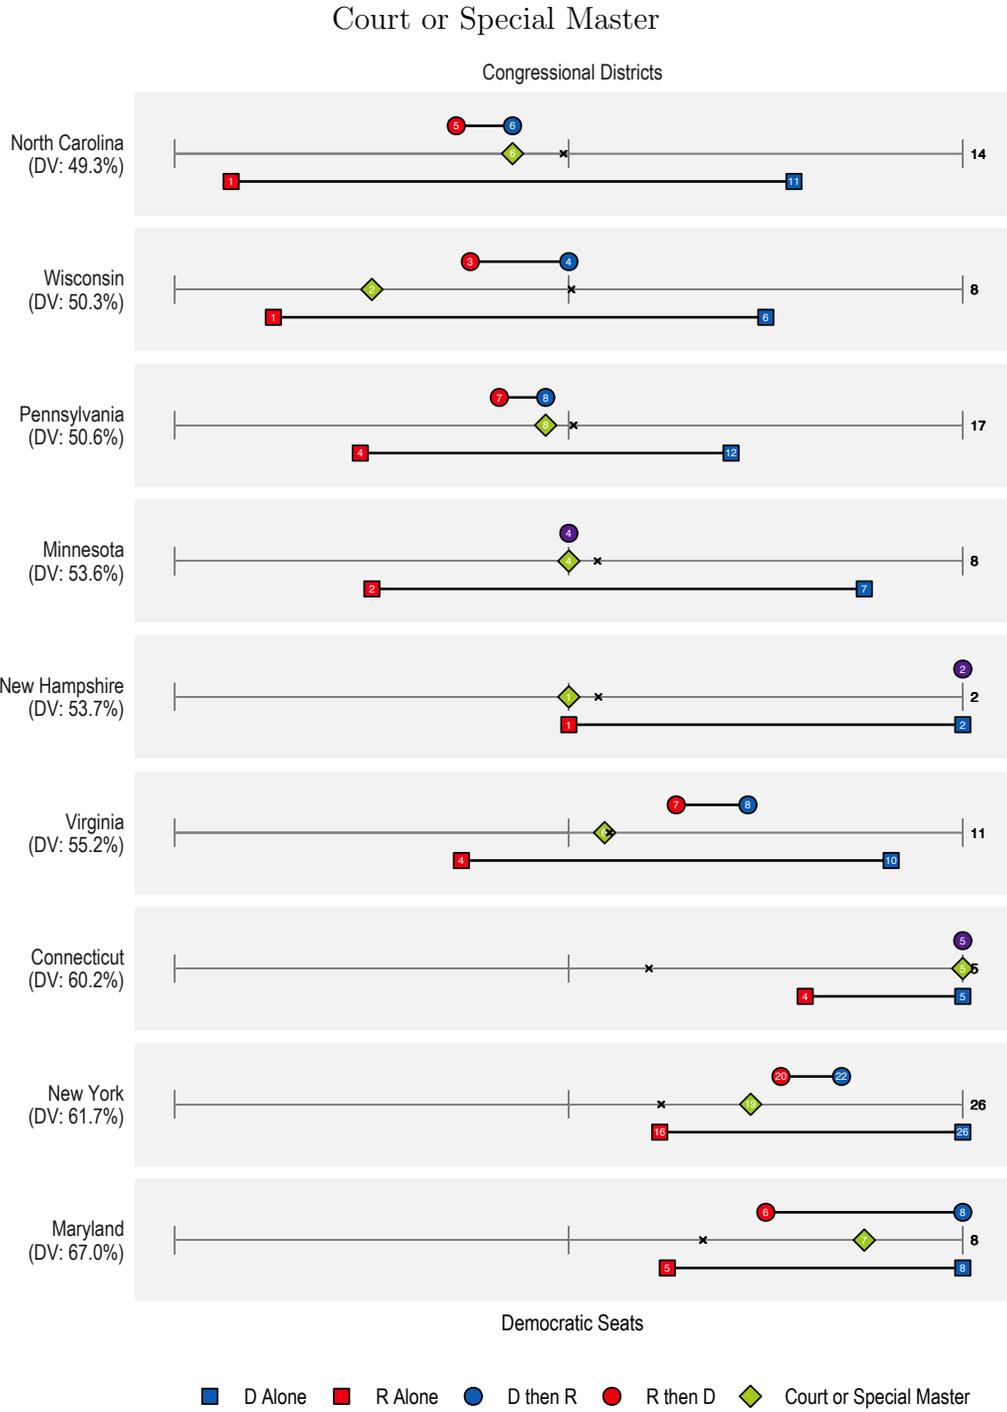


Figure A.7: **State by State Simulation Benchmarks:** This figure displays DCP (labeled “D/R then R/D”) results to several other benchmarks for state congressional districts where a court or special master guided redistricting in the post-2020 redistricting cycle. Points labelled “D/R Alone” denote simulated unilateral redistricting. Points labelled “Courts or Special Master” denote the actual outcome for the state in terms of projected seats after redistricting occurred. The small “X” mark denotes 2020 Presidential Election Democratic vote share, also reported in parentheses under each state name.

A.6.5 Simulation Results After Uniform Swing

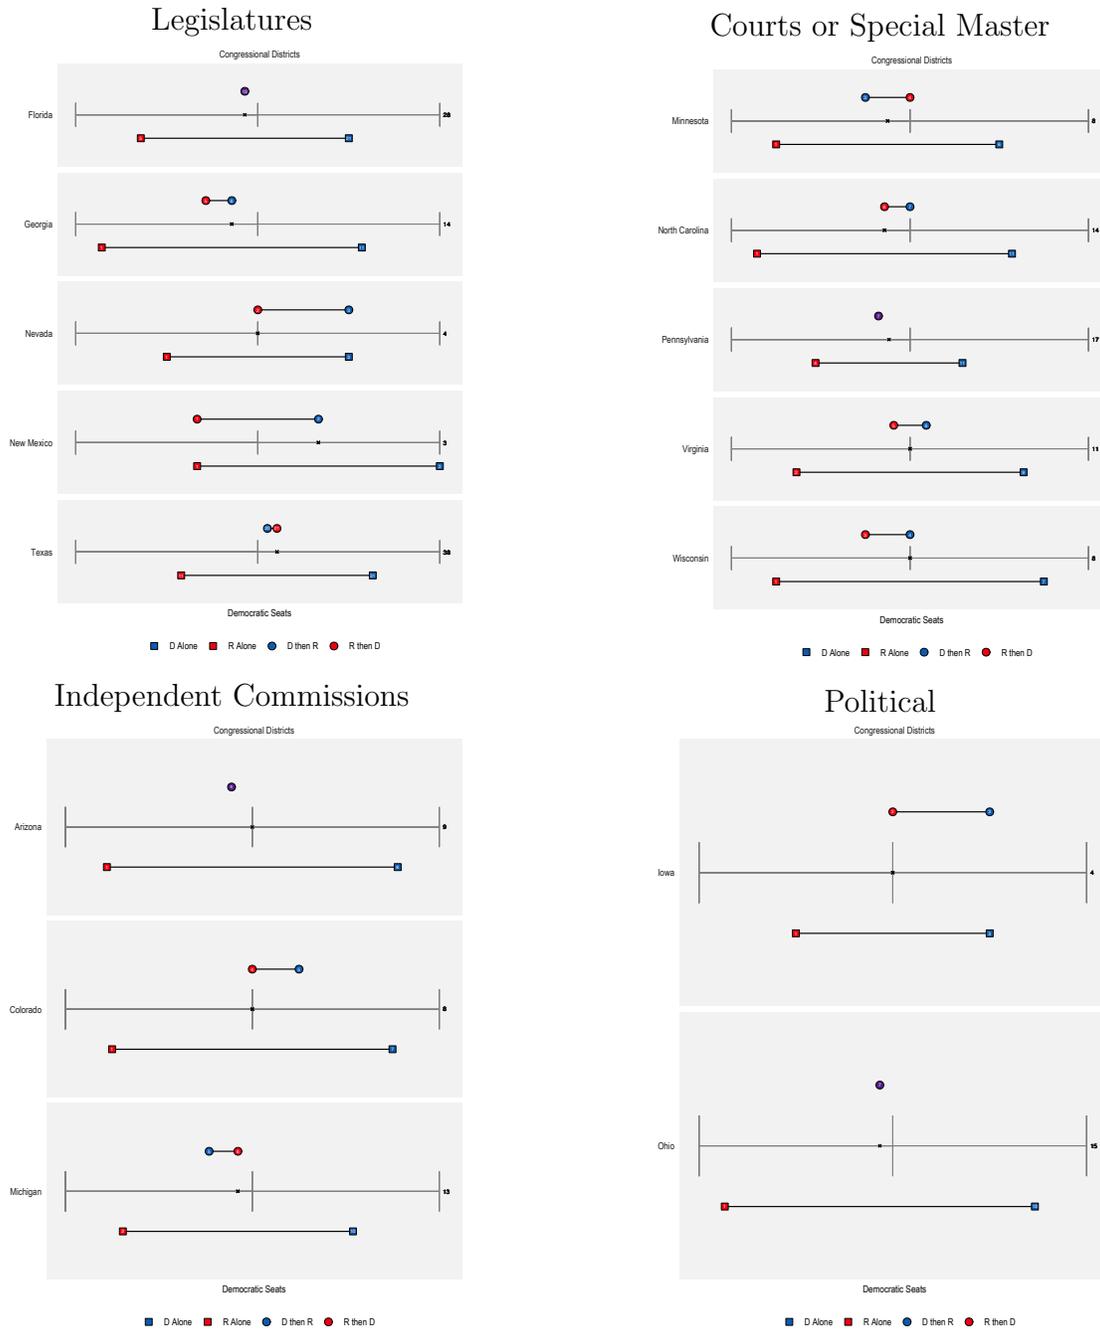


Figure A.8: This figure compares DCP (labeled “D/R then R/D”) results to unilateral redistricting for states where we have performed a uniform swing so that statewide 2020 Democratic Presidential vote share is evenly split between the parties. States included had an original Democratic vote share between 45% and 55%. Points labelled “D/R Alone” denote simulated unilateral redistricting. The small “X” mark denotes the geometric fairness target for the 50-50 maps (e.g., the average of the most extreme map in terms of Democratic seats won and Republican seats won).

State	Seats	R Alone	R Then D	D Then R	D Alone
FL	28	5	13	13	21
GA	14	1	5	6	11
NM	3	1	1	2	3
NV	4	1	2	3	3
TX	38	11	21	20	31

Table A.7: **Legislatures Redistricting:** Simulation Results After Uniform Swing

State	Seats	R Alone	R Then D	D Then R	D Alone
MN	8	1	4	3	6
NC	14	1	6	7	11
PA	17	4	7	7	11
VA	11	2	5	6	9
WI	8	1	3	4	7

Table A.8: **Court or Special Master Redistricting:** Simulation Results After Uniform Swing

State	Seats	R Alone	R Then D	D Then R	D Alone
AZ	9	1	4	4	8
CO	8	1	4	5	7
MI	13	2	6	5	10

Table A.9: **Independent Commission Redistricting:** Simulation Results After Uniform Swing

State	Seats	R Alone	R Then D	D Then R	D Alone
IA	4	1	2	3	3
OH	15	1	7	7	13

Table A.10: **Political Redistricting:** Simulation Results After Uniform Swing

A.6.6 Compliance with the Voting Rights Act

One potential concern for adopting the Define-Combine Procedure is the ability to use it to draw maps while complying with the federal Voting Rights Act, in particular by preserving majority-minority districts. If a set number of majority-minority districts are drawn, will DCP continue to produce similar results for each party if used for the remaining districts? To consider the possibility of VRA compliance, we ran a series of simulations in Georgia as follows:

1. Using 2020 census data and VTDs, use the shortburst algorithm to generate four majority-Black districts.
2. Remove the four majority-Black districts from map (effectively freezing these districts).
3. Run the unilateral and DCP simulation algorithms (as described above) to generate the boundaries of the remaining ten Congressional Districts, and record the results from each algorithm.

The above procedure was run ten times. In each simulation the boundaries of the four majority-Black districts were different (but were all concentrated in the Atlanta Metro area). Figure A.9 shows the results of the simulations. Across all ten simulations, there are substantial gaps in the number of Democratic seats with unilateral redistricting (4–5 seats), but much smaller gaps under DCP. In four of the simulations, both parties reached the same number of Democratic seats, in five simulations there was a one-seat difference, and in one simulation (likely due to the particular boundaries of the four majority-Black districts) there was a two-seat gap.

These results demonstrate that if the VRA districts can be drawn and then frozen in place, DCP can still produce an equitable result that reduces partisan gerrymandering in the remainder of the state.

Georgia Congressional Districts

Four Black Majority Districts Generated and Locked Before Unilateral or DCP Redistricting

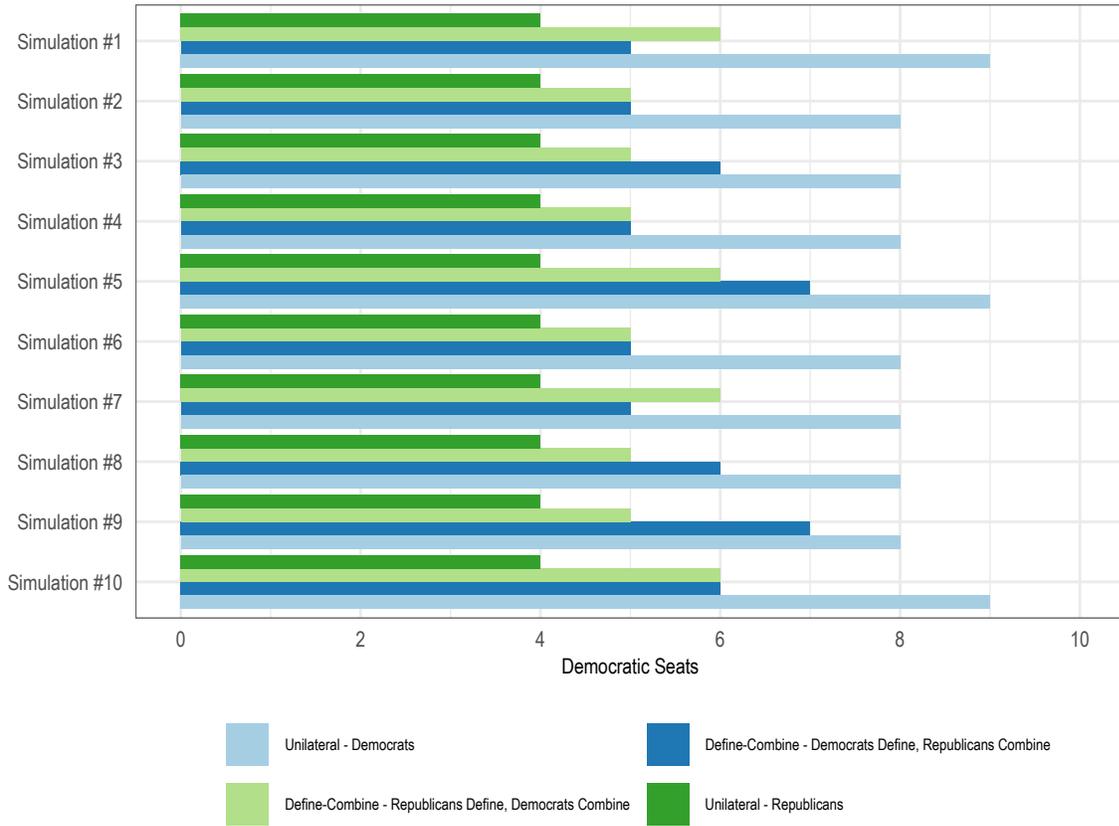


Figure A.9: Simulation Results for Georgia Congressional Districts, with Four Majority-Black Districts for VRA Compliance

A.7 Grid Maps

We now extend the analysis by simulating thousands of different distributions of voters on a grid map and then identifying the maps selected by each party under unilateral redistricting and under DCP.

A.7.1 Grid Map Simulations

The simulations proceed in four steps:

1. Define a grid of P precincts; each will have the same population.
2. Generate a random distribution of voters in each precinct. Instead of making each precinct either one Party A voter or one Party B voter, each precinct contains the same population size, but with a randomly selected percentage of voters supporting each party. First, we pick a target vote share m for Party A in the grid as a whole. We vary this across simulations in 2.5% increments from 30% to 70%. For each target vote share, we draw a vote share for each precinct from a truncated normal distribution with mean m .³⁸ We repeat this process 100 times for each level of m , resulting in 1,700 different distributions of voters.
3. Generate potential maps for the grid:
 - (a) Generate a set of possible maps of N districts, and a set of possible maps of $2N$ districts. For the simple 30-unit grid, we generated every possible map. For more complex grids, we generated a random sample of maps.
 - (b) For the set of $2N$ districts, generate all possible plans that combine pairs of contiguous districts. For the simple grid, we generated every possible combination, and for more complex grids we generated a random sample of combinations.
4. For each distribution of voters, examine the set of generated maps to identify:
 - (a) The best map for Party A, if Party A chooses a map unilaterally.
 - (b) The best map for Party B, if Party B chooses a map unilaterally.
 - (c) The map Party A would choose if it goes first under the Define-Combine Procedure.
 - (d) The map Party B would choose if it goes first under the Define-Combine Procedure.

³⁸The truncated normal distribution is bounded at 0 and 1 and has a standard deviation of 0.25.

For each identified map we calculate the number of seats won by each party.

A.7.2 Varying Partisan Composition of Voters

Figure A.10 presents the results for a 30-voter grid (5 districts of 6 voters each). The x-axis corresponds to the share of voters supporting Party A, and the y-axis to the share of seats won by Party A. The dotted lines illustrate the average across simulations when Party A (in blue) and Party B (in red) draw the district maps unilaterally. The solid lines display the averages for DCP, with the blue (red) line corresponding to Party A (B) as the definer followed by Party B (A) as the combiner. When one party dominates state-wide vote share (e.g., more than 75% of vote), the results remain similar no matter who controls the unilateral process and no matter if DCP is implemented. However, when both parties are competitive in terms of vote share, significant differences emerge. At $V = 0.5$, the advantage conferred by drawing maps unilaterally is $\delta_{.5}^U \equiv 0.8 - 0.2 = 0.6$, or three seats. Under DCP the first-mover advantage is $\delta_{.5}^D \equiv 0.6 - 0.4 = 0.2$, or one seat. All told, implementing DCP on a map evenly divided between the parties will reduce the advantage of the redistricters, as compared to their opponents, by a seat share of 0.4 or 2 seats. DCP does offer the defining party a first-mover advantage in this context of one additional seat over the combiner.

The definer's advantage in DCP declines as the size of the grid and/or the number of districts increases. Figure A.11 displays the average number of seats won for a larger hexagonal grid where 150 precincts are divided into 15 districts. While there is still a substantial gap in seat share when parties draw maps unilaterally, the seat shares under DCP converge, no matter who moves first. For example, when voters in the 150 precinct grid are split evenly between the parties and there is unilateral redistricting, the party drawing the map gets a seat premium roughly equal to half of all seats on average ($\delta_{.5}^U = 0.54$). In contrast, under DCP the seat share remains nearly the same no matter which party goes first ($\delta_{.5}^D = .007$).

Figure A.12 plots the values of δ_V^U and δ_V^D against vote share for the 150 precinct grid. Partisan advantage due to unilateral redistricting (δ_V^U) peaks when the state is evenly split. In contrast, the partisan advantage under DCP (δ_V^D) remains relatively low across the full distribution. While it increases slightly for vote shares in the range 0.37-0.45 and 0.55-0.63, the defining party still receives an advantage of less than one of sixteen seats, compared to an advantage equal to four of sixteen seats at those same vote shares under unilateral redistricting. Across all vote shares, DCP reduces the partisan advantage of going first and substantially limits the ability of each party to gerrymander.

Because we randomly generate these grid maps, the average geographic bias due to clustering of partisans is zero for any vote share. As a result, we may identify bias due to

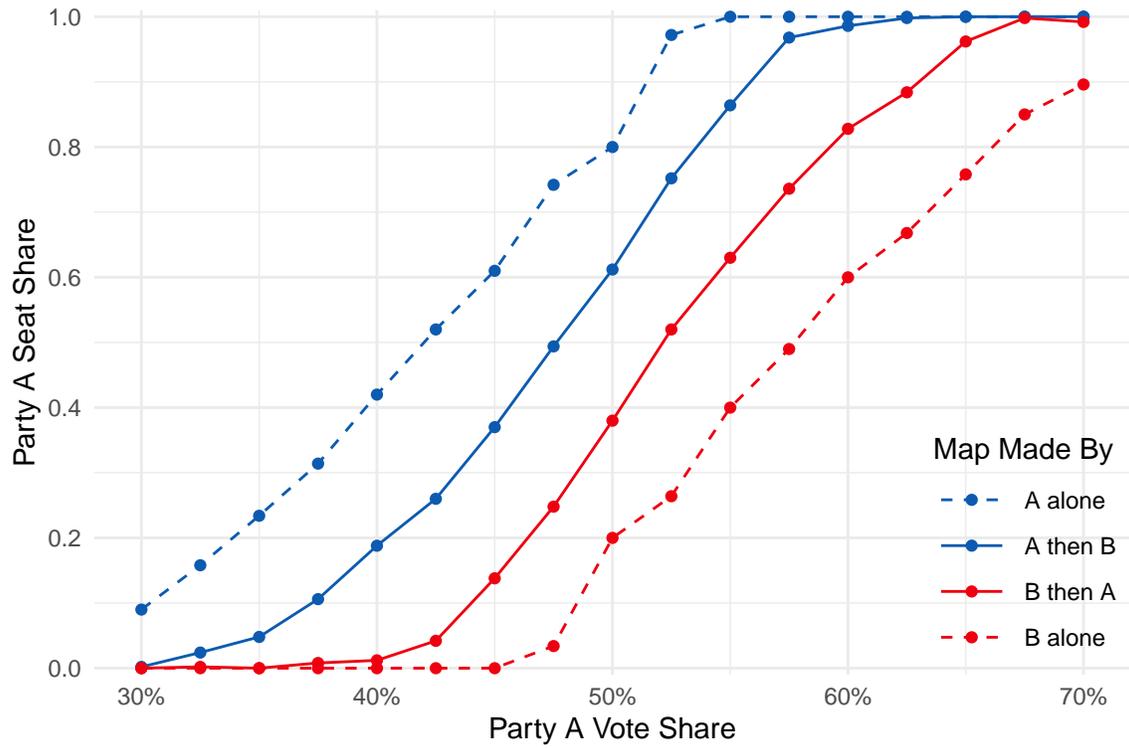


Figure A.10: Results for Voter Distribution Simulations on a Simple Grid. Each simulation uses a 5x6 grid.

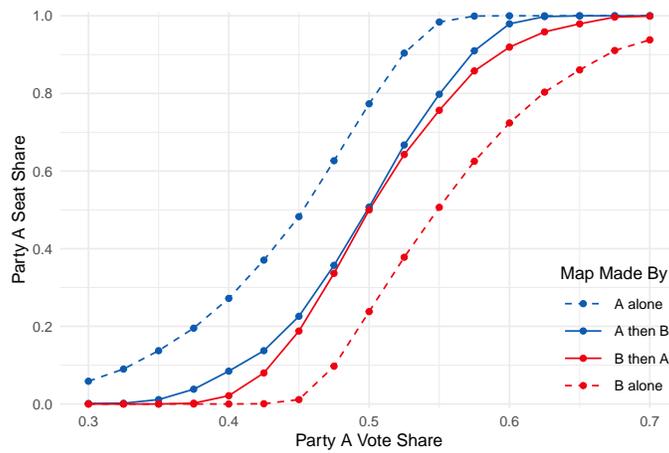


Figure A.11: Results for Voter Distribution Simulations on a Larger Grid (150 precincts; 15 districts)

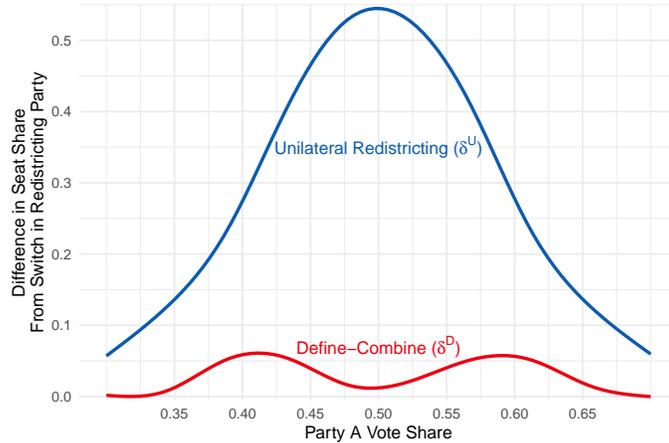


Figure A.12: Differences in Seat Shares from Switch in Redistricting Party on the 150 Precinct Grid

redistricting by a direct examination of the seat shares for unilateral redistricting and for DCP. For both maps simulated in this section, DCP reduces bias most when votes are evenly split; as a party’s vote share increases, the gap in biases due to redistricting procedure narrow, until converging for vote shares 0.6 and above.

The figures above present average numbers of wins for each party or the average values of δ^U and δ^D , across 1,700 different random distributions of voters across the grids (100 distributions for each mean level of vote share for Party A, from 30% to 70% in increments of 2.5%). However, we also would like to know how DCP performs not just on average but for every possible map. That is, are there any distributions of voters for which DCP does not represent a meaningful improvement over the unilateral case? To address this question, we examine the results from each separate distribution of voters. For each voter distribution on the 150 precinct grid, we calculated δ^U , δ^D , and the difference between them. If $\delta^D = \delta^U$, then, for that particular voter distribution, DCP fails to improve the outcome. Figure A.13 presents scatter plots showing, for each mean level of Party A vote share, the values of δ^U (on the x-axis) and δ^D (on the y-axis). Points are sized by the number of times that result is realized in each of the 100 simulations for each vote share level. In all cases, $\delta^D < \delta^U$. In other words, DCP improved the outcome in all 1,100 cases by a meaningful margin.

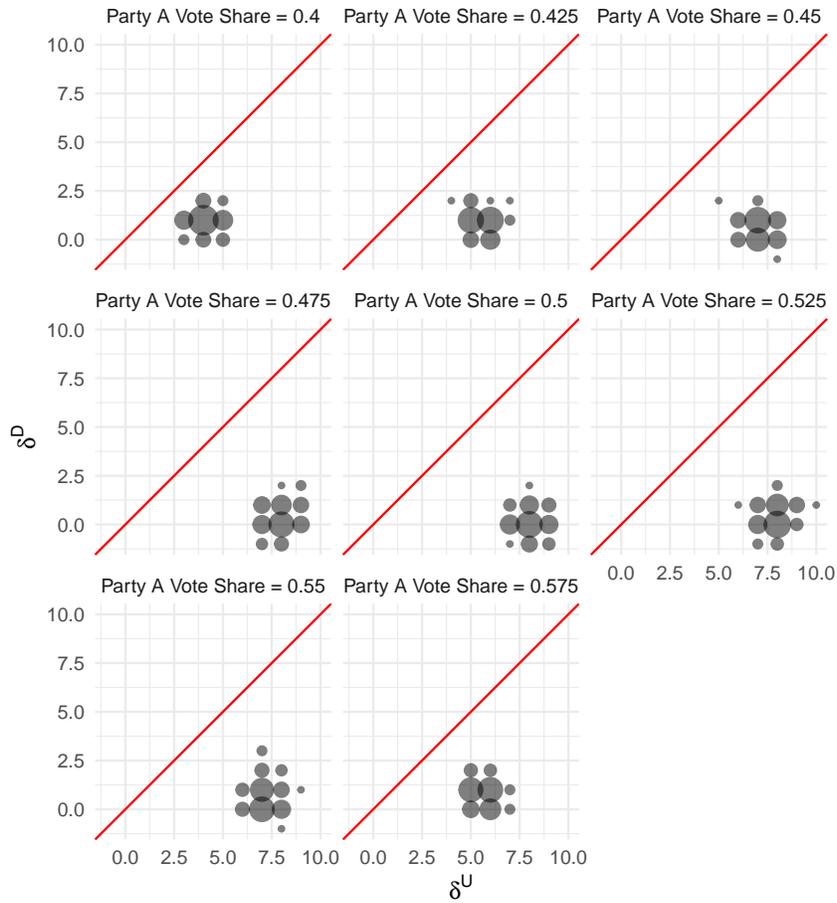


Figure A.13: Values of δ^U and δ^D for each simulated voter distribution on the 150 precinct grid.

A.7.3 Alternative Objectives to Maximizing Seats Won in Next Election

In the simulations above, both parties sought to maximize seats won. However, parties may have objectives other than maximizing current seat share. These differing objectives could include maximizing the odds of maintaining a majority, or preventing losses in future elections rather than just the impending one.

Katz, King and Rosenblatt (2020) notes that if a majority party controls the redistricting process but is worried about the future (i.e., “running scared”), then the party may favor plans that maximize partisan bias and minimize responsiveness to insulate the party from future partisan swings. In contrast, if a party expects to win a majority of the vote in future elections, they may seek to create districts that are microcosms of the state as a whole by maximizing responsiveness and minimizing bias (Katz, King and Rosenblatt, 2020). Further complicating matters, the objectives of competing parties could also differ, since a party seeking to maintain a partisan advantage will likely have goals that differ from the party seeking to gain one.

To explore the implications of asymmetric utility functions, we simulated redistricting games where the competing parties had differing objectives: Party A is confident about the future and maximizes responsiveness while Party B is running scared and seeks to maximize bias.³⁹Figure A.14 illustrates results when Party A earns a vote share over 0.5. The key point revealed by this exercise is that the main findings we obtained when both parties maximized current seats won still hold up. DCP reduces the level of bias as well as the definer’s advantage for state-wide vote shares between roughly 0.5 and 0.65. For vote shares above 0.65, the underlying advantage of Party A is so great that the redistricting procedure no longer matters and Party A simply wins all the seats.

A.7.4 Geographic Partisan Clustering

Geographic clustering of partisans affects the ability of parties to gerrymander. The clustering of Democratic voters in cities can lead to “unintentional gerrymandering”—even district maps drawn with no intention to gerrymander, such as randomly generated maps, still exhibit partisan biases disadvantaging Democrats due to partisan differences in geographic concentration of voters (Chen and Rodden, 2013). High levels of clustering disadvantage the clustered party since its votes are more likely to be inefficiently grouped together, leading to more “wasted” votes (Stephanopoulos and McGhee, 2015). Similar arguments may apply to

³⁹Party A maximizes utility by choosing the option that provides the maximum level of responsiveness. When multiple plans produce the same value of responsiveness, Party A uses wins as tiebreaker. Similarly, Party B maximizes bias and uses wins as a tiebreaker. Katz, King and Rosenblatt (2020) finds a tradeoff between bias and wins, such that maximizing one should generally reduce the other as well.

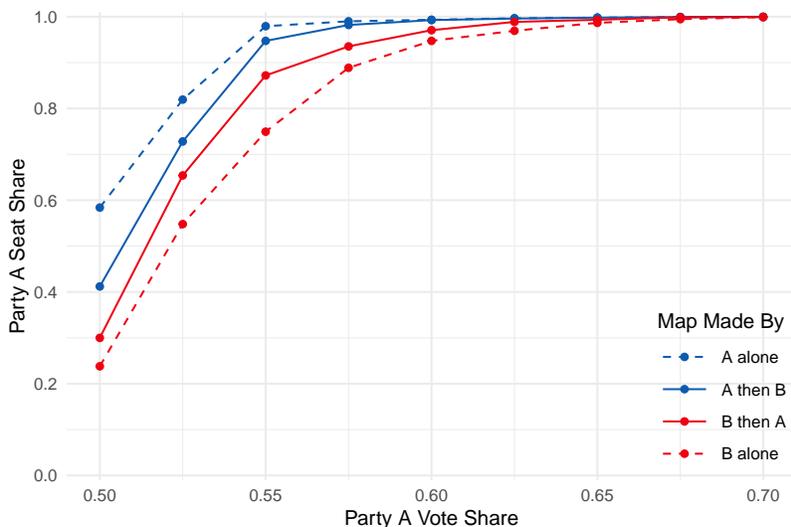


Figure A.14: Asymmetric Utility Functions: Party A maximizes responsiveness and Party B maximizes bias.

racial gerrymandering, given differential levels of geographic clustering of racial groups (Magleby and Mosesson, 2018). DCP’s effectiveness at reducing partisan gerrymandering, along with the size of the advantage conferred to the first versus second mover, might similarly depend in part on the level of geographic clustering. We explore that possibility here.

We examine two different elements of clustering: (1) overall clustering of partisans, and (2) differential clustering of one political party.

To examine overall clustering, we employ Moran’s I, a common measure of spatial autocorrelation (Moran, 1948) used in academic political science research and legal work on redistricting (Mayer, 2016).⁴⁰ Moran’s I ranges from -1 to 1 , with more positive values denoting increased clustering. For example, a 5-by-6 grid with all Party A voters located on the left and all Party B voters on the right yields a Moran’s I of 0.796 ; a grid with each party’s voters distributed evenly yields a Moran’s I of -1 (see Figure A.18). A Moran’s I of 0 indicates randomly dispersed voters with neither clustering nor a pattern of “even” dispersion. Appendix A.7.6 illustrates maps with varying levels of Moran’s I values and provides additional details on Moran’s I calculations.

We first randomly assign single voters to 5-by-6 grids of precincts, setting the probability that a voter supports Party A equal to 50% . For each randomly-drawn grid, we evaluate map-drawing under both unilateral redistricting and DCP, and we also calculate Moran’s I for each grid map. Figure A.15 plots the relationship between seats won by Party A and the

⁴⁰See Cho and Gimpel (2012) for a discussion of the use of various clustering measures in research on American Politics. Scholars recognize Moran’s I as a valid measure geographic clustering of partisans (Stephanopoulos, 2018).

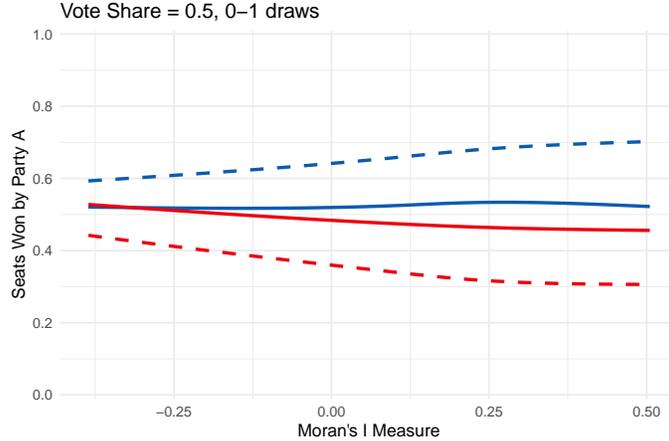


Figure A.15: Results by Clustering, Vote Shares = 50%

amount of geographic partisan clustering when Party A support is fixed at 50%.

As clustering increases, the definer’s advantage ($\delta_{.5}^U$) for unilateral redistricting increases slightly as well. Going from the lowest observed level of Moran’s I to the highest increases $\delta_{.5}^U$ by about 0.2. How does DCP perform under different levels of geographic partisan clustering? When clustering is low, DCP removes the definer/first-mover advantage (e.g., $\delta_{.5}^D = 0$). When clustering is high, even under DCP there remains a small but significant definer/first-mover advantage. Nonetheless, DCP’s performance varies only slightly due to clustering in this example, and at all levels of clustering DCP dramatically reduces the advantage of the definer compared maps drawn under unilateral redistricting. Even at the highest levels of observed clustering, DCP eliminates at least 80% of the definer’s advantage observed in unilateral redistricting (e.g., $1 - \frac{\delta_{.5}^D}{\delta_{.5}^U} \geq 0.8$). Overall, then, map-wide clustering does not appear to meaningfully alter our general results for the definer’s advantage.^{41,42} In terms of bias induced by the redistricting process, note that—for any level of Moran’s I—there is no underlying geographic bias on average since voter locations are assigned randomly. Under these conditions, DCP reduces bias as compared to URP independently of the level of clustering.

⁴¹Figures A.20 illustrate the relationship between seats won and clustering for different values of state-wide vote share V .

⁴²In addition, note that the range of clustering in our grid examples does not span the full range of possible values of Moran’s I. We do not observe many maps that have negative Moran’s I values simply because they are highly unlikely to occur by our random sampling procedure. (See Appendix A.7.6 for more details.) This is not a limitation of our analysis of the effects of clustering on gerrymandering or DCP for two reasons: first, in a case where partisans are very evenly dispersed (Moran’s I values close to -1), it becomes impossible to gerrymander because there is no way to “crack” or “pack” partisans together, so bias in these maps (under either unilateral redistricting or DCP) will be close to 0. Second, when using real data on partisanship to calculate measures of the dispersion of partisans, all states show significant levels of partisan geographic clustering.

Another way to examine the effect of geographic bias is to compare the partisan advantage for one party from the full sample of simulated maps to the results from choosing a map using URP or DCP. In our grid simulations, we simulate 100 random draws of voters for each vote share level. We can use the 100 independent draws when vote shares are perfectly split between the parties to calculate the average wins across all maps for each vote draw, and we can compare it to the selected maps for each vote draw. Figure A.16 plots the results. Across the 100 draws, geographic bias (the x-axis) varies by about 5% in favor of each party. Geographic bias favoring Party A correlates positively with Party A seat share under either redistricting procedure. As geographic bias decreases, we see less variation in outcomes, and expect both parties to evenly divide the seats.

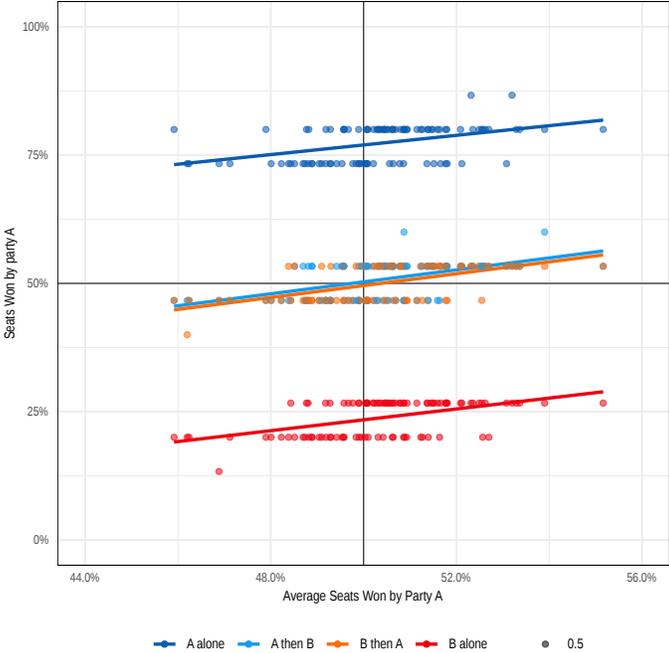


Figure A.16: Relationship Between Geographic Bias and Outcomes for Simulated Grids

A.7.5 Number of Districts and Definer’s Advantage

To explore the relationship between number of districts and the advantage conferred by controlling the redistricting process (e.g., as the one drawing the map under unilateral redistricting or as the Definer under DCP), we simulated potential maps for a hexagonal grid. For all maps, we had a fixed population (4000 hexagons or precincts), fixed vote shares (evenly split between the parties), no systematic voter clustering, and varying numbers of districts to be drawn. The number of districts ranged between 4 and 200. As a result, maps with fewer districts exhibit a higher population to district ratio while maps with more districts had a lower population to district ratio.

As Figure A.17 illustrates, unilateral redistricting exhibits a hump shape in terms of the advantage conferred to the party controlling the redistricting process. This advantage peaks for a state with eight districts and then declines. In contrast, under DCP the definer’s advantage is highest for small numbers of districts, turns slightly negative (i.e., indicating a combiner’s advantage / second-mover advantage) and then converges towards zero.

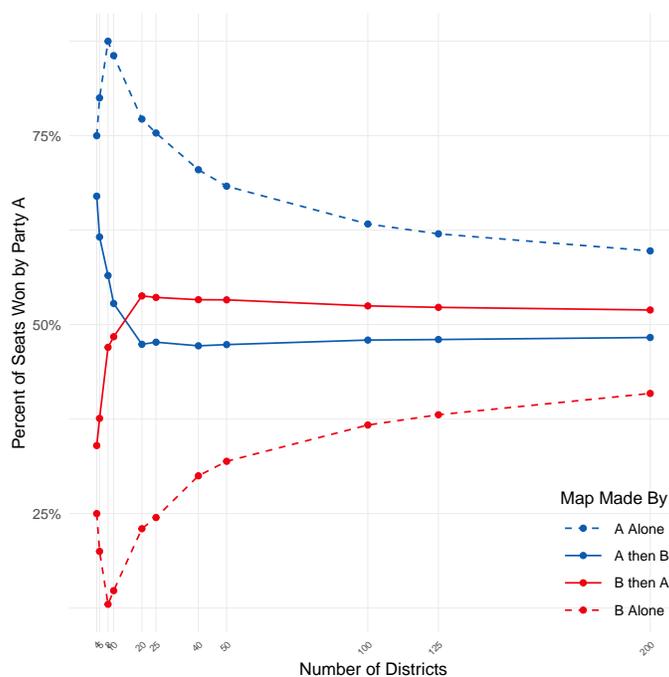


Figure A.17: Relationship Between Number of Districts and First- and Second-Mover Advantage. We generated a 63x64 grid of 4,000 hexagons, each with equal population, and used GerryChain to generate potential maps with different numbers of districts. We generated sets of 10,000 maps each, with 8, 10, 16, 20, 40, 50, 80, 100, 200, 250, and 400 subdistricts, resulting in 4, 5, 8, 10, 20, 25, 40, 50, 100, 125, and 200 final districts. We generated 25 different voter distributions, each with mean of 50% for each party, and calculated the maps selected by each party under the four methods...

These results raise several questions. First, why does the advantage conferred to the unilateral redistricter appear to be maximized at a relatively low (but not at the lowest) number of districts? To see why, first recall that voters split the vote evenly between both parties. As a result, *at least one* district will always need to register a majority for the redistricting party's opponent.⁴³ This leads to optimal gerrymandering where a party packs the opponent into as few districts as possible. For example, in a map with 4 districts and an evenly split vote, the redistricter packs half of the opposing party voters into one district and distributes the remaining half equally into the three remaining districts. This approach also works at 8 districts, where the unilateral redistricter can win 7 of 8 seats.

If a state had an evenly split vote and no contiguity constraint, then a redistricter could pack opposing party voters into one district and win all remaining districts no matter the number of districts. However, because districts must be contiguous, geography begins to limit how efficiently voters may be packed and cracked. As the number of districts increases, this problem increasingly limits the level of gerrymandering. In our example, once the unilateral redistricter must draw 200 districts, he or she is able to edge out only a 0.2 definer's advantage. As the number of districts converges to the number of individuals in the state, the advantage from controlling redistricting will diminish to zero.

Second, why does DCP begin to exhibit a second-mover/combiner's advantage once the number of districts surpasses some threshold? Note first that no matter the number of districts the combine step prevents the definer from fully taking advantage of the ability to pack and crack. Attempts to do so can be mitigated through the combine step. However, once the definer must draw over some threshold number of districts, then the second-mover/combiner has so many possible options (on average) that he or she can wrest some of the ability to pack and crack back from the definer. To take this to an extreme, consider a case where the definer is drawing as many sub-districts as there are people in a state. In this case, the second-mover/combiner faces a problem equivalent to unilateral redistricting (i.e., taking the full population N and creating $\frac{N}{2}$ districts); the combiner therefore retains advantages similar to those of a unilateral redistricter outlined earlier in this section.

⁴³The proof: Suppose each party comprised half of each district (and breaking ties randomly would lead to an expected number of wins for each party equal to half the number of districts). Then any change in the balance of voters between districts will create a majority for each party in at least one other district. Hence no other configuration of voters can be reached without each party winning at least one district.

A.7.6 Partisan Clustering Calculations

We use Moran’s I as a measure of the degree of geographic clustering among partisans, both in our simulated grid example and with the precinct-level election data for each of the states in our analysis. We use the following formula to calculate Moran’s I:

$$I = \frac{N}{W} \frac{\sum_i \sum_j w_{ij} (v_i - V)(v_j - V)}{\sum_i (v_i - V)^2}$$

Where N is the number of spatial units, v_i and v_j are the vote shares of grid square i and j respectively, V is the average of the vote share across the entire simple grid, w_{ij} are spatial weights, and W is equal to the sum of the weights $\sum_{ij} w_{ij}$. In the simple grid analysis presented in Appendix A.7.4, the vote shares v_i and v_j will either be 0 or 1 and we use “neighbor” weights such that $w_{ij} = 1$ if grid squares i and j are adjacent rook neighbors and 0 otherwise.

Figure A.18 shows examples of different configurations of voters along with the corresponding Moran’s I measure for each, holding the overall map vote share for each party at a constant 0.5. Figure A.18 (a) displays a high clustering scenario, where all voters from each party are packed on either side of the grid, with a Moran’s I of 0.796. Figure A.18 (b) shows the case where voters from each party are perfectly evenly spread out across the grid, resulting in a Moran’s I of -1 . Figure A.18 (c) shows a case where Moran’s I is approximately 0, indicating neither clustering nor a tendency towards even dispersion. And Figure A.18 (d) shows an example of a “city” of Party A voters (white dots) surrounded by Party B voters (black dots), and demonstrates that in this sort of geographic setup Moran’s I is 0.479, indicating significant geographic clustering.

Figure A.19 (a) and (b) demonstrate scenarios where the overall vote share for Party A (white dots) of the map is much lower than $V = 0.5$, but the geographic clustering of Party A voters remains high. Figure A.19 (c) and (d) show random draws from our simulation procedure, this time again with the vote shares of each party set to 0.5, but with only a slightly positive (c) or slightly negative (d) Moran’s I clustering measure.

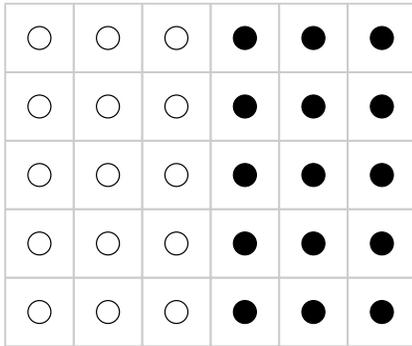
The range of Moran’s I in the clustering results presented in Appendix A.7.4 with our simple grid simulations are limited due to the fact that our random sampling procedure for the partisanship of each voter makes it highly unlikely that very negative values of Moran’s I will occur naturally. This is because, in general, there are many different possible ways to cluster voters — in one corner of the map, in another corner, in the middle, etc. — but only a small number of ways to have voters evenly dispersed. For example, in order to have a Moran’s I of -1 , voters of each partisan affiliation need to be evenly spread across the entire map (see Figure A.18 (b)). Given a probability of 50% that a voter will be for Party

A or Party B, this means that the probability of Moran's I being exactly -1 is equal to $2 \times (0.5^{30})$, which is approximately 1-in-500 million.⁴⁴ This is not a significant limitation of our clustering analysis because a map with clustering close to -1 become impossible to gerrymander and, when looking at real precinct-level map data, all states we use in our paper demonstrate significant geographic clustering of partisans.

Figure A.20 plots the relationship between seats won by Party A and Moran's I (as shown in Figure A.15), but with each plot depicting results for a different overall Party A vote share, ranging from 40% to 60%. In general, the pattern is similar to the case where statewide vote share is split 50-50; in all cases, $\delta_{.5}^D$ is substantially lower than $\delta_{.5}^U$, indicating a significant reduction in the advantage conferred to the party controlling the redistricting process and in bias due to redistricting. Thus, even if vote shares of each party vary in addition to level of geographic clustering of political parties, DCP still proves effective.

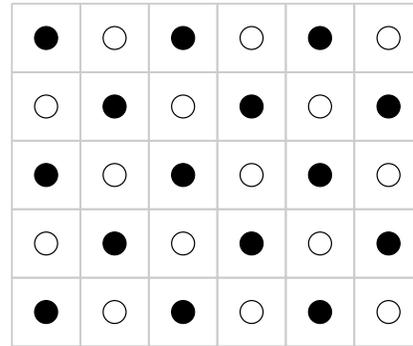
⁴⁴There are two possible ways to do have perfect dispersion - one starting with a black dot in the corner and alternating across the rest of the grid, and one starting with a white dot.

Moran's I = 0.796, Vote Share = 0.5



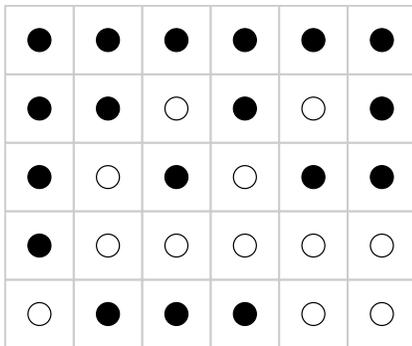
(a) High Clustering

Moran's I = -1.0, Vote Share = 0.5



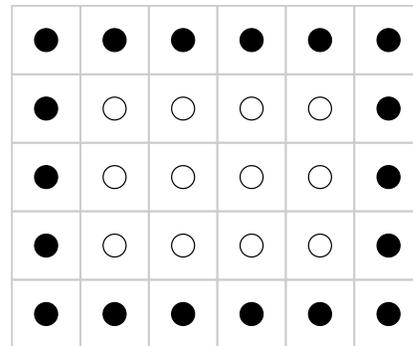
(b) No Clustering

Moran's I = 0.003, Vote Share = 0.5



(c) Randomly Spread

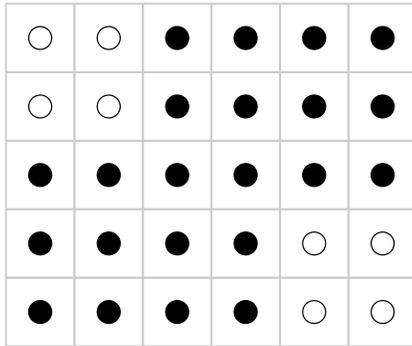
Moran's I = 0.479, Vote Share = 0.5



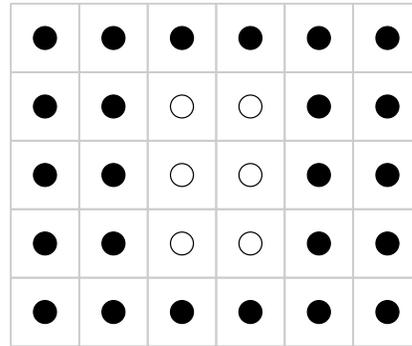
(d) "City" Example

Figure A.18: Ranges of Moran's I with Different Vote Distributions

Moran's I = 0.531, Vote Share = 0.266



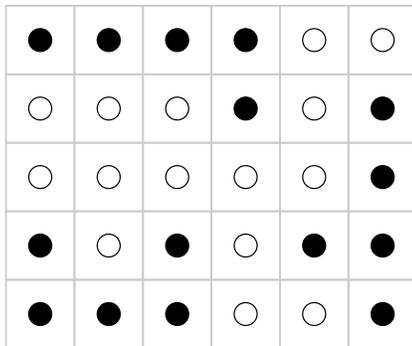
Moran's I = 0.531, Vote Share = 0.20



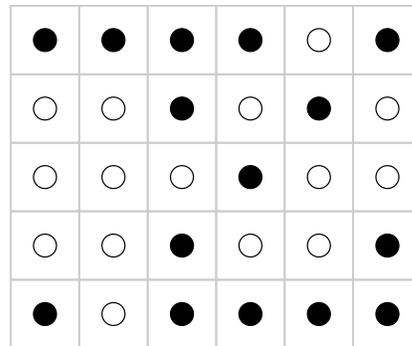
(a) High Moran's I, Low Vote Share

(b) High Moran's I, Lower Vote Share

Moran's I = 0.143, Vote Share = 0.5



Moran's I = -0.102, Vote Share = 0.5



(c) Random - Positive Clustering

(d) Random - Negative Clustering

Figure A.19: Samples of Moran's I with Different Vote Distributions

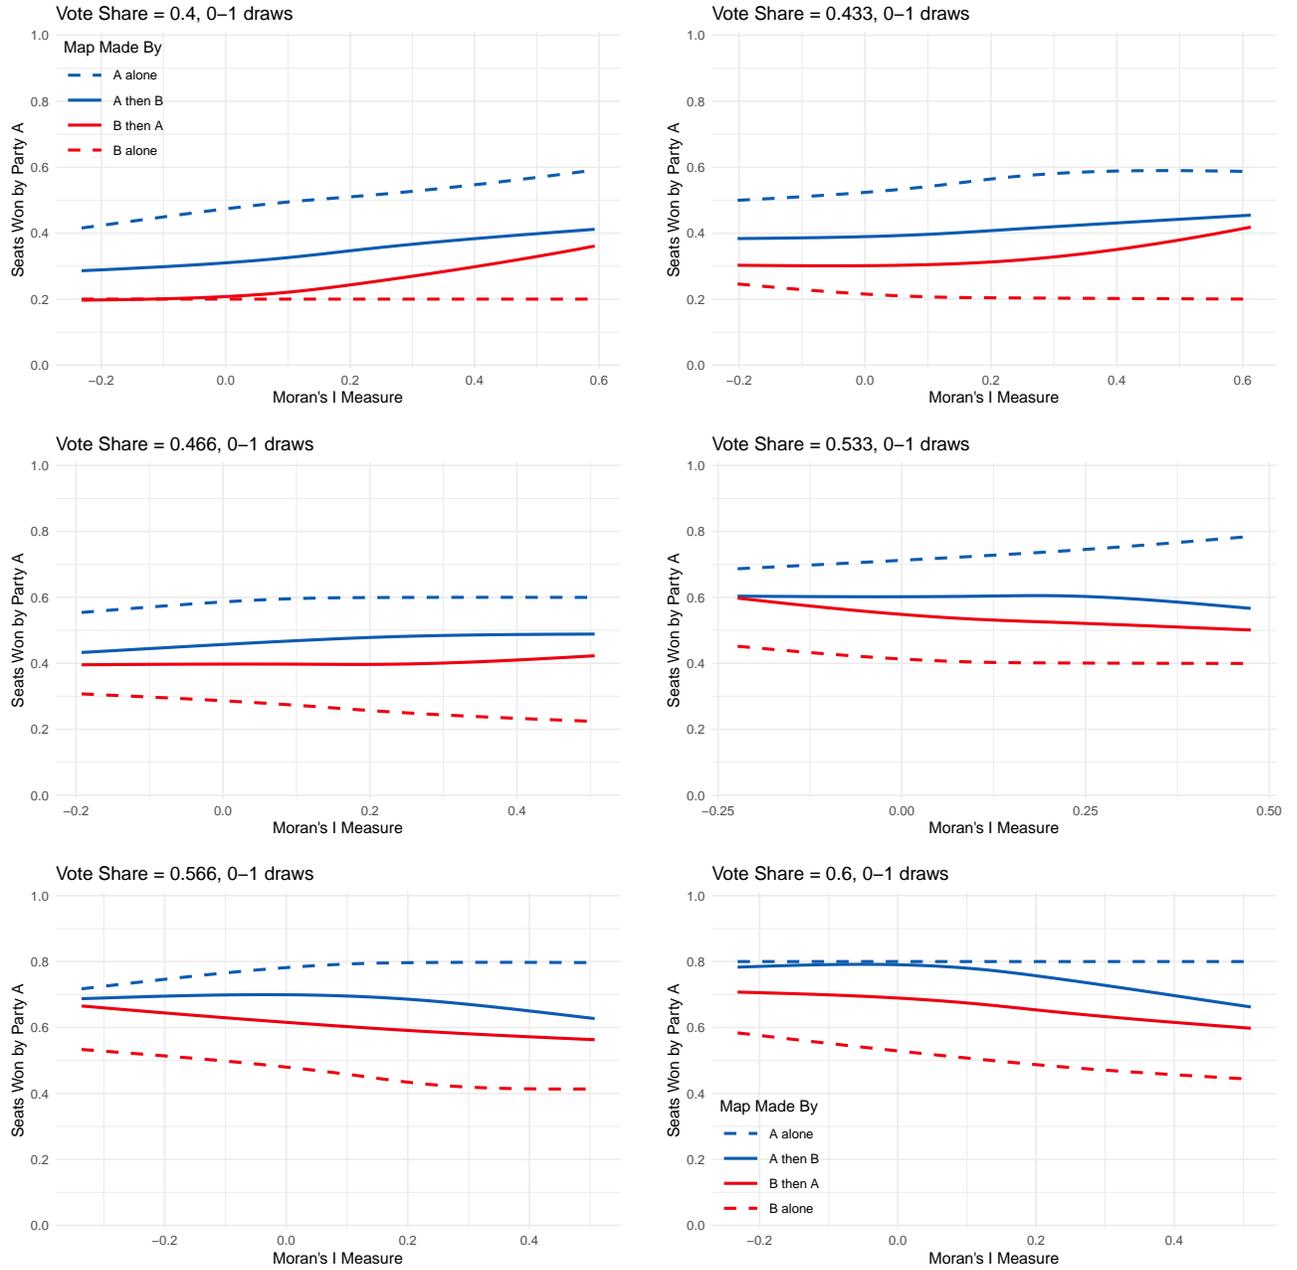


Figure A.20: Define-Combine Results, by Vote Share and Moran's I, 40%-60%

References

- Alexeev, Boris and Dustin Mixon. 2019. “Partisan Gerrymandering with Geographically Compact Districts.” *Journal of Applied Probability* 55(4):1046–1059.
- Altman, Micah and Michael P. McDonald. 2011. “BARD: Better Automated Redistricting.” *Journal of Statistical Software* 42(4).
- Benadè, Gerdus, Ariel Procaccia and Jamie Tucker-Foltz. 2021. “You can have your cake and redistrict it too.” *ACM Transactions on Economics and Computation* .
- Brams, Steven. 2020. “Making Partisan Gerrymandering Fair: One Old and Two New Methods.” *Social Science Quarterly* 101(1):68–72.
- Brams, Steven and Alan Taylor. 1996. *Fair Division: From cake-cutting to dispute resolution*. Cambridge University Press.
- Carson, Jamie L. and Michael H. Crespin. 2004. “The effect of state redistricting methods on electoral competition in United States House of Representatives races.” *State Politics & Policy Quarterly* 4(4):455–469.
- 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 David Cottrell. 2016. “Evaluating partisan gains from Congressional gerrymandering: Using computer simulations to estimate the effect of gerrymandering in the U.S. House.” *Electoral Studies* 44:329–340.
- Chen, Jowei and Jonathan Rodden. 2013. “Unintentional Gerrymandering: Political Geography and Electoral Bias in Legislatures.” *Quarterly Journal of Political Science* 8:239–269.
- 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.
- Cho, Wendy Tam and James G. Gimpel. 2012. “Geographic Information Systems and the Spatial Dimensions of American Politics.” *Annual Review of Political Science* 15:443–460.
- Cho, Wendy Tam and Yan Liu. 2016. “Toward a Talismanic Redistricting Tool: A Computational Method for Identifying Extreme Redistricting Plans.” *Election Law Journal* 15(4):351–366.

- Cirincione, Carmen, Thomas A. Darling and Timothy G. O'Rourke. 2000. "Assessing South Carolina's 1990s Congressional Districting." *Political Geography* 19(2):189–211.
- Duchin, Moon. 2018a. "Gerrymandering Metrics: How to Measure? What's the Baseline?". <https://arxiv.org/abs/1801.02064>.
- Duchin, Moon. 2018b. "Outlier Analysis for Pennsylvania Congressional Redistricting.". <https://mggg.org/uploads/md-report.pdf>.
- Edwards, Barry C., Michael H. Crespin, Ryan D. Williamson and Maxwell Palmer. 2017. "Institutional Control of Redistricting and the Geography of Representation." *Journal of Politics* 79(2):722–726.
- Ely, Jeffrey. 2019. "A Cake-Cutting Solution to Gerrymandering." p. 27. <https://www.jeffely.com/post/a-cake-cutting-solution-to-gerrymandering>.
- Fifield, Benjamin, Michael Higgins, Kosuke Imai and Alexander Tarr. 2020. "A New Automated Redistricting Simulator Using Markov Chain Monte Carlo." *Journal of Computational and Graphical Statistics* 0(0):1–14.
- Friedman, John N. and Richard T. Holden. 2008. "Optimal gerrymandering: sometimes pack, but never crack." *The American Economic Review* 98(1):113–144.
- Grofman, Bernard and Gary King. 2007. "The Future of Partisan Symmetry as a Judicial Test for Partisan Gerrymandering after *LULAC v. Perry*." *Election Law Journal* 6(1):2–35.
- Henderson, John A., Brian T. Hamel and Aaron M. Goldzimer. 2018. "Gerrymandering Incumbency: Does Nonpartisan Redistricting Increase Electoral Competition?" *The Journal of Politics* 80(3):1011–1016.
- Katz, Jonathan, Gary King and Elizabeth Rosenblatt. 2020. "Theoretical Foundations and Empirical Evaluations of Partisan Fairness in District-Based Democracies." *American Political Science Review* 114(1):1–15.
- King, Gary and Robert X. Browning. 1987. "Democratic Representation and Partisan Bias in Congressional Elections." *American Political Science Review* 81(4):1251–1273.
- Landau, Z., O. Reid and I. Yershov. 2009. "A fair division solution to the problem of redistricting." *Social Choice and Welfare* 32(3):479–492.
- Levitt, Justin. 2020. "All About Redistricting: Who Draws The Lines?" <https://redistricting.lls.edu/who.php>. Accessed 2/16/20.

- Magleby, Daniel and Daniel Mosesson. 2018. “A New Approach for Developing Neutral Redistricting Plans.” *Political Analysis* 26(2):147–167.
- Mayer, Kenneth. 2016. “Rebuttal Report: Response to Expert Reports of Sean Trende and Nicholas Goedert.”
URL: <https://perma.cc/DF5W-X8RU>
- McDonald, Michael P. 2004. “A Comparative Analysis of Redistricting Institutions in the United States, 2001-02.” *State Politics and Policy Quarterly* 4(4):371–395.
- McDonald, Michael P. 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.
- Miller, Peter and Bernard Grofman. 2013. “Redistricting Commissions in the Western United States.” *UC Irvine Law Review* 3(3):637–668.
- Moran, Patrick A.P. 1948. “The Interpretation of Statistical Maps.” *Journal of the Royal Statistical Society Series B*(10):243–251.
- Osborne, Martin J et al. 2004. *An introduction to game theory*. Vol. 3 Oxford university press New York.
- Pegden, Wesley, Ariel Procaccia and Dingli Yu. 2017. “A Partisan Districting Protocol with Provably Nonpartisan Outcomes.” <https://arxiv.org/abs/1710.08781>.
- Redistricting Commissions: What Works*. 2018. Technical report Brennan Center for Justice.
URL: <https://www.brennancenter.org/sites/default/files/analysis/Redistricting%20Commissions%20-%20What%20Works.pdf>
- Schaffner, Brian and Stephen Ansolabehere. 2015. 2010-2014 Cooperative Congressional Election Study Panel Survey. Technical report.
URL: <https://doi.org/10.7910/DVN/TOE8I1>
- Stephanopoulos, Nicholas. 2018. “The Causes and Consequences of Gerrymandering.” *William and Mary Law Review* 59(5).
URL: <https://scholarship.law.wm.edu/wmlr/vol59/iss5/14>
- Stephanopoulos, Nicholas and Eric McGhee. 2015. “Partisan Gerrymandering and the Efficiency Gap.” *University of Chicago Law Review* 82:831–900.
- Stephanopoulos, Nicholas and Eric McGhee. 2018. “The Measure of a Metric: The Debate over Quantifying Partisan Gerrymandering.” *Stanford Law Review* 70(5):1503–1566.