Supplementary Material for

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

Maxwell Palmer, Benjamin Schneer, and Kevin DeLuca

Contents

A 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 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).[1] 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.[2] In total, 31 states[3] 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 19 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.[4]

[1] Though commission-based states did register significantly higher approval rates than legislature-based states.

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

[3] Four of these 31 states have only have one Congressional districts and do not actually engage in Congressional redistricting.

[4] Selection methods for independent commissions vary significantly—in Arizona, Idaho, Montana, and

All told, 10 of the 19 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. One of the nine states with truly independent commissions, Alaska, uses its commission to draw state legislative districts but since it has only one Congressional district its commission does 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 to create a fairer map appealing to a neutral (and presumably more moderate) commission member.[5]

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 Crespin, 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.

[5]Only 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 et al., 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" (Brennan Center for Justice, 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" (Brennan Center for Justice, 2018).

Last of all, the establishment of an independent redistricting commission is not a realistic option 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.[6] Of the 24 states that do, eight have independent redistricting commissions already. Many 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 issues 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.

One of the largest obstacles to effective judicial intervention is that courts lack effective

[6]See `https://www.ncsl.org/research/elections-and-campaigns/initiative-and-referendum-processes.aspx`.

guidelines and standards to adjudicate partisan gerrymandering litigation. At a minimum, courts must decide (1) how to measure and evaluate partisan gerrymandering,[7] (2) how to compare multiple maps,[8] 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.[9] 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."[10] Only some states, however, have existing state laws

[7]There 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).

[8]Courts sometimes rely on simulated or counterfactual election results 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 et al., 2000; Altman and McDonald, 2011; Chen and Rodden, 2013; Chen and Cottrell, 2016; Cho and Liu, 2016; Magleby and Mosesson, 2018; Duchin, 2018; Fifield et al., 2020).

[9]Amendment 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/`.

[10]For 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 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.

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.[11] Notably, of the 41 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 degradation of political representation attributable to partisan gerrymandering.

A.2 Other Proposed Solutions

Some of the most promising alternative redistricting solutions 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 et al., 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 et al., 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; even if interest in such approaches did exist, the difficulties of implementing them 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, which has proven to be a stumbling block in the past

[11] See `https://www.ncsl.org/research/redistricting/free-equal-election-clauses-in-state-constitutions.aspx`

(Landau et al., 2009). Both approaches abstract from real-world geographies and do not place constraints on how voters are assigned to districts (Landau et al., 2009; Pegden et al., 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 more 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 convex[12] 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. We designed the Define-Combine Procedure 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.

[12]For 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
			Advisory (4)	Political (3)	Backup (3)	Commissions (9)
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[2]
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/`
[2] New York is classified by Levitt as an Independent Commission state; however, there are significant limits to the independence of the process because the stage legislature can consider but then vote not to accept maps created by the commission. After considering and rejecting two commission-produced maps, the legislature may offer its own maps.
*State only has one U.S. House district; state legislative redistricting authority used for classification.

B Analytic Solutions

In this section, we develop a simplified analytic model of redistricting under the unilateral process and Define-Combine Procedure, characterize the optimal behavior under each redistricting procedure, and sketch proofs for this behavior. Finally, based on the solutions to the model we compare URP and DCP outcomes along metrics used to analyze the performance and fairness of redistricting procedures in real-world redistricting.

B.1 General Setup

There are two political parties, party A and party B, who need to partition a state into N districts, indexed by d. There are P total voters across the state, who belong to either party A or party B and will vote for their party with certainty. Each party knows the partisan identity of all voters. In this model, when drawing maps political parties face no geographic, compactness or contiguity constraints on districts, meaning that parties can directly choose the voters to assign to each district regardless of where the voters live or the final shape of the district. In the analytic model, we also assume that each party has at least $\frac{1}{2N}$ voters (e.g., enough voters to comprise at least one subdistrict), and we assume that the values of P and N are large. We also allow parties to win districts with vote shares of $\frac{1}{2} + \epsilon$ voters.

Let V_A be the votes in the state for Party A, and $V_B = P - V_A$ be the votes for Party B. For each district d, let ν_{dA} and ν_{dB} be the number of votes in the district for each party. Political parties maximize their utility, which depends only on the share of seats won, $U_i = \frac{\sum_{d=1}^{N} s_d}{N}$, with s_d defined as in the main text:

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

B.2 Unilateral Redistricting

In unilateral redistricting, one party is in charge of drawing N districts as the first mover and the other party has no power to affect the outcome as the second mover. Suppose party A defines all districts,[13] which determines the vote shares of each party in each district and therefore determines the number of seats won by each party. Given no geographic constraints, party A in this case simply maximizes the number of districts, d, where $\nu_{dA} > \nu_{dB}$.

[13]We consider the optimal strategies of parties under each condition from the perspective of party A to make the explanations more concrete. By definition, the utility of party B is always: $U_B = 1 - U_A$ because this is a zero-sum game.

Theorem 1 *The maximum utility of the district-drawing party A under unilateral redistricting is the following piecewise function of its share of voters, V_A:*

$$U_A = \left\{ \begin{array}{ll} 2V_A & \textit{if } \frac{1}{2N} \leq V_A < \frac{1}{2} \\ 1 & \textit{if } V_A \geq \frac{1}{2} \end{array} \right\}$$

Proof: First, consider the case where $V_A \geq 0.5$. Because there are no geographic or compactness constraints, party A can win every district by dividing their share of V_A voters equally into N districts so that $\nu_{dA} > 0.5 > \nu_{dB}, \forall d$. In this case, party A wins all of the seats, which means $U_A = 1$, which is its maximum value.[14] Suppose that party A could achieve $U_A = 1$, the maximal utility, by creating at least one district such that $\nu_{db} > 0.5 > \nu_{dA}$. Party A would lose this district, implying that utility could be increased by creating a district such that $\nu_{dA} > 0.5 > \nu_{dB}$. However, since party A already has utility $U_A = 1$ it is a contradiction that utility could be increased by drawing a new district.

Second, consider the case where $V_A < 0.5$. Party A will not have enough voters to win all of the seats in this case but will still efficiently spread out voters across the districts to win a majority of votes in as many districts as possible. The number of voters in each district is P/N, the number of voters needed to win[15] a single district is $\frac{(P/N)}{2} = P/(2N)$, and the total number of voters that party A has is PV_A. Thus, the total number of districts where party A can win a majority of votes is equal to $\frac{PV_A}{(P/(2N))} = 2NV_A$, and their utility will be equal to their share of districts won, $\frac{2NV_A}{N} = 2V_A$.

Suppose a party deviated from this strategy, allocating $\frac{P+\theta}{2N}$ voters to a district while maintaining or increasing total utility. But then there must exist at least one other district with $\frac{P-k\theta}{2N}$ with $0 < k \leq 1$. However, by definition $P/(2N) > \frac{P-k\theta}{2N}$ are required to win the district. This marginal district cannot be lost while maintaining or increasing total utility, yielding a contradiction.

∎

[14] While there are many boundaries of districts that would result in party A winning all the seats in the case where $V_A > 0.5$, these all result in the same value of utility for party A, so in all cases $U_A = 1$ if $V_A > 0.5$.

[15] For clarity in the analytic solutions, we treat a 50% district the same as a 50%+1 district, so that a party can win when drawing a district with 0.5 share of the vote. In other words, ties are resolved in favor of the redistricting party. This is an accurate approximation when P is large.

B.3 DCP Redistricting

Under the Define-Combine Procedure, suppose party A "defines" $2N$ subdistricts, and then party B gets to "combine" pairs of those subdistricts to construct the final N required districts. Let districts be indexed by d, as before, and subdistricts be indexed by j.

Theorem 2 *The maximum utility of the first-mover or "define" party A under define-combine procedure redistricting is the following piecewise function of its share of voters, V_A:*

$$U_A = \left\{ \begin{array}{ll} V_A & \text{if } V_A < \frac{1}{3} \\ 4V_A - 1 & \text{if } \frac{1}{3} \leq V_A < \frac{1}{2} \\ 1 & \text{if } V_A \geq \frac{1}{2} \end{array} \right\}$$

Proof: Similar to the strategy in the unilateral case, when $V_A \geq 0.5$, the optimal strategy for party A is to define each of the $2N$ subdistricts so that $\nu_{jA} \geq 0.5, \forall j$. This would force party B in the define stage to create N final districts, where in every case $\nu_{dA} > \nu_{dB}$. So again, if $V_A \geq 0.5$, then $U_A = 1$ and $U_B = 0$ as party A wins every seat. The logic of the proof is identical to the unilateral case, as we can abstract away from the combine stage when all districts are identical.

In the case where $V_A < 0.5$, there are two strategies that party A could implement.

- **"Self Packing"**: Party A could "pack" itself by drawing as many subdistricts where $\nu_{jA} = 1$, in order to guarantee some minimum number of seats for itself. Any subdistrict with which a $\nu_{jA} = 1$ subdistrict is paired will turn out to be a district that party A wins.

- **"Cracking"**: Party A could draw as many subdistricts as possible that have $\nu_{jA} = 0.5$. If at least half of the subdistricts can be drawn with $\nu_{jA} = 0.5$, then some of these subdistricts will have to be paired together by party B and party A will win those final districts.

Proposition 1 *Under DCP, the "definer" will choose the strategy "self packing" with $U_A = V_A$ if $V_A < \frac{1}{3}$, and will choose the strategy "cracking" with $U_A = 4V_A - 1$ if $\frac{1}{3} < V_A < \frac{1}{2}$.*

Under the "self packing" strategy, the number of subdistricts that minority party A can draw with $\nu_{jA} = 1$ is equal to its number of voters, PV_A, divide by the number of voters needed to be 100% of a *sub*district, which is $\frac{P}{(2N)}$. Therefore, the maximum number of subdistricts that party A can pack with $\nu_{jA} = 1$ is $\frac{PV_A}{P/(2N)} = 2NV_A$.

However, in response to this strategy, party B in the "combine" stage will "pack" party A subdistricts together, creating many final districts where A wins 100% of the vote share and therefore wastes many of its votes. Specifically, party A will only end up winning half of the maximum possible seats it could win, $2NV_A/2 = NV_A$ seats, due to the optimal behavior of party B during the "combine" stage. With this strategy party A will always win the number of seats proportional to its share of the total voters, V_A, and this is true of this strategy at any value of V_A.[16]

When considering the "cracking" strategy, one reason it may be optimal for party A is that with enough subdistricts that are $\nu_{jA} = 0.5$, party B will have to combine some of these subdistricts and party B will end up losing those final districts while wasting many of their own votes. Specifically, as long as $\nu_{jA} = 0.5$ in more than half of the subdistricts, party B will have to combine some $\nu_{jA} = 0.5$ districts and party A is guaranteed to win at least one seat. More generally, however, party B can pair $\nu_{jA} = 0.5$ subdistricts with any district where $\nu_{jA} < 0.5$ and party A will waste many of their votes and lose that final district.[17] In fact, if only half or less of subdistricts can be drawn with $\nu_{jA} = 0.5$, then party B can combine each $\nu_{jA} = 0.5$ subdistrict with a subdistrict where $\nu_{jA} < 0.5$ and party A will lose every seat. Hence, the effectiveness of this strategy depends on how many subdistricts party A can draw with $\nu_{jA} = 0.5$, which is simply a function of its share of the overall vote, V_A.

The point where party A can draw at least half of the subdistricts with $\nu_{jA} = 0.5$ is when $NPV_A > \frac{P}{4N}$, which is equivalent to $V_A > \frac{1}{4}$. Therefore, if $V_A \leq \frac{1}{4}$ then party A wins no seats and $U_A = 0$ using the "cracking" strategy. However, if $V_A > \frac{1}{4}$, then the number of seats party A will ultimately win depends on how many subdistricts party A can draw with $\nu_{jA} = 0.5$.

In total, party A can draw $\frac{PV_A}{P/(4N)} = 4NV_A$ subdistricts where $\nu_{jA} = 0.5$ (this is twice the number of districts it can draw with $\nu_{jA} = 1$, as derived above). Given $V_A > \frac{1}{4}$, party B is not able to combine all $\nu_{jA} = 0.5$ subdistricts with $\nu_{jA} < 0.5$ subdistricts, and hence party A will win at least some districts. But party B can still create some winning final districts for itself during the combine stage. The number of districts that party A wins will be equal to one half of the *remaining* subdistricts where $\nu_{jA} = 0.5$, after party B runs out of subdistricts where $\nu_{jA} < 0.5$. The number of subdistricts where $\nu_{jA} < 0.5$ is equal to

[16] This is because party B's optimal behavior when faced with many $\nu_{jA} = 1$ subdistricts is to maximally waste party A's voters by combining all of the $\nu_{jA} = 1$ subdistricts together for a final map where A wins fewer districts with 100% of the vote.

[17] Because party A is the minority party with $V_A < 0.5$, there has to be at least one subdistrict where $\nu_{jA} < 0.5$.

$2N - 4NV_A$, and so the number of districts that party A will ultimately win is equal to $\frac{4NV_A-(2N-4NV_A)}{2} = 4NV_A - N$, which means $U_A = 4V_A - 1$:

$$U_A = \left\{ \begin{array}{ll} 0 & \text{if } V_A < \frac{1}{4} \\ 4V_A - 1 & \text{if } \frac{1}{4} < V_A < 0.5 \end{array} \right\}$$

Now, given the options of party A at the define stage, what is the optimal strategy and the utilities gained from each? At $V_A \geq 0.5$, party A can win all the seats so their utility is $U_A = 1$. So the only question is what party A does in the case where A is the minority party with $V_A < 0.5$. Party A can either "self pack" and get utility V_A or engage in "cracking" and get utility 0 if $V_A < \frac{1}{4}$ or $2V_A$ if $V_A \geq \frac{1}{4}$.

Clearly, when $V_A < \frac{1}{4}$ party A will "self pack". The question is whether the party changes strategies when $V_A \geq \frac{1}{4}$, and at what value of V_A might the switch occur. To solve this, we find a point V_A^*, where the utility from "cracking" surpasses the utility from the "self-pack" strategy. This point is the threshold value where party A switches strategies. To find V_A^*, set $V_A = 4V_A - 1$, which gives $V_A^* = \frac{1}{3}$. Therefore, when $V_A < \frac{1}{3}$, party A will choose the strategy "self packing" with $U_A = V_A$, and when $\frac{1}{3} < V_A < \frac{1}{2}$, party A will choose the strategy "cracking" with $U_A = 4V_A - 1$.

∎

The utility for party A's optimal strategy as the first mover using the Define-Combine Procedure is displayed as the blue line in Figure B.1, which shows a kink at $\frac{1}{3}$ where the first mover will switch from the "self-packing" to the "cracking" strategy. The red line in Figure B.1 shows the utility received by party A as the second mover, while the green line shows the average of the first- and second-mover utility that party A gets.

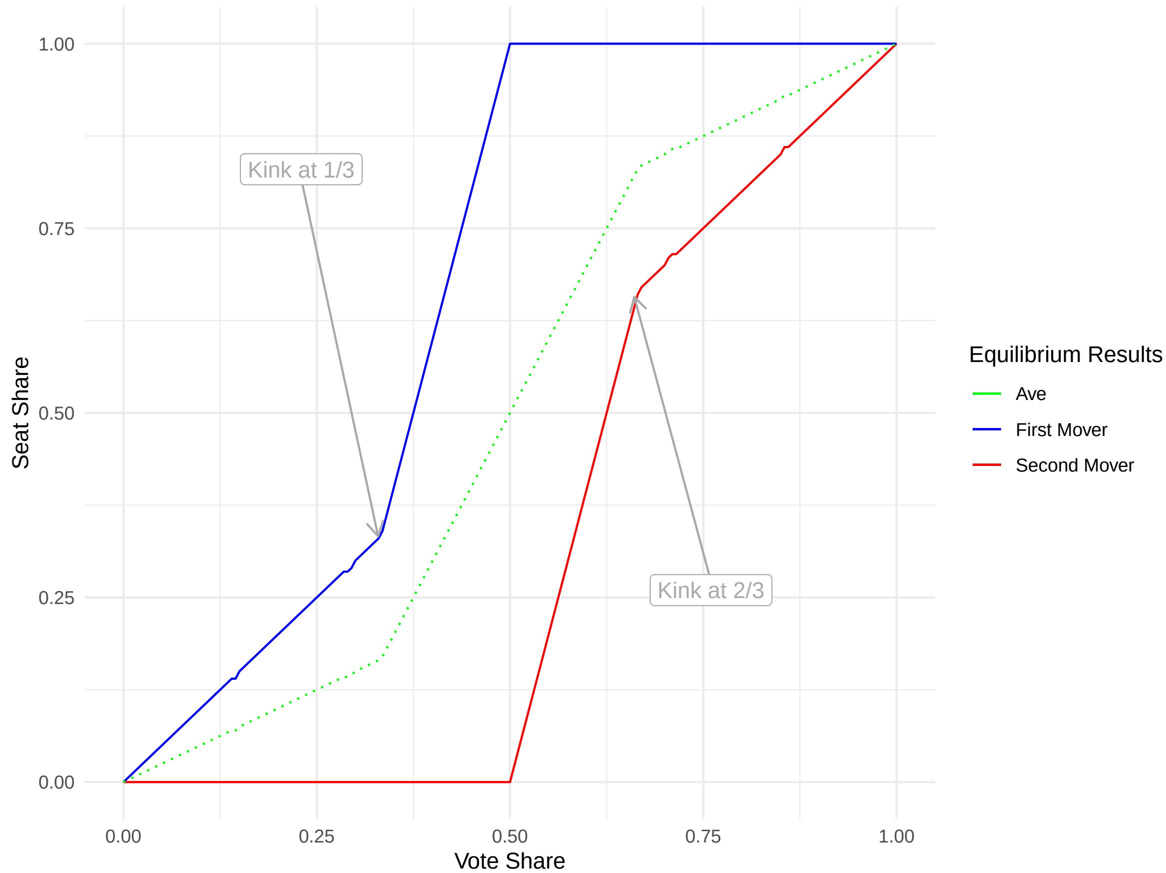

Figure B.1: **DCP Analytical Seats-Votes Curve:** This figure displays the seats-votes curve for a player under equilibrium behavior (e.g., both players behaving optimally) for the Define-Combine Procedure without geographic constraints. The blue line reports the seats won for a given statewide vote share for the first-moving player under DCP; the red line reports the seats won for a given statewide vote share for the second-moving player under DCP. Importantly, there are kinks in each curve at the vote share where it is optimal to switch from a "self-packing" strategy to a "cracking" strategy.

B.4 URP versus DCP: Intuition

To provide a graphical intuition, we compare the translation of votes into seats under optimal behavior for the first mover in DCP and the first mover in URP as well as for the second mover in DCP and the second mover in URP in Figure B.2. Importantly, the payoffs for the redistricters vary for statewide vote shares below 0.5 for first movers and above 0.5 for second movers.

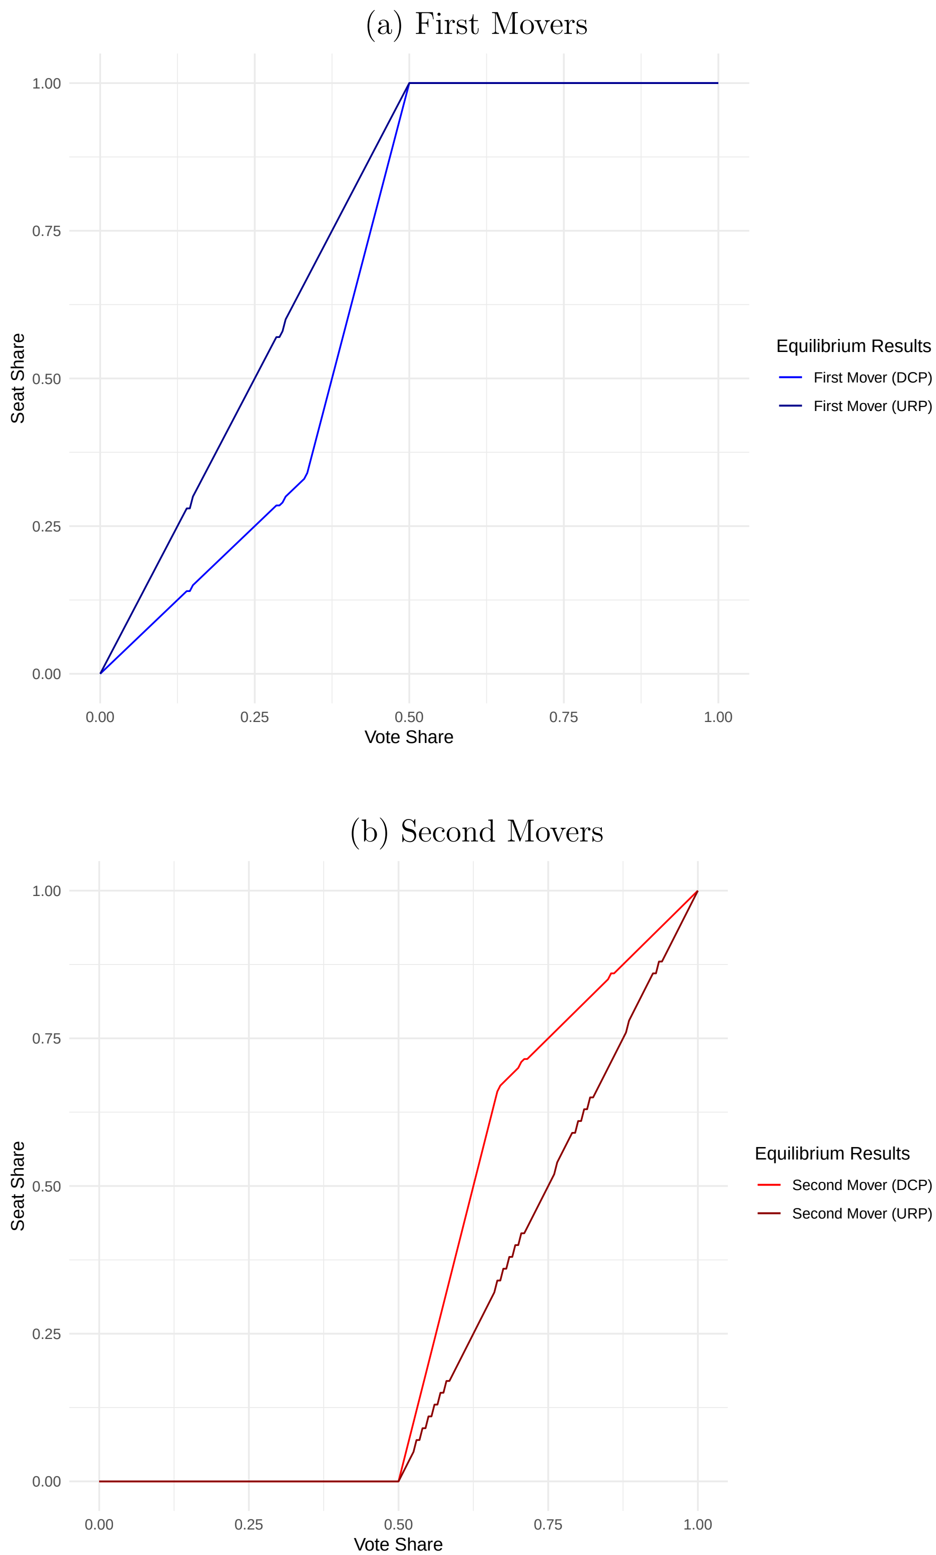

Figure B.2: **Comparing URP to DCP:** The figures above compare the equilibrium payoffs in terms of seat share for (a) first movers and (b) second movers under DCP and URP. For first movers, the payoffs deviate for statewide vote shares below 0.5. For second movers, the payoffs deviate for statewide vote shares above 0.5.

The analytical solutions illustrate that DCP's main advantage versus URP, in the case with no geographic constraints, comes in the form of protections for a majority party from an aggressive minority party. With no geographic constraints, we see that a minority party can win $2V$ seats under URP. However, when the minority party is instead the first mover under DCP, the party instead wins either a proportional share of seats (V, when $V < \frac{1}{3}$) or a share of seats that is strictly less than the URP seat share (when $\frac{1}{3} \leq V_A < \frac{1}{2}$), as is illustrated in Figure B.2(a).

On the other hand, when no geographic constraints are imposed, we find similar outcomes for DCP versus URP when the majority party is also the first mover since the majority party can always draw all subdistricts to contain a majority of its supporters. Note, however, that as soon as geographic constraints are added DCP will again tend towards moderating the extent of the advantage for the majority party. Specifically, DCP requires the first mover to draw smaller (sub)districts than under URP, which makes it always at least weakly more difficult to combine voters from disparate locations and helps limit the extent of packing and cracking that can occur under DCP. These properties, however, do not apply in the case with no geography. We discuss them further in Appendix Section J.

B.5 URP versus DCP: Analytical Benchmark Comparisons

Based on the analytical solutions for each model, we can also compare the performance of DCP to URP along many of the same metrics as for our real-world simulations.

First, we consider the metric of Definer's Advantage (δ), the payoff accrued from being the first mover as compared to the second mover for each redistricting procedure. Thus, we compare the equilibrium payoff for first movers minus second movers for each redistricting procedure. Recall that under DCP and URP with no geographic constraints, payoffs for each party depend on the statewide vote share.

Under DCP we have:

$$\delta = \left\{ \begin{array}{ll} V_A & \text{if } V_A < \frac{1}{3} \\ 4V_A - 1 & \text{if } \frac{1}{3} \leq V_A < \frac{1}{2} \\ 1 - (4V_A - 2) & \text{if } \frac{1}{2} \leq V_A < \frac{2}{3} \\ 1 - V_A & \text{if } V_A \geq \frac{2}{3} \end{array} \right\}$$

Under URP we have:

$$\delta = \left\{ \begin{array}{ll} 2V_A & \text{if } V_A < \frac{1}{2} \\ 1 - (-1 + 2V_A) & \text{if } V_A \geq \frac{1}{2} \end{array} \right\}$$

By inspection, it is clear that for all values of V_A, the Definer's Advantage (δ) has larger magnitude under URP than DCP. That is, the difference between moving first and moving second matters more under unilateral redistricting, with the first mover being advantaged.

Next, we consider the extent to which the optimal solutions under these two procedures deviate from the benchmark of proportionality. A proportional translation of seats into votes simply means that for a vote share V_A party A would earn a seat share also equal to V_A. To calculate the deviation from proportionality, we can find the difference between the seat share allocation for a party behaving optimally and a proportional seat share. Thus, under DCP as first mover, we have:

$$\text{Deviation from Proportionality} = \left\{ \begin{array}{ll} 0 & \text{if } V_A < \frac{1}{3} \\ 3V_A - 1 & \text{if } \frac{1}{3} \le V_A < \frac{1}{2} \\ 1 - V_A & \text{if } V_A \ge \frac{1}{2} \end{array} \right\}$$

Under DCP as second mover, we have:

$$\text{Deviation from Proportionality} = \left\{ \begin{array}{ll} V_A & \text{if } V_A < \frac{1}{2} \\ 3V_A - 2 & \text{if } \frac{1}{2} \le V_A < \frac{2}{3} \\ 0 & \text{if } V_A \ge \frac{2}{3} \end{array} \right\}$$

Finally, under URP, we have:

$$\text{Deviation from Proportionality} = \left\{ \begin{array}{ll} V_A & \text{if } V_A < \frac{1}{2} \\ 1 - V_A & \text{if } V_A \ge \frac{1}{2} \end{array} \right\}$$

For all values of V_A, the Deviation from Proportionality for DCP is less than or equal to the Deviation from Proportionality for URP. When V_A is less than one half, URP and DCP for the second mover are equivalent. However, when vote share V_A is greater than one half, DCP for the second mover is either more proportional or equivalent to the proportional solution depending on the exact vote share. Similarly, comparing first-mover DCP to URP yields a similar result, but with the improvements from DCP for this measure coming from vote shares that are more than one half.

Next, we examine how each measure deviates from a fairness target, which is defined as the midpoint between the outcomes for a party being the first mover and the second mover.

Under DCP, we have:

$$\text{Deviation from Fairness Target} = \left\{ \begin{array}{ll} \frac{V_A}{2} & \text{if } V_A < \frac{1}{3} \\ \frac{4V_A - 1}{2} & \text{if } \frac{1}{3} \leq V_A < \frac{1}{2} \\ \frac{4V_A - 3}{2} & \text{if } \frac{1}{2} \leq V_A < \frac{2}{3} \\ \frac{1 - V_A}{2} & \text{if } V_A \geq \frac{2}{3} \end{array} \right\}$$

Under URP, the Deviation from Fairness Target is equal to V_A for all vote shares. As a result, by inspection, the Deviation from the Fairness Target under URP is always less than or equal to the Deviation under DCP for all values of V_A.

When we evaluate the Partisan Bias for each procedure, we find that the levels of partisan bias are equivalent for URP and DCP. Specifically, at vote share $V_A = 0.5$, the seat share is the same regardless of redistricting procedure.

To sum up, we find that across a range of different metrics assessing the performance of a redistricting procedure, Define-Combine performs better than or equal to Unilateral redistricting across all metrics and regardless of vote share.

B.6 Imposing Geographic Constraints

Two types of geographic constraints arise from a contiguity requirement: (1) some subdistricts can be paired with only a subset of other subdistricts; (2) some voters can be paired only with a subset of other voters in the state.

Generally, researchers have found that there are not closed form expressions to characterize the optimal redistricting behavior while imposing arbitrary geographic constraints (Liu et al., 2016). The number of total potential maps in any real-world redistricting context is also too large and makes exhaustive searches computationally infeasible, even with supercomputers (see subsection I). As a result, there is a large literature on simulation methods to approximate solutions. To this end, the primary results in our paper apply a simulation algorithm to characterize optimal behavior in real-world states with geographic constraints, and in subsection H, we use simulations and a simplified grid map to assess how various constraints affect the relative performance of DCP.

C 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.[18] 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.
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.[19] Across 10 separate chains for each state, we ultimately simulate and score 200,000 plans, selecting the most extreme result for each party.[20]

Second, we use a "nested shortburst" algorithm to simulate the Define-Combine Procedure. For each state, chamber, and party we run 20 separate sets of simulations. In each simulation, we:

1. Generate a random starting map of $2N$ districts.
2. Run the shortburst burst algorithm 100 times, with 20 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 find the best combine response for the combining party in two steps:

1. Generate 50 random pairings of the subdistricts.
2. Choose the pairing that maximizes seats for the combining party. If more than one pairing maximizes seats, then we choose up to five unique pairings.
3. We each pairing, we run 25 shortbursts of the pairing algorithm and generate 10 different combinations in each burst.

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

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

[20]If 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.

For each chain, we save the final define and combine districts.

The second implementation of the algorithm maximizes seats for the second party across 300–1,300 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.[21] 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.[22]

All of the simulations have constraints for nearly equal population (we allow for 1% deviations), contiguity, and a reasonable level of compactness.

Section C.1 presents pseudo code that outlines the algorithms described above. The full R code and data to run our simulations is included with the replication data.

[21]For 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. This is both fully complete and easier to implement for smaller states, but not feasible for states with more districts.

[22]An 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.

C.1 Pseudo Code for Simulations

Unilateral Redistricting

```
load input map data for N districts
load base maps for N districts
select a random base map
run redist_shortburst:
    - 2000 steps, with 10 maps per step
    - maximize number of seats won + maximum vote share in districts with
      vote shares less an 50%
save final map
```

Define-Combine Procedure

```
load input map data for 2N districts
load base maps for 2N districts
select a random base map
run redist_shortburst for define-stage:
    - 100 steps, with 20 maps per step
    - maximize number of seats won by defining party
    - to score seats won for each map:
        - if N>7:
            - create new map object where each unit is a subdistrict
            - calculate adjacencies of subdistricts
            - generate 50 random pairings of adjacent subdistricts
            - choose subdistrict pairing that is best for combining party
              (up to 5)
            - run redist_shortburst for combine-stage:
                - 25 steps, 10 maps per step
                - minimize seats won by defining party
        - if N<=7:
            - create new map object where each unit is a subdistrict
            - calculate adjacencies of subdistricts
            - enumerate all possible pairings of adjacent subdistricts
            - select pairing that minimizes seats won by defining party
save final map
```

D Measurement of Redistricting Plans

Consider an electoral system with seats-votes function $S_M(\nu_1, \ldots, \nu_N)$ for a map M, which takes as an input district-level vote shares $\nu_1, \ldots, \nu_N$ and yields as an output a seat share.[23] The statewide vote share V is the average of district-level vote shares (importantly, elections with identical statewide average vote share V but different realizations of $\nu_1, \ldots, \nu_N$ could result in different winning candidates). Conditioning on a statewide vote share V, we can find the average seat share by taking the expected value of the function, over the joint distribution for $\nu_1, \ldots, \nu_N$—e.g., $E\left(S_M(\nu_1, \ldots, \nu_N) \mid V\right) = S_M(V)$. 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 $\delta^U_{.5} = S_{\widetilde{M}_A}(0.5) - S_{\widetilde{M}_B}(0.5)$. Similarly, $\delta^D_{.5} = S_{\widehat{M}_A}(0.5) - S_{\widehat{M}_B}(0.5)$ indicates the definer's advantage under DCP. A large positive value for $\delta^U_{.5}$ indicates that the party controlling the URP can reap a significant electoral advantage through gerrymandering; a large positive value for $\delta^D_{.5}$ 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.

[23]Notation used here is similar to Katz, King, and Rosenblatt (Katz et al., 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, $|S_{\widetilde{M}_A} - V| > |S_{\widehat{M}_A}(0.5) - V|$ denotes a case where the URP map for Party A deviates from proportionality more than the DCP map. Note that while proportionality has some clear benefits as an intuitive benchmark, lack of proportionality does not necessarily imply a lack of fairness. Courts have "sometimes recognized proportionality as a virtue, but have explicitly rejected strict proportionality as a standard" (Bernstein and Duchin, 2017).

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 et al. (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 $|S_M(V) - \theta| < |S_N(V) - \theta|$.

Fifth, to determine the Efficiency Gap, we calculate the number of "wasted votes" in each district for each party. We define wasted votes for Party A W_A as the total of all votes in districts lost by Party A as well as all votes over 0.5 in districts won by Party A. Then, Efficiency Gap $= \frac{W^A - W^B}{N}$.

E Additional Simulation Results

E.1 Simulation Tables

State	Seats	R Alone	R Then D	D Then R	D Alone
AL	7	0	0	1	2
AR	4	0	0	0	1
FL	28	4	12	11	18
GA	14	2	5	6	11
IL	17	9	11	12	14
KS	4	0	0	2	2
KY	6	0	0	1	2
LA	6	0	0	1	3
MA	9	9	9	9	9
MO	8	0	2	3	5
MS	4	0	0	2	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	0	3	4
TN	9	0	1	2	4
TX	38	8	19	17	28
WV	2	0	0	0	0

Table E.1: **Legislatures Redistricting:** Simulation Results with columns 3 through 6 reporting the number of Democratic seats won.

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	7	11
NH	2	1	2	2	2
NY	26	15	18	22	26
PA	17	4	7	8	12
VA	11	5	6	8	10
WI	8	1	3	4	7

Table E.2: **Court or Special Master Redistricting:** Simulation Results with columns 3 through 6 reporting the number of Democratic seats won.

State	Seats	R Alone	R Then D	D Then R	D Alone
AZ	9	1	4	5	8
CA	52	41	47	46	52
CO	8	3	5	6	8
ID	2	0	0	0	0
MI	13	2	6	6	11
MT	2	0	0	0	0
WA	10	4	6	8	10

Table E.3: **Independent Commission Redistricting:** Simulation Results with columns 3 through 6 reporting the number of Democratic seats won.

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	2	2	4
ME	2	1	1	2	2
NJ	12	6	9	10	12
OH	15	1	5	6	11
UT	4	0	0	2	2

Table E.4: **Political Redistricting:** Simulation Results with columns 3 through 6 reporting the number of Democratic seats won.

E.2 Simulation Results After Uniform Swing

We conducted a set of simulations for competitive states in the 2020 Presidential election (two-party vote shares from 45–55%). For each state, we applied a uniform swing calculated using the state's results in the 2020 Presidential election. Figure E.1 shows the states included and the results from both URP and DCP.

Of the 174 seats in the thirteen states included, Democrats win 136 (78.2%) when drawing the maps unilaterally in every state. When Republicans draw all maps unilaterally, Democrats win 28 seats (16.1%). The simulations reveal a possible swing of 108 seats between parties through unilateral redistricting.

When Democrats draw the Define-stage map, they win 83 (47.7%) seats; when Republicans draw the Define-stage map, Democrats win 79 (45.4%) 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 12 seats, eliminating almost 90% of the swing in seats theoretically possible under unilateral partisan control of redistricting in these states.

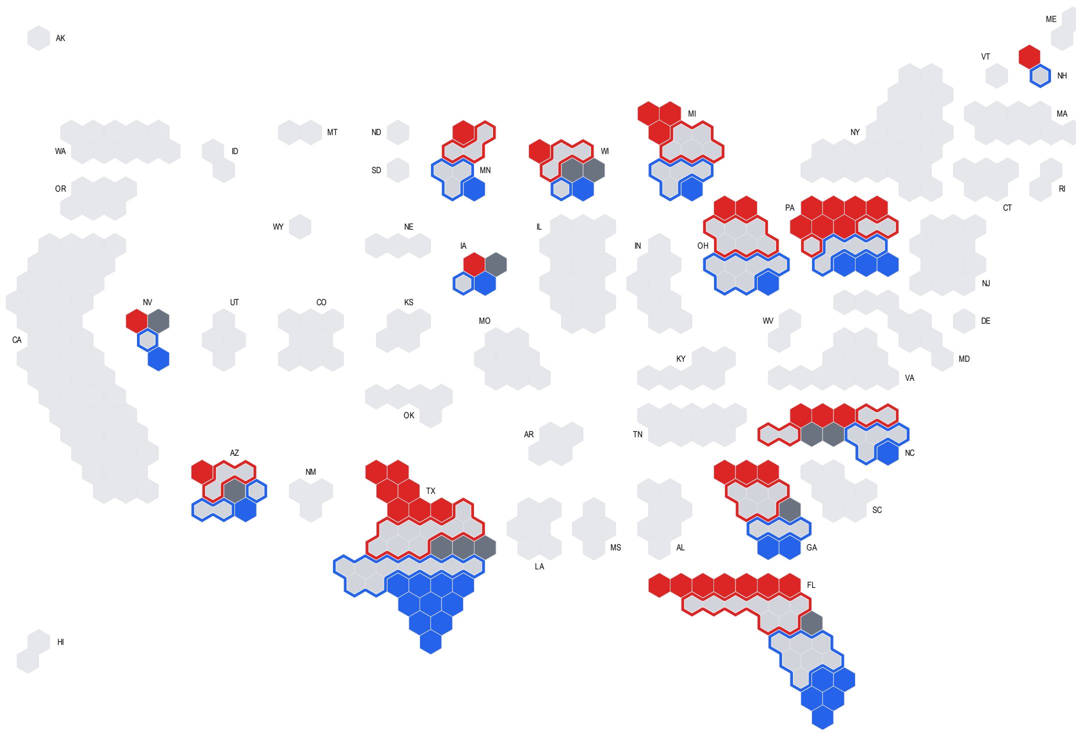

Figure E.1: **Uniform Swing Simulations, Map of Differences in Results by Party:** This map shows how the partisan division of competitive states (only states with two-party vote shares from 45-55%) differ under URP and DCP. Hexagons colored blue and red are seats that are always won by Democrats and Republicans, respectively, in both methods regardless of which party controls the process. Hexagons shaded in dark or light gray could be won by either party if they controlled URP. Hexagons in dark gray could be won by either party under DCP. The hexagons outlined in blue and red would be won by Democrats and Republicans, respectively, under DCP.

State	Seats	R Alone	R then D	D then R	D Alone
FL	28	5	13	12	21
GA	14	2	5	6	11
NV	4	1	2	3	3
TX	38	10	22	19	31

Table E.5: **Legislatures Redistricting:** Simulation Results After Uniform Swing with columns 3 through 6 reporting the number of Democratic seats won.

State	Seats	R Alone	R then D	D then R	D Alone
MN	8	1	4	4	7
NC	14	1	5	7	11
NH	2	0	1	1	1
PA	17	3	7	7	10
WI	8	1	2	4	7

Table E.6: **Court or Special Master Redistricting:** Simulation Results After Uniform Swing with columns 3 through 6 reporting the number of Democratic seats won.

State	Seats	R Alone	R then D	D then R	D Alone
AZ	9	1	4	5	8
MI	13	1	5	5	10

Table E.7: **Independent Commission Redistricting:** Simulation Results After Uniform Swing with columns 3 through 6 reporting the number of Democratic seats won.

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 E.8: **Political Redistricting:** Simulation Results After Uniform Swing with columns 3 through 6 reporting the number of Democratic seats won.

E.3 Simulation Results After National Swing

We conducted a set of simulations for a hypothetical case where the two-party vote is 50% nationwide. Starting from the 2020 presidential election, we applied a common swing that increased the percentage of the Republican vote by 2.27 percentage points in each precinct.[24] Figure E.2 shows the results from both URP and DCP.

Of the 429 seats in the thirteen states included, Democrats win 316 (73.7%) when they draw the maps unilaterally in every state. When Republicans draw all maps unilaterally, Democrats win 108 seats (25.2%). The simulations reveal a possible swing of 208 seats between parties through unilateral redistricting.

When Democrats draw the Define-stage map, they win 225 (52.4%) seats; when Republicans draw the Define-stage map, Democrats win 194 (45.2%) 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 45 seats, eliminating almost 80% of the swing in seats theoretically possible under unilateral partisan control of redistricting in these states.

[24]The 2.27 percentage point swing was calculated using the two-party vote across the 44 states in our analysis.

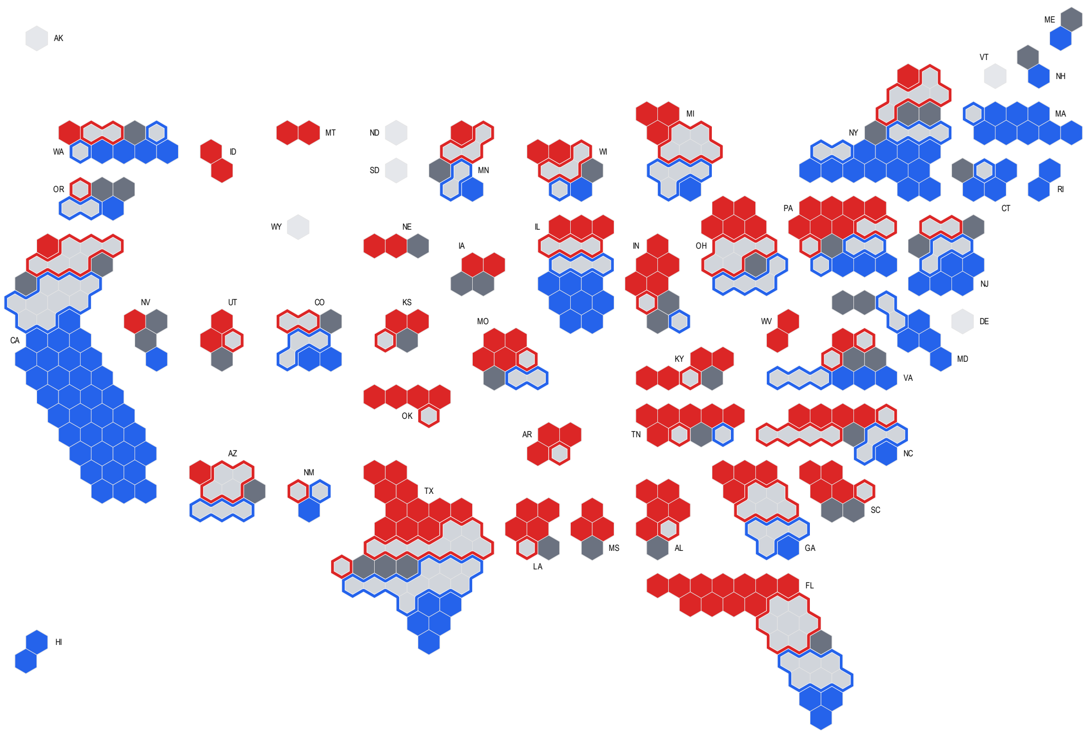

Figure E.2: **National Swing Simulations, Map of Differences in Results by Party:** This map shows how the partisan division of states differ under URP and DCP. Hexagons colored blue and red are seats that are always won by Democrats and Republicans, respectively, in both methods regardless of which party controls the process. Hexagons shaded in dark or light gray could be won by either party if they controlled URP. Hexagons in dark gray could be won by either party under DCP. The hexagons outlined in blue and red would be won by Democrats and Republicans, respectively, under DCP.

E.4 Simulation Results for State Legislatures

While our results in this article focus on how DCP could apply to congressional redistricting, the same logic applies to state legislatures, as well as other district-based bodies, such as city councils, county commissions, or elected judges. Here, we present results for three state senates: North Carolina, Pennsylvania, and Wisconsin. In each chamber we find that DCP substantially reduces potential bias compared to unilateral redistricting.

In the North Carolina State Senate, 26 of the 50 seats could swing based on party control of unilateral redistricting. In contrast, only two are up for grabs under DCP. Similarly in Pennsylvania, 20 of the 50 seats could swing under URP, and only one under DCP. In Wisconsin, 15 of the 33 seats could swing under URP, but none under DCP.

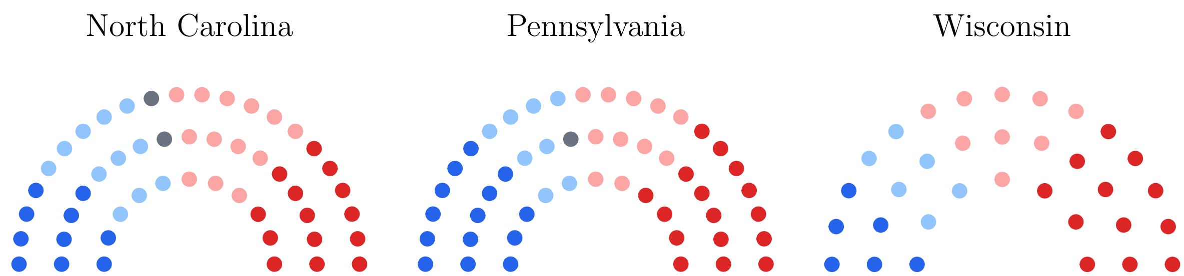

Figure E.3: **State Legislative Simulations Results:** This map shows how the partisan division of state senates differ under URP and DCP. Circles colored dark blue and red are seats that are always won by Democrats and Republicans, respectively, in both methods regardless of which party controls the process. Circles shaded in light blue, light red, or gray could be won by either party if they controlled URP. Circles in dark gray could be won by either party under DCP. The circles in light blue and light red would be won by Democrats and Republicans, respectively, under DCP.

E.5 Simulation Results When Randomizing First Mover

We report results for the average outcome if states were to randomize the first mover in a redistricting procedure. To estimate this, consider a large number of repeated draws of the first mover for both URP and DCP. As the sample grows large, the result converges to the average of the outcomes when each party moves first. The interpretation and how to compare these results is markedly different than previous tables. Specifically, one can think of the average of unilateral districting outcomes (URP) as a fairness target exhibiting a roughly neutral allocation of seats given the distribution of voters and geography of a state. Thus, the DCP average would ideally be located close to the average under unilateral redistricting.

State	**Seats**	**URM Ave**	**DCP Ave**
AL	7	1.0	0.5
AR	4	0.5	0.0
FL	28	11.0	11.5
GA	14	6.5	5.5
IL	17	11.5	11.5
KS	4	1.0	1.0
KY	6	1.0	0.5
LA	6	1.5	0.5
MA	9	9.0	9.0
MO	8	2.5	2.5
MS	4	1.0	1.0
NE	3	0.5	0.5
NM	3	2.0	2.5
NV	4	2.0	2.5
OK	5	0.5	0.0
OR	6	4.0	4.0
RI	2	2.0	2.0
SC	7	2.0	1.5
TN	9	2.0	1.5
TX	38	18.0	18.0
WV	2	0.0	0.0

Table E.9: **Legislatures Redistricting:** Simulation Results with Randomized First Mover. Columns 3 and 4 report the average number of Democratic seats won.

State	Seats	URM Ave	DCP Ave
CT	5	4.5	5.0
MD	8	6.5	7.0
MN	8	4.5	4.0
NC	14	6.0	6.0
NH	2	1.5	2.0
NY	26	20.5	20.0
PA	17	8.0	7.5
VA	11	7.5	7.0
WI	8	4.0	3.5

Table E.10: **Court or Special Master Redistricting:** Simulation Results with Randomized First Mover. Columns 3 and 4 report the average number of Democratic seats won.

State	Seats	URM Ave	DCP Ave
AZ	9	4.5	4.5
CA	52	46.5	46.5
CO	8	5.5	5.5
ID	2	0.0	0.0
MI	13	6.5	6.0
MT	2	0.0	0.0
WA	10	7.0	7.0

Table E.11: **Independent Commission Redistricting:** Simulation Results with Randomized First Mover. Columns 3 and 4 report the average number of Democratic seats won.

State	Seats	URM Ave	DCP Ave
HI	2	2.0	2.0
IA	4	1.5	1.0
IN	9	2.0	2.0
ME	2	1.5	1.5
NJ	12	9.0	9.5
OH	15	6.0	5.5
UT	4	1.0	1.0

Table E.12: **Political Redistricting:** Simulation Results with Randomized First Mover. Columns 3 and 4 report the average number of Democratic seats won.

F 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 F.1 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).

For partisan bias, DCP clearly reduces bias compared to the adopted plans in states with legislature drawn maps (from 10.4% to 6%). 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 under these circumstances, 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 5 in the main text, which plots 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 almost all states that did not ultimately default to maps drawn by a court or special master, the outcome from DCP is closer to matching the fairness target than the actual outcome achieved in the state. The figure illustrates how each actual redistricting outcome, represented by a point where a line with an arrow originates, deviate further from a

Metric	Category	URP	DCP	Adopted Plan
Definer's Advantage	Court or Special Master	0.505	0.121	0.131
Definer's Advantage	Independent Commission	0.396	0.052	0.146
Definer's Advantage	Legislature	0.446	0.118	0.204
Definer's Advantage	Political	0.542	0.146	0.292
Deviation from Proportionality (Actual State Party Control)	Court or Special Master	0.263	0.125	0.107
Deviation from Proportionality (Actual State Party Control)	Independent Commission	0.289	0.181	0.172
Deviation from Proportionality (Actual State Party Control)	Legislature	0.258	0.154	0.217
Deviation from Proportionality (Actual State Party Control)	Political	0.285	0.185	0.268
Deviation from Fairness Target (Actual State Party Control)	Court or Special Master	0.253	0.076	0.071
Deviation from Fairness Target (Actual State Party Control)	Independent Commission	0.198	0.031	0.073
Deviation from Fairness Target (Actual State Party Control)	Legislature	0.223	0.070	0.105
Deviation from Fairness Target (Actual State Party Control)	Political	0.271	0.073	0.146
Partisan Bias	Court or Special Master	0.309	0.074	0.047
Partisan Bias	Independent Commission	0.364	0.091	0.061
Partisan Bias	Legislature	0.286	0.060	0.104
Magnitude of Efficiency Gap	Court or Special Master	0.281	0.095	0.056
Magnitude of Efficiency Gap	Independent Commission	0.248	0.094	0.089
Magnitude of Efficiency Gap	Legislature	0.214	0.113	0.128
Magnitude of Efficiency Gap	Political	0.250	0.129	0.156

Table F.1: **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%.

fairness target than the simulated DCP outcome in most states where redistricting occurred by independent commissions, legislatures, or politician-controlled commissions. The two exceptions we observe are in North Carolina (Court or Special Master) and Washington (Independent Commission).

G State-by-State Shortburst Comparisons, 2020 Presidential Election Democratic Vote Share

Figures G.1, G.2, and G.3 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,[25] 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.

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 ten of eleven Congressional districts. In contrast, Republicans could draw maps where Democrats win only five Congressional districts—a gap of 5 seats. Under DCP, the gap shrinks dramatically to 2 seats depending on the party that moves first. Overall, one state exhibits a four seat gap and one state exhibits a three seat gap; ten of forty-four states exhibit a two seat gap, nineteen states exhibit a one seat gap, and thirteen states exhibit a zero seat gap under DCP. In twenty-nine of the thirty-two states where URP produces more than a one seat difference in seats depending on the redistricting party, DCP narrows the range of outcomes as compared to URP, regardless of which party defines subdistricts and which party combines them. Only California, Florida and Texas exhibit a second-mover or combiner's advantage.

[25]These 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

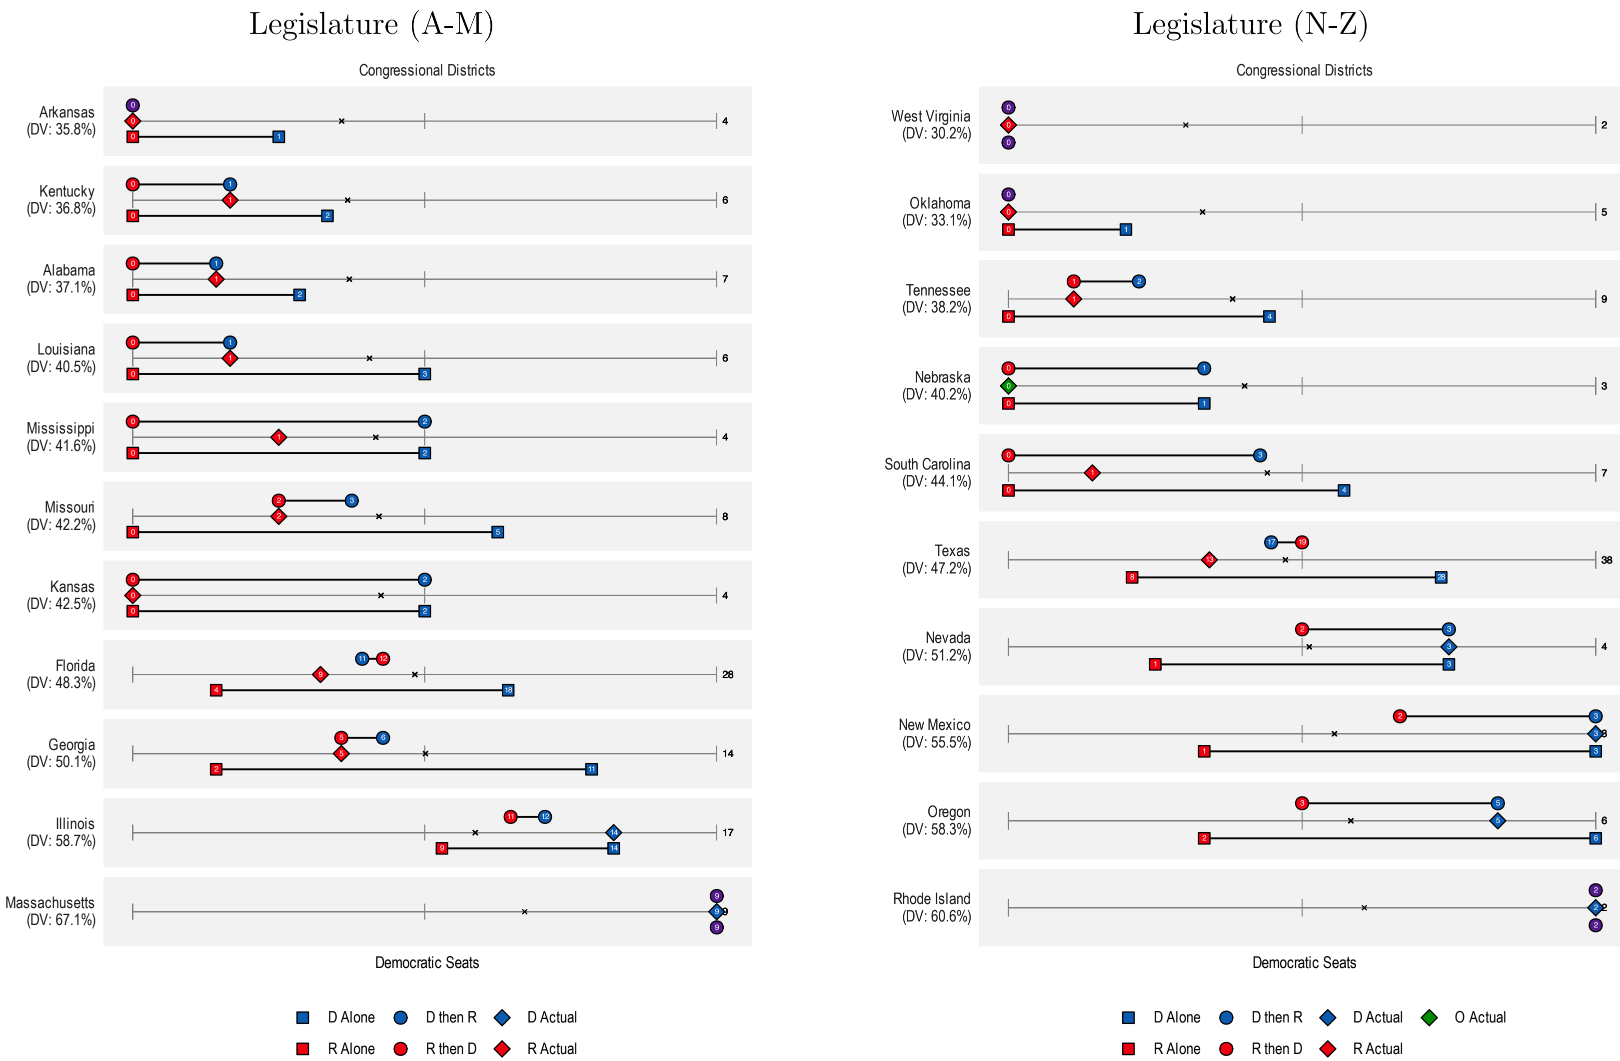

Figure G.1: **State by State Simulation Benchmarks:** This figure displays DCP (labelled "D/R then R/D") results to several other benchmarks for state Congressional districts where legislatures performed redistricting in the post-2020 redistricting cycle. Points labelled "D/R Alone" denote unilateral 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.

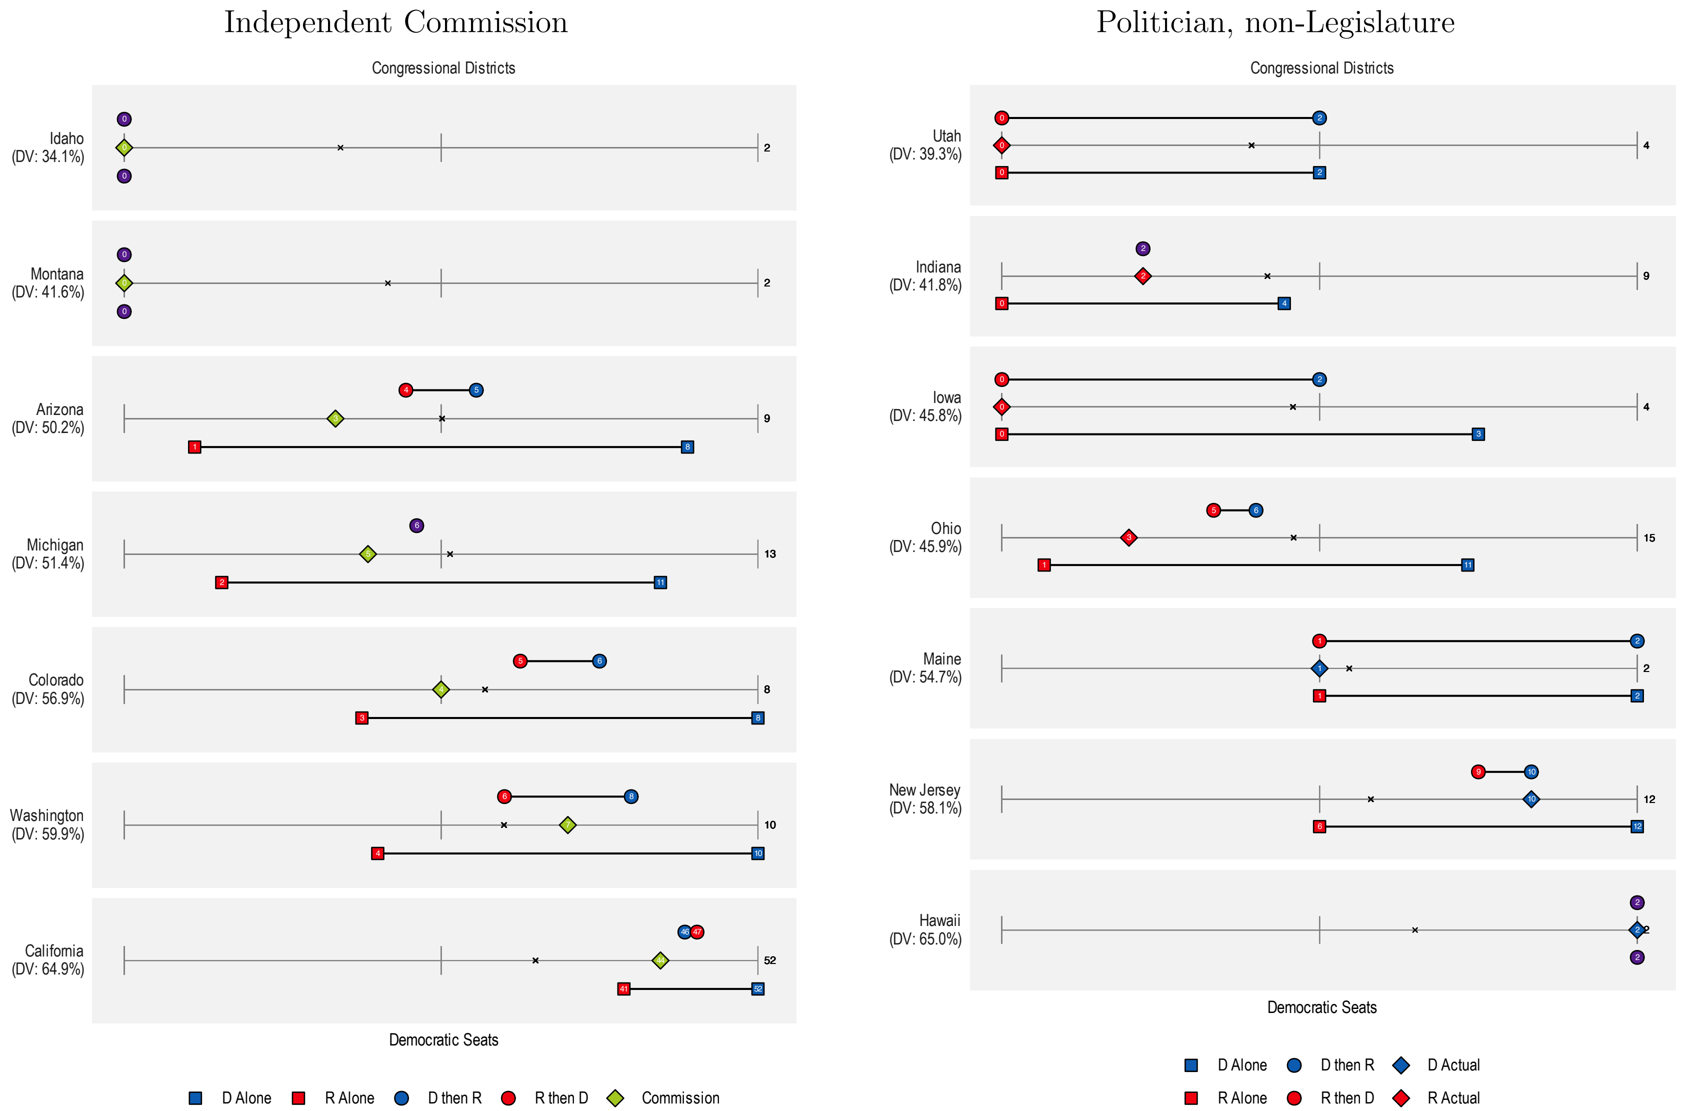

Figure G.2: **State by State Simulation Benchmarks:** This figure compares DCP (labeled "D/R then R/D") results to several other benchmarks for state Congressional districts where commissions or politicians guided redistricting in the post-2020 redistricting cycle. Points labelled "D/R Alone" denote 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.

Court or Special Master

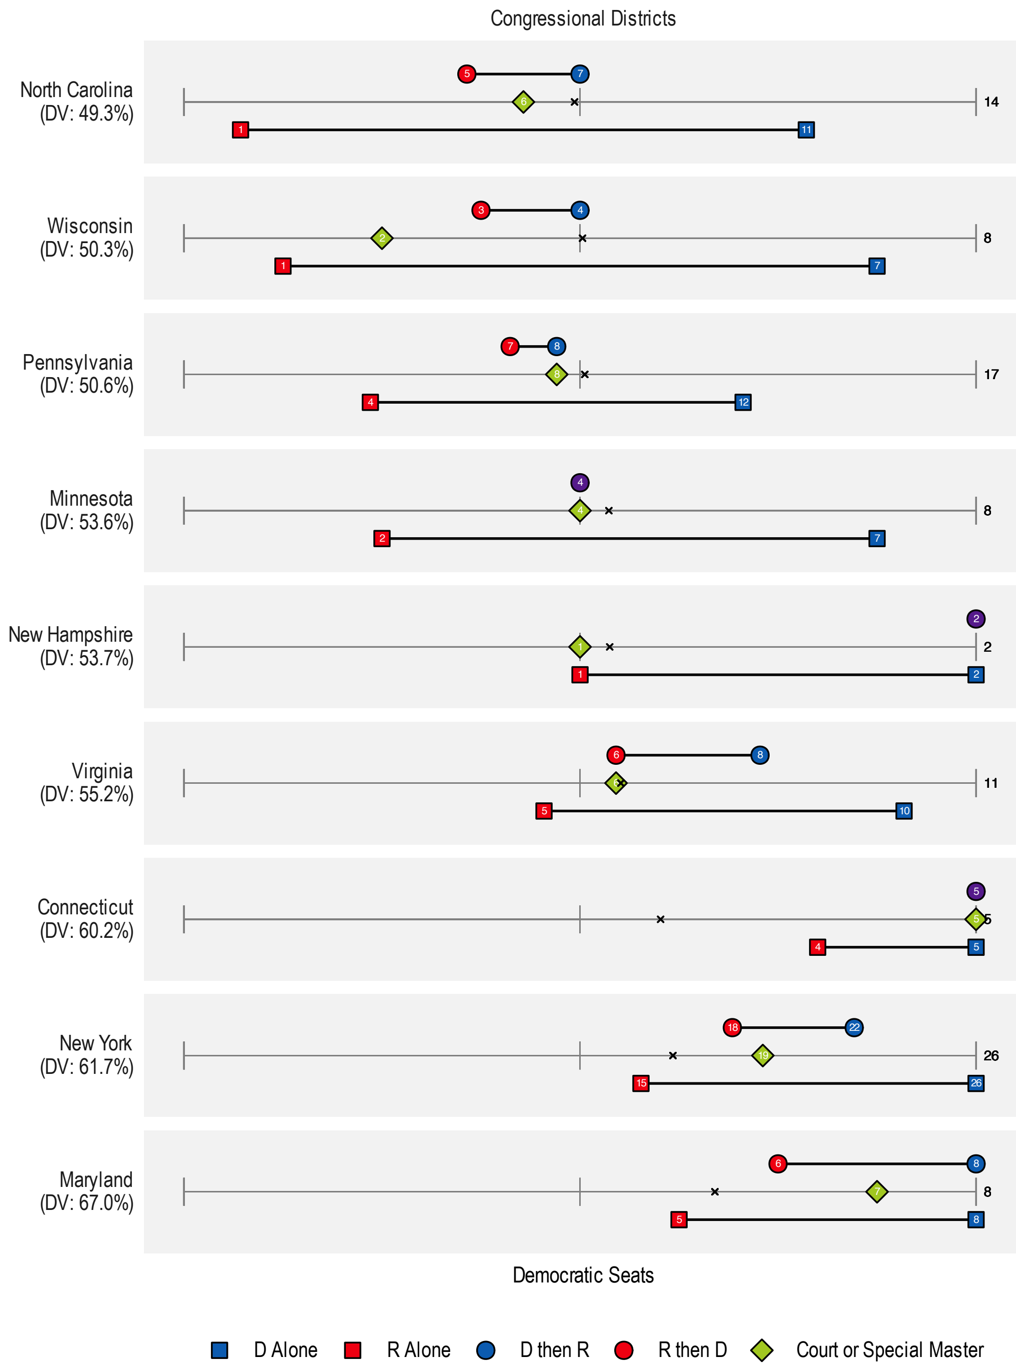

Figure G.3: **State by State Simulation Benchmarks:** This figure displays DCP (labeled "D/R then R/D") results to several other benchmarks for state Congressional districts where a court or special master guided redistricting in the post-2020 redistricting cycle. Points labelled "D/R Alone" denote simulated 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.

H Grid Maps

To test the robustness of DCP and to explore its properties while still accounting for geography, we use simulated grid maps that allow for (1) varying levels of statewide partisanship, (2) varying objectives for the redistricting parties, and (3) voter distributions with varying degrees of geographic clustering. We explore how each of these extensions influences the properties of the maps chosen under URP and DCP. Overall, DCP continues to reduce gerrymandering dramatically when varying the degree of statewide partisanship, the parties' objective functions, and geographic clustering.

H.1 Grid Map Simulations

To simulate grid maps, we define a grid of equal-population precincts, randomly assign partisanship (Party A or Party B) to voters in each precinct while fixing map-wide partisan composition at a specific value, and then solve for the URP and DCP maps that each party would select in redistricting. We report the average across the randomly-generated voter distributions to characterize the results. We repeat this approach for a rectangular 5x6 grid map (e.g., thirty precincts with ten subdistricts and five final districts), which allows us to evaluate all possible Define and Combine stage maps, as well as a larger hexagonal map with 150 precincts, thirty subdistricts and fifteen districts.

Here we recount the specifics of our approach. 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.[26] 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.

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

(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.

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

H.2 Varying Partisan Composition of Voters

DCP produces significant improvements in terms of reduced Definer's Advantage across a range of different statewide vote shares. Grid simulations show that partisan advantage conferred to a unilateral redistricter peaks when the parties evenly split the statewide vote. In contrast, the partisan advantage under DCP remains low across the full distribution of possible statewide vote shares. Crucially, these results obtain when evaluating every possible map at both the Define and Combine stages. Thus, in both our state-by-state simulations and in fully solvable grid maps, DCP reduces the advantage conferred to the redistricting party as well as partisan bias from redistricting.

Figure H.1 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.

In this case, DCP definitively reduces the extent to which a minority party that moves first can gerrymander their opponent (comparing the solid and dotted blue lines for Party A vote shares below 50%). DCP also reduces the extent of bias due to gerrymandering when the first mover holds a statewide majority, but for a narrower range of values.

When one party dominates statewide 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 $\delta^U_{.5} \equiv 0.8 - 0.2 = 0.6$, or three seats. Under DCP the first-mover advantage is $\delta^D_{.5} \equiv 0.622 - 0.374 = 0.248$, or about 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 almost 2 seats. DCP does offer the defining party a first-mover advantage in this context of an 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 H.2 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 almost equal to half of all seats on average ($\delta^U_{.5} = 0.41$). In contrast, under DCP the seat share remains nearly the same no matter which party goes first ($\delta^D_{.5} = 0$).

Figure H.3 plots the values of δ^U_V and δ^D_V against vote share for the 150 precinct grid. Partisan advantage due to unilateral redistricting (δ^U_V) peaks when the state is split evenly. In contrast, the partisan advantage under DCP (δ^D_V) is low across the full distribution of partisan vote share scenarios. 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 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 0.7 and above.

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 in our simulations 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 H.4 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 cases by a meaningful margin.

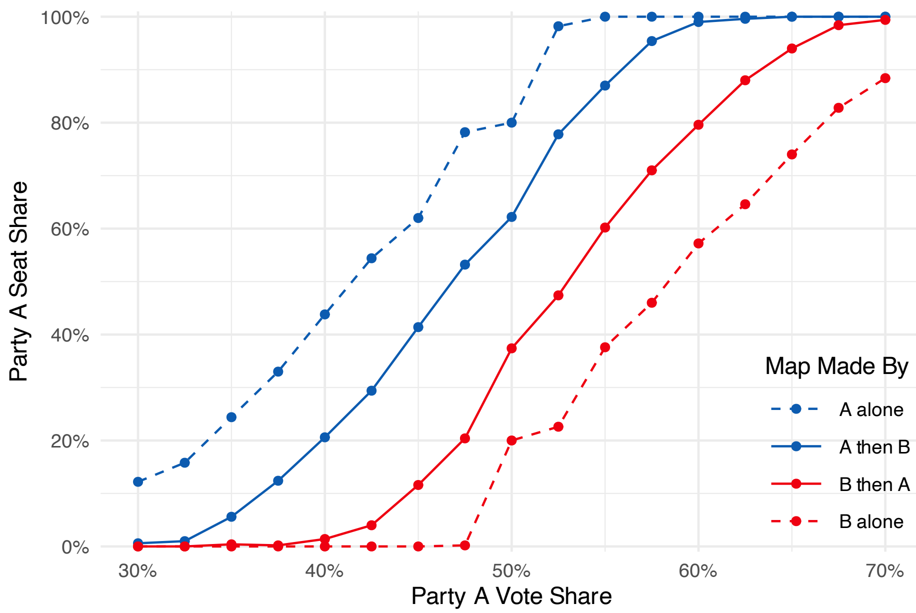

Figure H.1: Results for Voter Distribution Simulations on a Simple Grid. Each simulation uses a 5x6 grid.

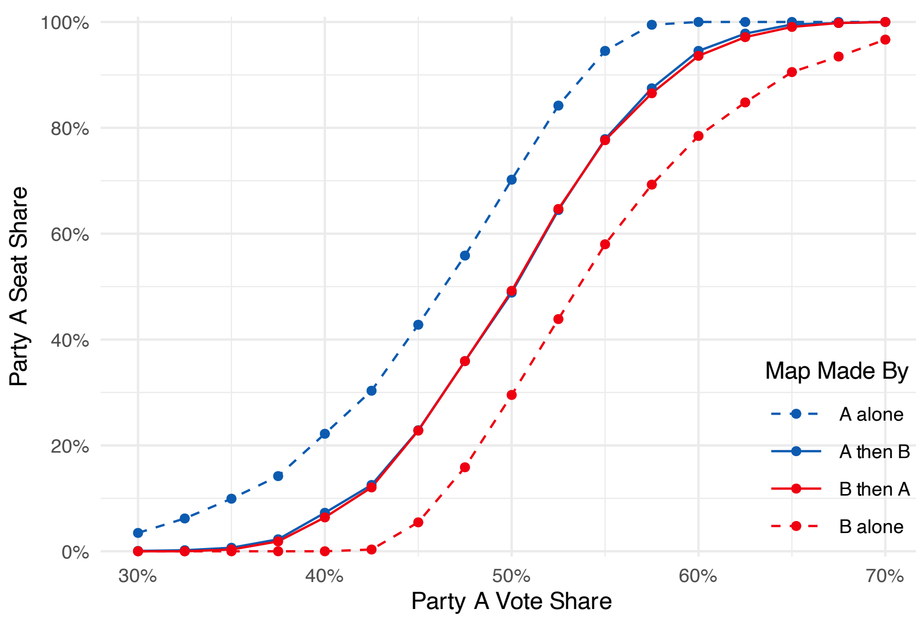

Figure H.2: Results for Voter Distribution Simulations on a Larger Grid (150 precincts; 15 districts)

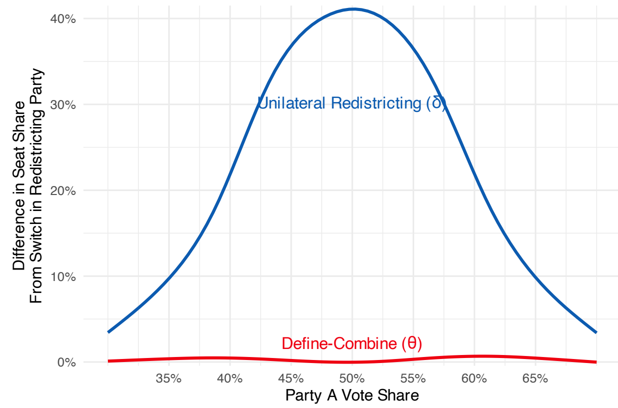

Figure H.3: Differences in Unilateral and DCP Advantage by Partisan Composition

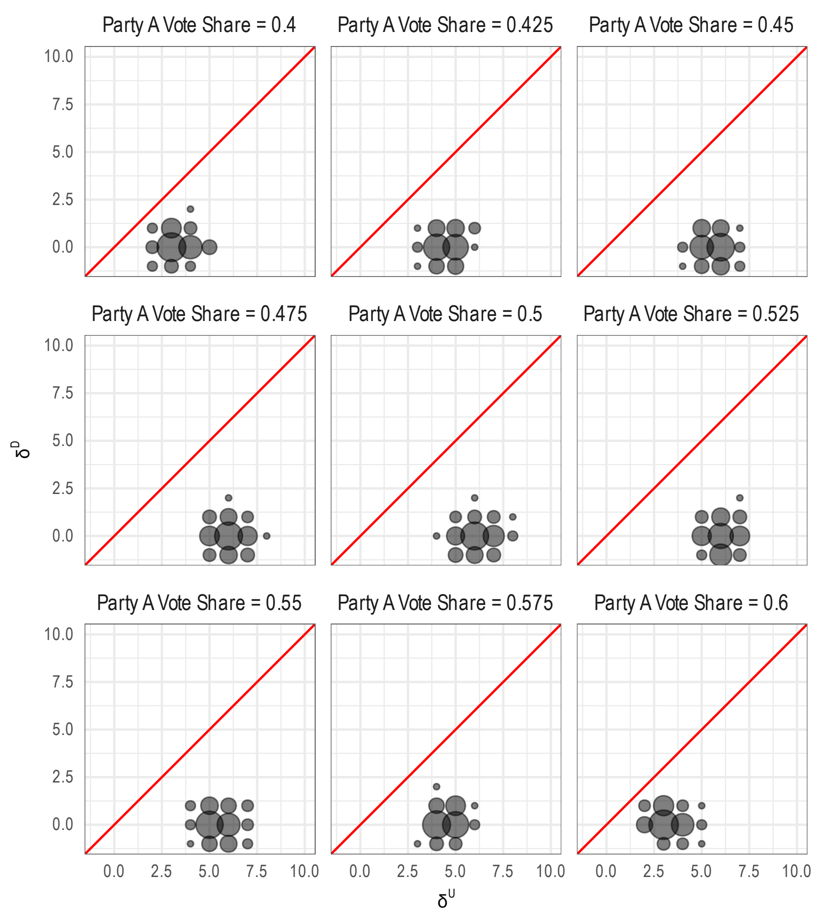

Figure H.4: Values of δ^U and δ^D for each simulated voter distribution on the 150 precinct grid.

H.3 Alternative Objectives to Maximizing Seats Won in Next Election

Because they face electoral uncertainty, parties may have objectives other than simply maximizing seats won in the next election. For example, parties that are "running scared" may seek to maximize partisan bias and minimize responsiveness to insulate against future partisan swings. Conversely, parties optimistic about the future may favor plans maximizing responsiveness and minimizing bias (Katz et al., 2020).

When parties hold differing objectives—one party maximizing responsiveness and the other bias—our same core results obtain. We plot these results in Figure H.5, which illustrates results when Party A maximizes responsiveness and Party B maximizes bias.[27] DCP reduces the advantage conferred to the unilateral redistricting party for relatively competitive statewide vote shares (e.g., between 0.45 and 0.55). The reductions in partisan bias from implementing DCP operate similarly. 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 statewide vote shares. However, 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.

These results are also informative for how parties with utility functions sensitive to electoral uncertainty may respond for both URP and DCP. As Figure H.5 illustrates, in a scenario where a party seeks more responsiveness and less bias (e.g., Party A), the party trades off winning a large majority of districts in order to do better if the vote swings in its favor. Specifically, this is reflected by the fact that Party A wins roughly 60% of seats for a 50-50 district under URP, leaving roughly 20% of seats on the table in order to draw more responsive districts. Similarly, when seeking more bias and less responsiveness, Party B trades off winning every possible seat in order to draw safer seats. Crucially, DCP still narrows the gap in seats between the parties as compared to URP even when parties have objectives that are not focused only on maximizing seats won in just the next election.

More broadly, this intuition should extend to cases where both parties have other, opposing objective functions. For example, if parties differ about which incumbents to protect, DCP is likely to reduce the extent of the incumbency advantage as well.

[27] Party 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 et al. (2020) finds a tradeoff between bias and wins, such that maximizing one should generally reduce the other as well.

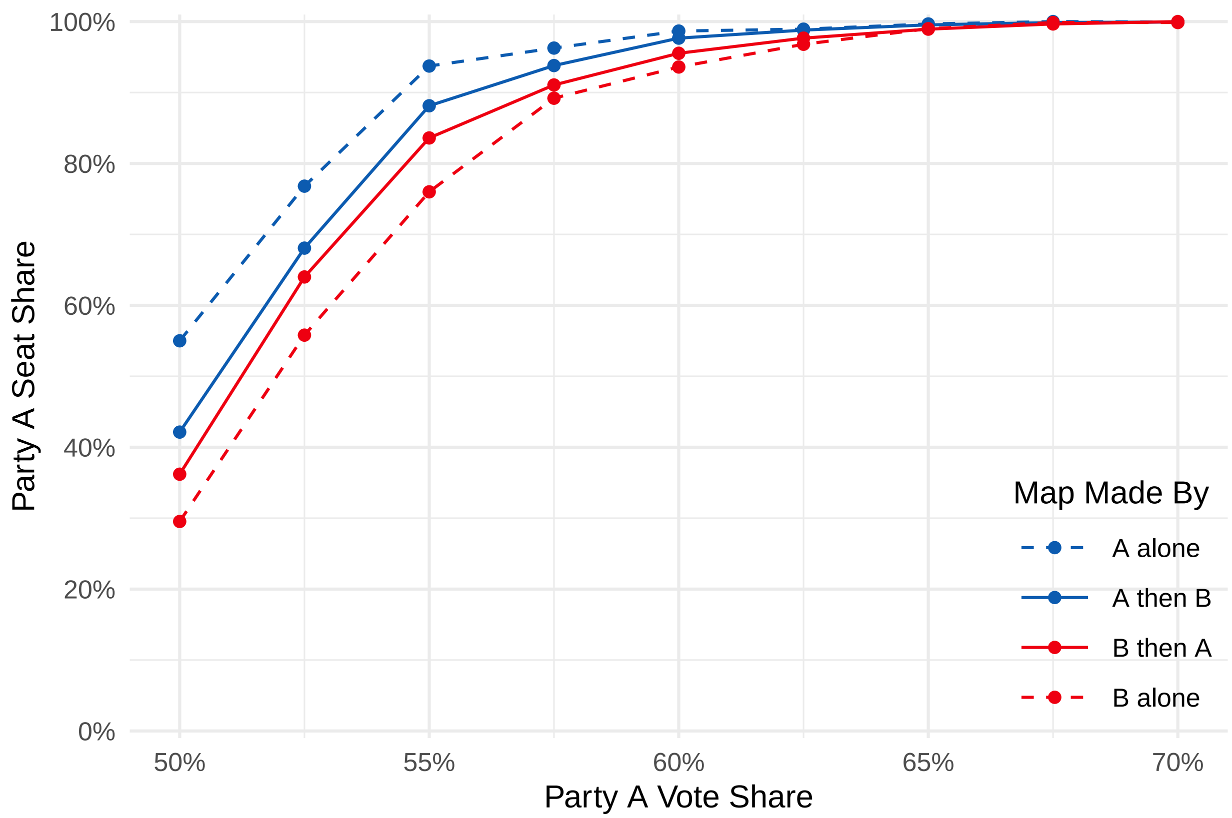

Figure H.5: Simulation Results with Alternative Objective Functions. Party A is maximizing responsiveness, while Party B is maximizing bias.

H.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 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 again use grid maps to isolate the effects of different geographic constraints on the redistricting process—ones that are too hard or impossible to model analytically. We examine two different elements of geographic constraints: (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).[28] 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 H.9). A Moran's I of 0 indicates randomly dispersed voters with neither clustering nor a pattern of "even" dispersion. Appendix H.5 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 H.6 plots the relationship between seats won by Party A and the amount of geographic clustering, when the vote is exactly split between the parties, and Figure H.7 plots the relationship between seats won by Party A and the amount of geographic clustering while varying the support for Party A between 40% and 60%.

As clustering increases, the definer's advantage ($\delta^U_{.5}$) for unilateral redistricting increases

[28]See 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).

slightly as well. Going from the lowest observed level of Moran's I to the highest increases $\delta^U_{.5}$ by a bit under 0.2. How does DCP perform under different levels of geographic partisan clustering? When clustering is low, DCP removes most of the definer/first-mover advantage. When clustering is higher, 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., $1 - \frac{\delta^D_{.5}}{\delta^U_{.5}} \geq 0.8$). Overall, then, map-wide clustering does not appear to meaningfully alter our general results for the definer's advantage.[29,30] 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.

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 H.8 plots the results. Across the draws, geographic bias (the x-axis) varies by about 6% 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.

Of particular note, these results also suggest that the average or mid-point of seats won for either URP or the first stage of DCP maps captures useful information about a map's geographic bias and could be used to assess such bias alongside existing measures (Chen and Rodden, 2013; Benadè et al., 2021; McCartan and Imai, 2023).[31]

[29]Figures H.7 illustrate the relationship between seats won and clustering for different values of statewide vote share V.

[30]In 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 H.5 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.

[31]This idea is also endorsed implicitly by our use of the URP midpoint as a benchmark or fairness target

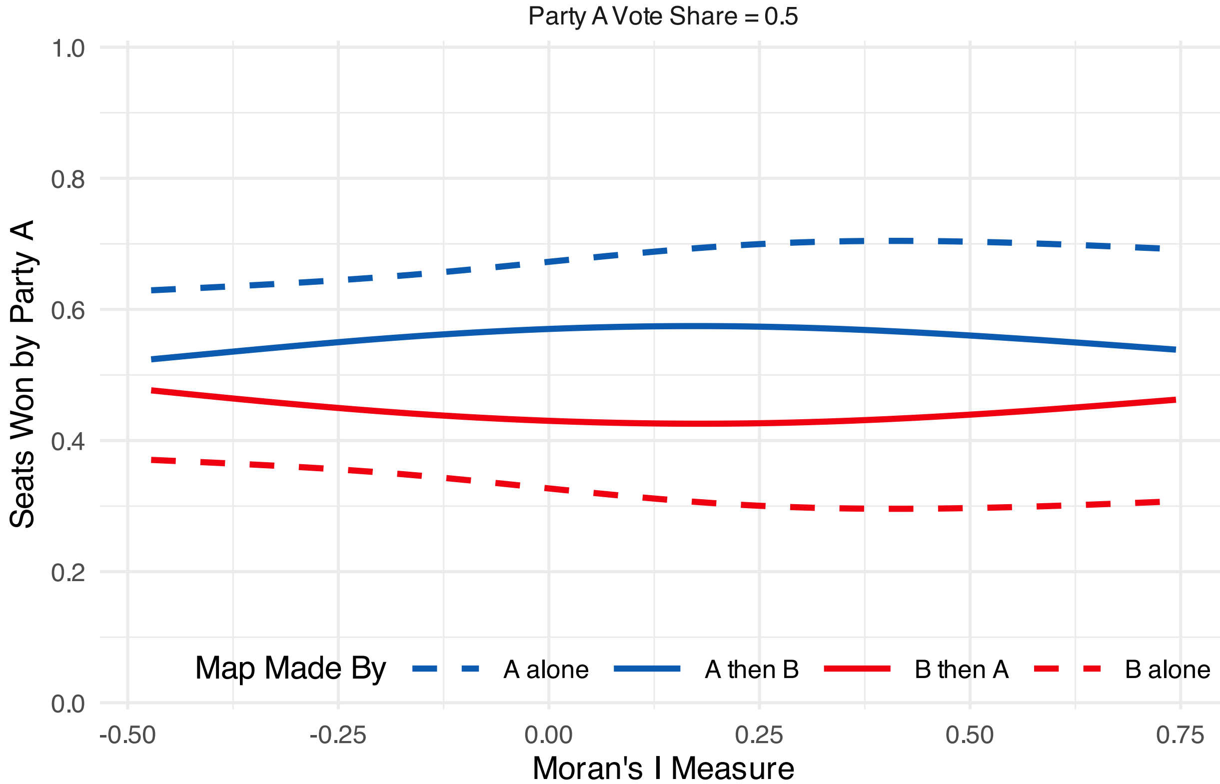

Figure H.6: Define-Combine Results, by Vote Share and Moran's I, Vote Share 50% for each Party.

in Section 4.

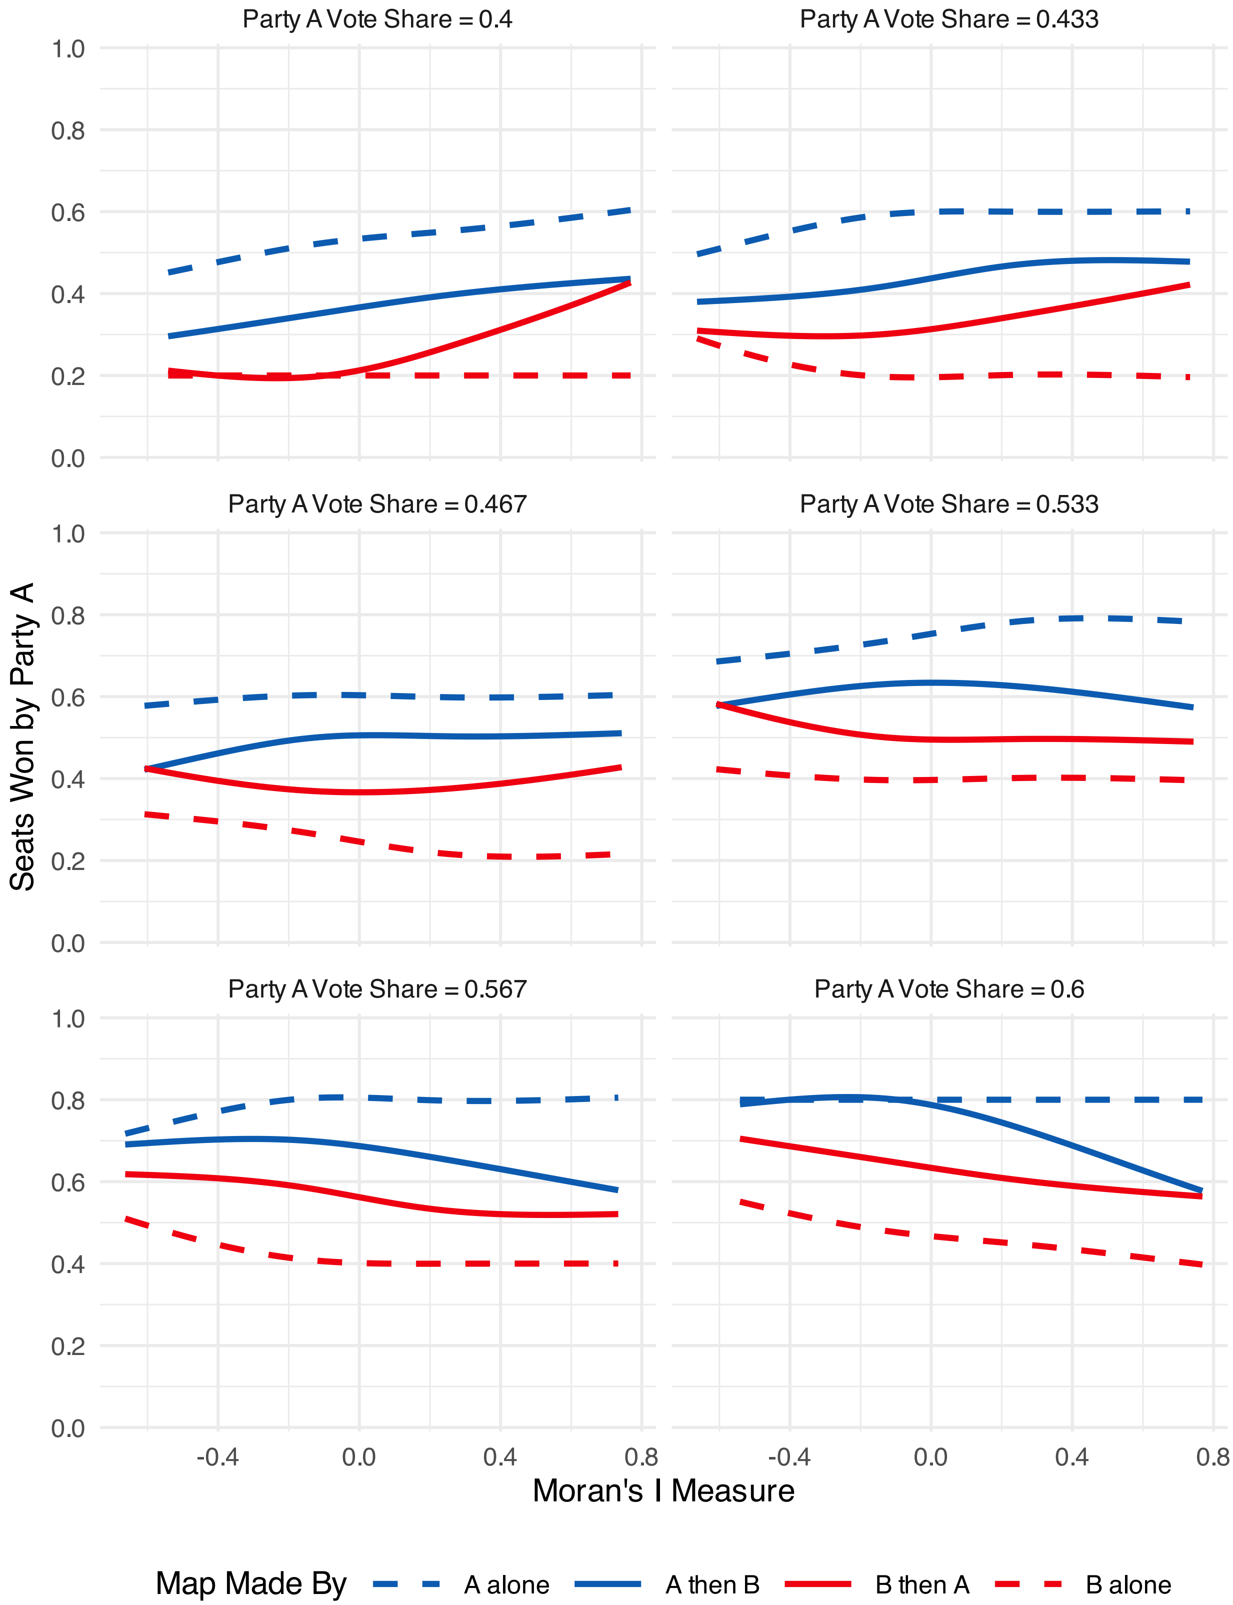

Figure H.7: Define-Combine Results, by Vote Share and Moran's I, 40%-60%

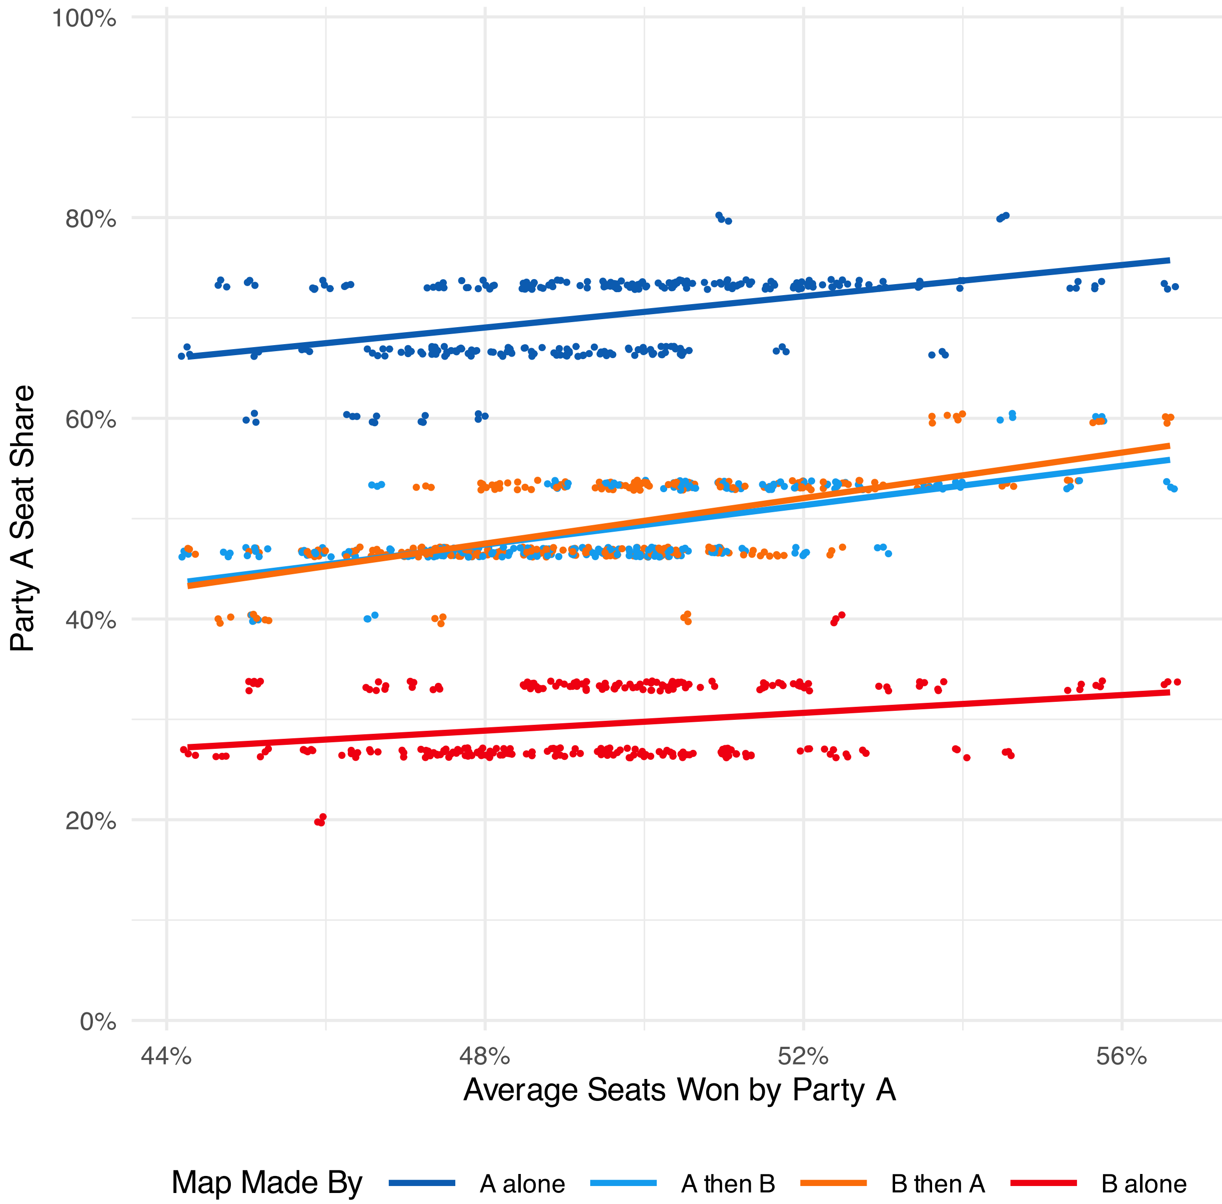

Figure H.8: Relationship Between Geographic Bias and Outcomes for Simulated Grids

H.5 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:

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

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 H.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 H.9 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 H.9 (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 H.9 (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 H.9 (c) shows a case where Moran's I is approximately 0, indicating neither clustering nor a tendency towards even dispersion. And Figure H.9 (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 H.10 (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 H.10 (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 H.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 H.9 (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.[32] This is not a significant limitation of our clustering analysis because a map with clustering close to -1 becomes 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 H.7 plots the relationship between seats won by Party A and Moran's I, but with each plot depicting results for a different overall Party A vote share. In general, the pattern is similar to the case where statewide vote share is split 50-50; in all cases, the magnitude of $\delta^{D}_{.5}$ is substantially less than $\delta^{U}_{.5}$, 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.

[32]There 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.

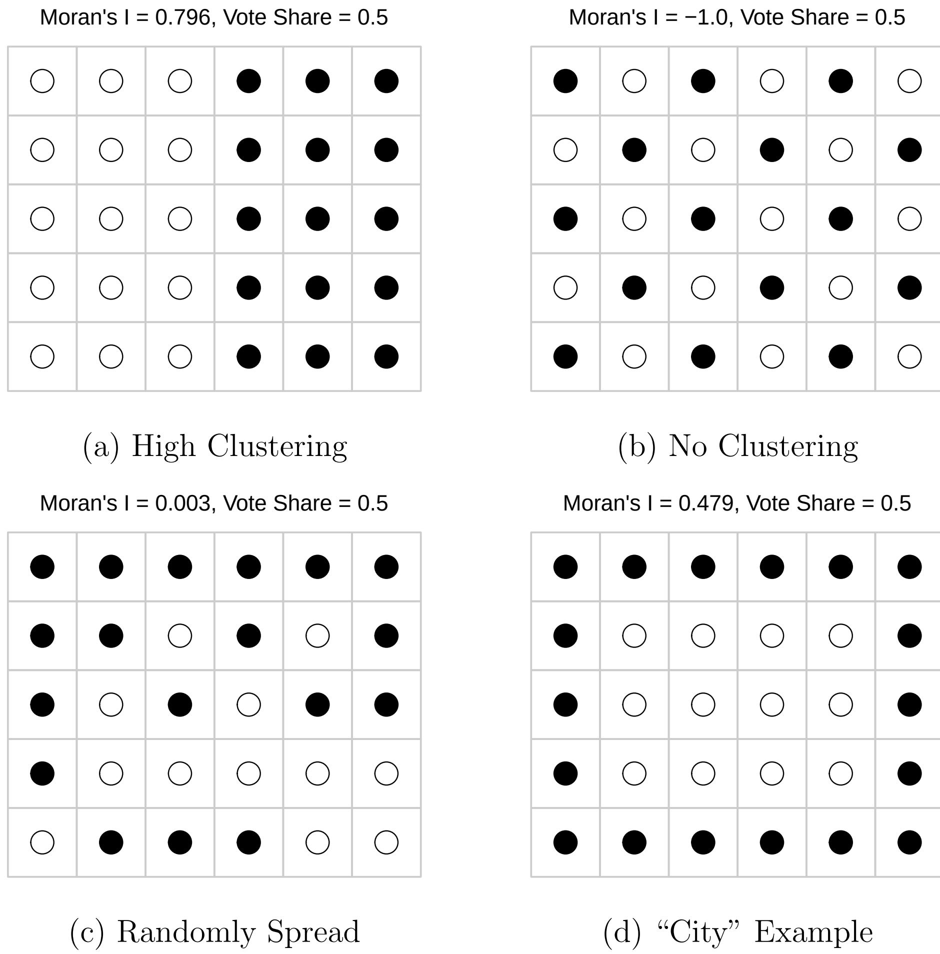

Figure H.9: Ranges of Moran’s I with Different Vote Distributions

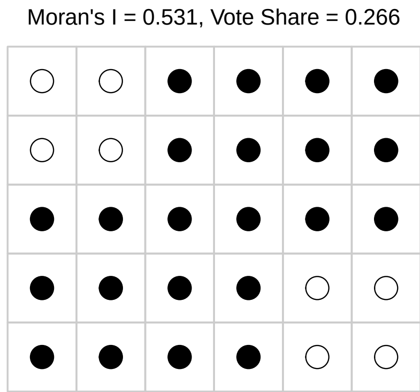

(a) High Moran's I, Low Vote Share

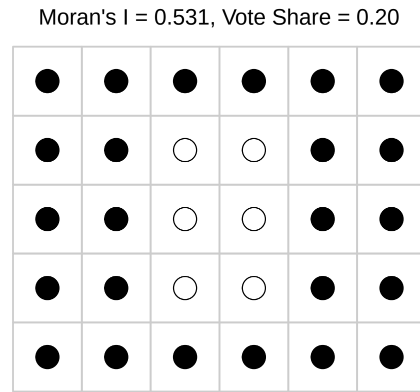

(b) High Moran's I, Lower Vote Share

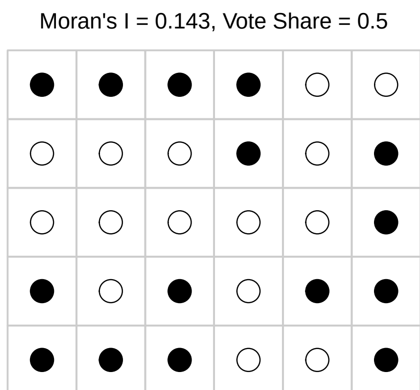

(c) Random - Positive Clustering

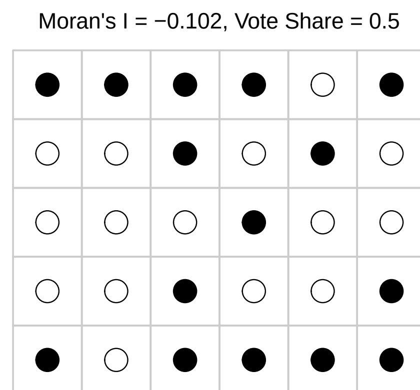

(d) Random - Negative Clustering

Figure H.10: Samples of Moran's I with Different Vote Distributions

I Calculations for Number of Possible Maps

This Appendix section provides calculations for the number of possible maps at each stage of the Define-Combine Procedure, assuming that there are no geographic or compactness constraints. We start with calculations for the second stage, which is more computationally feasible, then discuss the first stage.

I.1 Second Stage

At the second stage, for a state with N Congressional districts there will be $2N$ subdistricts that have been defined in the first stage by the definer. The number of possible final combine stage maps pairing subdistricts is smaller than the number of first stage maps using subdistrict components, such as precincts or census blocks. However, it is still so large as to make exhaustive search algorithms infeasible in many cases.

For any single define-stage map, there will be $2N - 1$ choose 2 ways to combine the subdistricts to make the first district:

$$\binom{2N-1}{2} = \frac{(2N-1)!}{(2N-3)!2!}$$

For the second final-stage districts, there are $2N - 1 - 2$ remaining subdistricts of which 2 need to be chosen:

$$\binom{2N-3}{2} = \frac{(2N-3)!}{(2N-5)!2!}$$

We repeat this process until $N - 1$ districts have been created,[33] amounting to:

$$\binom{2N-1}{2} \times \binom{2N-3}{2} \times \binom{2N-5}{2} \times ... \times \binom{N-1}{2} \times ... \times \binom{3}{2}$$

But there are also $(N - 1)!$ ways to choose this same set of districts, since order does not matter. So divide all of the potential combinations by $(N - 1)!$:

$$\frac{\binom{2N-1}{2} \times \binom{2N-3}{2} \times \binom{2N-5}{2} \times ... \times \binom{N-1}{2} \times ... \times \binom{3}{2}}{(N-1)!}$$

Writing as factorials yields:

$$\frac{\frac{(2N-1)!}{(2N-3)!2!} \times \frac{(2N-3)!}{(2N-5)!2!} \times \frac{(2N-5)!}{(2N-7)!2!} \times ... \times \frac{(N-1)!}{(N-3)!2!} \times ... \times \frac{3!}{1!2!}}{(N-1)!}$$

[33] Once $N - 1$ districts have been defined, the last district is also defined by the remaining subdistricts.

Note that:

$$\frac{(2N-1)!}{(2N-3)!2!} = \frac{(2N-1)(2N-2)(2N-3)!}{(2N-3)!2!} = \frac{(2N-1)(2N-2)}{2!}$$

After substituting this expression in at each instance, it yields:

$$\frac{\frac{(2N-1)(2N-2)}{2!} \times \frac{(2N-3)(2N-4)}{2!} \times \frac{(2N-5)(2N-6)}{2!} \times ... \times \frac{(N-1)(N-2)}{2!} \times ... \times \frac{3!}{1!2!}}{(N-1)!}$$

$$\rightarrow$$

$$\frac{\frac{(2N-1)!}{2!}}{(N-1)!} = \frac{(2N-1)!}{(N-1)!2!}$$

Table I.1 calculates the total number of combine stage maps for a given (1) define stage map, depending on the number of final districts (N) that need to be delineated. As demonstrated by the table, the number of combine-stage maps for even just a single define stage map very quickly becomes computationally unfeasible to enumerate, at least without geographic and compactness and other constraints. Even a relatively small state like Kentucky (6 districts) will have over 166,000 possible combine-stage maps for *each* define stage map.

Table I.1: Possible Combine Stage Maps Given A Single Define Stage Map

N	$2N$	Total Combinations
1	2	1
2	4	3
3	6	30
4	8	420
5	10	7,560
6	12	166,320
7	14	4,324,320
8	16	129,729,600
9	18	4,410,806,400
10	20	167,610,643,200
11	22	7,039,647,014,400
12	24	323,823,762,662,400
13	26	16,191,188,133,120,000
14	28	874,324,159,188,480,000
15	30	50,710,801,232,931,800,000
16	32	3,144,069,676,441,770,000,000
17	34	207,508,598,645,157,000,000,000
18	36	29,051,203,810,322,000,000,000,000
27	54	5,299,992,308,438,690,000,000,000,000,000,000,000,000,000
36	72	41,152,928,000,957,700,000,000,000,000,000,000,000,000,000,000,000,000,000,000
53	106	Approximately 6.7×10^{100}

Notes: N indicates the number of Congressional districts; $2N$ is the number of subdistricts that would be created in the "define" stage of DCP; and the "Total Combinations" column displays the number of possible final maps that could be created by the "combiner" in the second stage for a single define-stage map.

I.2 First Stage

Without geographic constraints, the number of possible first stage districts isn't really even feasible to fit on a page. The number of possible district maps can is a *Stirling number of the second kind* (Liu et al., 2016), which is the number of ways to divide N parts into K sets (or districts, in this case):

$$S(N, K) = \sum_{i=1}^{K} \frac{(-1)^{K-i} i^{N-1}}{(i-1)!(K-i)!}$$

These end up being extremely large in the context of redistricting, using precincts or census blocks as parts that need to be allocated into districts, and in general the number of possible maps makes it computationally infeasible to generate or "score" every possible map (Liu et al., 2016). Assuming constraints such as equal population will constrain the total number of possible valid maps significantly, but it will still be orders of magnitude too large to exhaustively search computationally.

As an example, take a medium sized state, such as Alabama, which has only 7 Congressional districts made up of 2,111 precincts.[34] Rounding to $2,114$ precincts to make it a number divisible by 14, and assuming that all precincts are of equal population, there are $\binom{2{,}113}{151}$ ways to construct the first district, $\binom{1{,}962}{151}$ ways to construct the second district, and so on, divided by $(14-1)!$. Using the same formula derived for the combine stage maps above, this amounts to $\frac{(P-1)!}{(N-1)!151!}$, where P is the number of precincts. The resulting total is well above 10^{5400} possible maps, and for *each* of these maps there are 4,324,320 possible combine-stage maps (see Table I.1).

At both the define and combine stages, the total number of possible map combinations makes it impossible to do an exhaustive search unless significant constraints are placed on the set of allowed maps.

J Comparing Analytical Results, Grid Maps, and Real-World Simulations

The analytical model without geographic constraints does not capture all the ways in which DCP may improve upon URP precisely because it does not take geography into account. By combining the insights from the analytical model and the fully solved grid simulations, we can illustrate how DCP interacts with *geographic constraints* to constrain gerrymandering further.

[34] Source: Election Administration and Voting Survey 2020 Comprehensive Report

Specifically, comparing the grid simulations to the analytical model highlights a puzzle: without geographic constraints, the analytical model suggests the largest advantage conferred to the first mover (e.g., Definer's Advantage) occurs in states with the most competitive parties statewide (e.g., a 50-50 state). However, in the simulations, this result does not carry through. Why? In a world with geographic constraints, DCP requires the Definer to create smaller subdistricts than the final districts (e.g., half the size in population); as a result, it constrains how much parties may combine disparate populations to gerrymander. Without geographic constraints, the parties can better take advantage of this cracking approach in states where parties are evenly split; however, with even mild geographic constraints, this is no longer true.

To study this trade-off further, we can examine grid maps with variation in the degree of clustering and one party with a clear majority.[35] Figure H.6 depicts seat shares by party across varying levels of partisan clustering (a negative value of Moran's I suggests more evenly mixed parties and a positive value of Moran's I suggests more clustered parties). No partisan clustering amounts to a case where the majority party Definer can implement the optimal strategy from the analytical model, drawing districts with a narrow party majority and winning all districts. As the level of clustering increases, the geographic constraints begin to bind. Specifically, the majority (Definer) party can no longer implement the optimal strategy from the analytical model: geographic clustering combined with the smaller districts required by DCP prevents the majority from drawing as many subdistricts where the party wins narrowly. The plot of seat share against geographic clustering in Figure H.7 accords with this story. For the lowest levels of geographic clustering, Party A, which has a 60% majority of voters statewide, can win a seat share when acting as the Definer equivalent to the seat share from unilateral redistricting. However, as geographic clustering increases, the seat share for the Definer declines as does the bias induced by redistricting. In contrast, the seat share for a unilateral redistricter with a statewide majority (e.g., Party A) remains the same regardless of the level of clustering.

This insight, based upon the analytical results in Appendix B and the fully solvable grid maps, matches well with the simulation results from real-world redistricting problems. The median state has a Moran's I of 0.8, and even the state with the lowest Moran's I in our sample, New Hampshire, still has a value of 0.37. As a result, U.S. states have sufficient partisan clustering for geographic constraints to interact with DCP to constrain the extent of gerrymandering not just for a minority party that finds itself as the initial map drawer, but also for a majority party.

[35] Appendix Figure H.9 depicts examples of different levels of clustering in a grid map.

References

Alexeev, B. and D. Mixon (2019). Partisan gerrymandering with geographically compact districts. *Journal of Applied Probability 55*(4), 1046–1059.

Altman, M. and M. P. McDonald (2011). Bard: Better automated redistricting. *Journal of Statistical Software 42*(4).

Benadè, G., A. Procaccia, and J. Tucker-Foltz (2021). You can have your cake and redistrict it too. *ACM Transactions on Economics and Computation*.

Bernstein, M. and M. Duchin (2017). A formula goes to court: Partisan gerrymandering and the efficiency gap. *Notices of the AMS 64*(9), 1020–1024.

Brams, S. (2020). Making partisan gerrymandering fair: One old and two new methods. *Social Science Quarterly 101*(1), 68–72.

Brams, S. and A. Taylor (1996). *Fair Division: From cake-cutting to dispute resolution.* Cambridge University Press.

Brennan Center for Justice (2018, July). Redistricting Commissions: What Works. Technical report, Brennan Center for Justice.

Carson, J. L. and M. 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, J. and D. 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, J. and J. Rodden (2013). Unintentional gerrymandering: Political geography and electoral bias in legislatures. *Quarterly Journal of Political Science 8*, 239–269.

Cho, W. T. and J. G. Gimpel (2012). Geographic information systems and the spatial dimensions of american politics. *Annual Review of Political Science 15*, 443–460.

Cho, W. T. and Y. Liu (2016). Toward a talismanic redistricting tool: A computational method for identifying extreme redistricting plans. *Election Law Journal 15*(4), 351–366.

Cirincione, C., T. A. Darling, and T. G. O'Rourke (2000). Assessing south carolina's 1990s congressional districting. *Political Geography 19*(2), 189–211.

Duchin, M. (2018). Gerrymandering metrics: How to measure? what's the baseline? https://arxiv.org/abs/1801.02064.

Edwards, B. C., M. H. Crespin, R. D. Williamson, and M. Palmer (2017). Institutional control of redistricting and the geography of representation. *Journal of Politics 79*(2), 722–726.

Ely, J. (2019). A Cake-Cutting Solution to Gerrymandering. pp. 27. https://www.jeffely.com/post/a-cake-cutting-solution-to-gerrymandering.

Fifield, B., M. Higgins, K. Imai, and A. Tarr (2020). A new automated redistricting simulator using markov chain monte carlo. *Journal of Computational and Graphical Statistics 0*(0), 1–14.

Grofman, B. and G. 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, J. A., B. T. Hamel, and A. M. Goldzimer (2018). Gerrymandering incumbency: Does nonpartisan redistricting increase electoral competition? *The Journal of Politics 80*(3), 1011–1016.

Katz, J., G. King, and E. Rosenblatt (2020). Theoretical foundations and empirical evaluations of partisan fairness in district-based democracies. *American Political Science Review 114*(1), 1–15.

King, G. and R. 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, March). A fair division solution to the problem of redistricting. *Social Choice and Welfare 32*(3), 479–492.

Levitt, J. (2020). All about redistricting: Who draws the lines? https://redistricting.lls.edu/who.php. Accessed 2/16/20.

Liu, Y. Y., W. K. T. Cho, and S. Wang (2016). Pear: a massively parallel evolutionary computation approach for political redistricting optimization and analysis. *Swarm and Evolutionary Computation 30*, 78–92.

Magleby, D. and D. Mosesson (2018). A new approach for developing neutral redistricting plans. *Political Analysis 26*(2), 147–167.

Mayer, K. (2016). Rebuttal report: Response to expert reports of sean trende and nicholas goedert.

McCartan, C. and K. Imai (2023). Sequential monte carlo for sampling balanced and compact redistricting plans. *Annals of Applied Statistics*.

McDonald, M. P. (2004). A comparative analysis of redistricting institutions in the united states, 2001-02. *State Politics and Policy Quarterly 4*(4), 371–395.

McDonald, M. P. and R. E. Best (2015). Unfair partisan gerrymanders in politics and law: A diagnostic applied to six cases. *Election Law Journal 14*(4), 312–330.

Miller, P. and B. Grofman (2013). Redistricting commissions in the western united states. *UC Irvine Law Review 3*(3), 637–668.

Moran, P. A. (1948). The interpretation of statistical maps. *Journal of the Royal Statistical Society Series B*(10), 243–251.

Pegden, W., A. Procaccia, and D. Yu (2017). A partisan districting protocol with provably nonpartisan outcomes. `https://arxiv.org/abs/1710.08781`.

Schaffner, B. and S. Ansolabehere (2015). 2010-2014 Cooperative Congressional Election Study Panel Survey. `https://doi.org/10.7910/DVN/TOE8I1`, Harvard Dataverse, V11.

Stephanopoulos, N. (2018). The causes and consequences of gerrymandering. *William and Mary Law Review 59*(5).

Stephanopoulos, N. and E. McGhee (2015). Partisan gerrymandering and the efficiency gap. *University of Chicago Law Review 82*, 831–900.

Stephanopoulos, N. and E. McGhee (2018). The measure of a metric: The debate over quantifying partisan gerrymandering. *Stanford Law Review 70*(5), 1503–1566.