

On Measuring Two-Party Partisan Bias in Unbalanced States

John F. Nagle and Alec Ramsay

ABSTRACT

Assuming that partisan fairness and responsiveness are important aspects of redistricting, it is important to measure them. Many measures of partisan bias are satisfactory for states that are balanced with roughly equal proportions of voters for the two major parties. It has been less clear which metrics measure fairness robustly when the proportion of the vote is unbalanced. We have addressed this by analyzing past election results for four states with Democratic preferences (CA, IL, MA, and MD), three states with Republican preferences (SC, TN, and TX) and comparing those to results for four nearly balanced states (CO, NC, OH, and PA). We used many past statewide elections in each state to build statistically precise seats for votes and rank for votes graphs to which many measures of partisan bias were applied. In addition to providing values of responsiveness, we find that five of the measures of bias provide mutually consistent values in all states, thereby providing a core of usable measures for unbalanced states. Although all five measures focus on different aspects of partisan bias, normalization of the values across the eleven states provides a suitable way to compare them, and we propose that their average provides a superior measure which we call composite bias. Regarding other measures, we find that the most seemingly plausible symmetry measure fails for unbalanced states. We also consider deviations from the proportionality ideal, but using it is difficult because the political geography of a state can entangle responsiveness with total partisan bias. We do not attempt to separate intentional partisan bias from the implicit bias that results from the interaction of the map drawing rules of a state and its political geography, on the grounds that redistricting should attempt to minimize total partisan bias whatever its provenance.

Keywords: Area: United States; redistricting, partisan bias, fairness, responsiveness, political geography

1. INTRODUCTION

John F. Nagle is a professor emeritus of physics and biological sciences at Carnegie Mellon University in Pittsburgh, Pennsylvania, USA. Alec Ramsay is principal at Dave's Redistricting in Seattle, Washington, USA. We thank B. Grofman, J. Cervas, M.D. McDonald, G. King, G.R.J. Wallin, S.-H. Wang, M. Duchin, N. Stephanopoulos, E. McGhee, and G.S. Warrington for informative discussion. We especially thank Stephen Wolf, whose data made this study possible.

© John F. Nagle and Alec Ramsay, 2020; Published by Mary Ann Liebert, Inc. This Open Access article is distributed under the terms of the Creative Commons Attribution Noncommercial License (<http://creativecommons.org/licenses/by-nc/4.0/>) which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and the source are cited.

IT IS WELL RECOGNIZED, not only by political scientists and politicians, that redistricting Congress and state legislatures is important. Indeed, ordinary citizens have engaged in drawing maps using free public software.¹ Not surprisingly, there are many

¹Some examples are Dave's Redistricting (Bradlee 2020), Districtr (Metric Geometry and Gerrymandering Group 2018), and JudgeIt (Gelman et al. 2012). In Pennsylvania, the Committee of 70 has provided Draw the Lines software (Thornburgh 2018) for people to draw maps to compete for substantial prizes; it unfortunately only uses registration at the overly coarse census tract level instead of election results at the precinct level.

criteria that can be considered in judging maps and different people place different weights on the different criteria. This article focuses on two of those criteria that we think are especially important, namely, fairness and responsiveness.²

Our measures of bias are for total partisan bias, of which overtly intended gerrymandering is just one part. There is also underlying bias due to political geography which is sometimes called unintentional gerrymandering (Chen and Rodden 2013). That bias is being quantified using ensembles of maps drawn by computers that use the criteria that a state prescribes in its rules.³ That provides a benchmark that is then subtracted from the bias of the adopted plan to estimate intended gerrymandering bias. While this is what is traditionally required for court cases (Grofman 2019; McDonald et al. 2018), we emphasize that unintentional bias is no longer unintentional if it can be reliably demonstrated that it occurs when following the state's rules or other informal criteria. This bias then becomes systemic bias due to the state's map drawing criteria interacting with its political geography. If there is such systemic bias, there may still be plans that are outliers in the ensemble that are nonetheless fair, and we believe a redistricting commission should adopt such a plan rather than an average plan. However, if no relatively fair plan is possible within the state's rules, then we believe the state should change its election laws to make fair plans possible. Even if a state's rules cannot be changed in time, a redistricting commission and concerned map-drawing citizens could still try to minimize total bias and achieve responsiveness.⁴ As has been reviewed by Stephanopoulos (2013), numerous states have had language promoting fairness and responsiveness, but these efforts have been hampered by uncertainties in defining and measuring these concepts, and that is what this article is about.

This poses the question: Is it possible to test whether a proposed redistricting plan for 2021 will be fair and responsive to the voters? Our approach to this question is to analyze the enacted redistricting plans for 2011 using many statewide election results applied to those maps.⁵ This article presents evidence that the use of past election results can indeed produce quite precise information that can be used for evaluating a plan before it is implemented.⁶ We suggest that this exercise can be relevant for the next round

of redistricting after the 2020 census, at least in those states that have redistricting commissions not intent on political advantage.

Even in states controlled by one party, the findings in this article may help challenge unfair and unresponsive plans in court. However, we again emphasize that our work is not focused on challenging plans in courts that, in lieu of precise laws regarding fairness and responsiveness, typically require evidence of partisan intent to overturn a map.⁷ Political geography, such as high density of Democrats in cities, can also create unfair and unresponsive maps when conventional redistricting criteria such as compactness and not splitting political subdivisions are adhered to (Rodden 2019; Nagle 2019). Our goal in this article is to elucidate measures of total bias and unresponsiveness, whatever their provenance.⁸ Map drawers can then choose how to balance these metrics with other criteria⁹ and a state legislature can be better informed about whether its state election law should be modified.

For states that are evenly balanced between two dominant parties, a fair plan is clearly one

²Responsiveness is often called competitiveness, but while closely related at the statewide level we prefer responsiveness to voters over competitiveness for parties.

³Of course, implementation of the rules in the computer code generally requires weighting of the different criteria, population equality, compactness, splitting political subunits, communities of interest, and minority protection. Remarks made by Katz et al. (2020, 176) and McGann et al. (2016, 110) regarding the use of computer drawn map ensembles to evaluate bias are especially pertinent.

⁴This echoes the discussion of McGann et al. (2016, 222) that "Electoral fairness is not something that occurs 'naturally'; it has to be actively pursued," and Altman and McDonald (2018, 108) that a solution may be to explicitly incorporate political goals into the redistricting criteria.

⁵Statewide elections are the same for all precincts whereas results for Congress or state legislatures are subject to incumbency differences as well as uncontested elections.

⁶Although this conclusion will not come as a surprise to many political scientists (e.g., McGann et al. 2016; McDonald et al. 2018; Gelman et al. 2012; Grofman and King 2007; Wang et al. 2018; Stephanopoulos and McGhee 2018), the methods and evidence presented here may still be of interest, especially for the express focus on unbalanced states.

⁷See Best et al. (2020) for insightful concerns regarding reliance on a partisan intent standard.

⁸Following the authors in the previous footnote 6. In particular, "the absence of intentional unfairness is not the same as fairness" (Katz et al. 2020, 170).

⁹Criteria balancing has recently been discussed by Altman and McDonald (2018).

that is likely to result in half the seats for half the votes.¹⁰ It has been rather more challenging to decide what a fair plan is for a state in which one party routinely obtains considerably more than half the votes (McGhee 2014; Wang 2016; Nagle 2017). It has appeared reasonable to assert that symmetry is required in the sense that if party A wins fraction S_A of the seats with fraction V of the vote, then party B should win the same fraction $S_B = S_A$ of the seats if it received the same vote fraction (e.g., McGann et al. 2016, 56). A difference in seats $\Delta S = S_A - S_B$ between these two seat fractions would then appear to be a reasonable measure of bias at vote V , with $\Delta S = 0$ being no bias. This has been recently named the β measure of bias (Katz et al. 2020, 166). It obtains the value $\Delta S(\langle V \rangle)$, where the $\langle V \rangle$ means that we calculate ΔS from the most likely statewide vote $\langle V \rangle$ determined by the average over many statewide contests and from its counterfactual counterpart when the vote fraction is reversed to $1 - \langle V \rangle$.¹¹

In order to calculate $\Delta S(\langle V \rangle)$, one must obtain a seats-votes curve $S(V)$ which estimates the fraction of seats S for any statewide vote fraction V (Katz et al. 2020, 165). We have calculated $S(V)$ curves for the congressional plans of 11 states. By using many past election results, our $S(V)$ curves have quite small estimated uncertainties. Details of our methodology are presented in Supplementary Appendix A.

The seats-votes curves are quite interesting, and they are appropriate for evaluating the responsiveness of a plan. However, section 3.2 of this article shows that the β measure of bias is highly misleading for states with a dominant party. This came as a shock to us, as symmetry and this way of evaluating it would appear to be so fundamental to what should be considered fair. We will show how and why the β measure fails by examining the vote shares v_j for congressional districts j in the form of a different kind of graph that we call rank-vote $r(v)$ graphs. The $r(v)$ graphs are also quite precise because we again use many election results. They are presented in sections 2–4.

The $r(v)$ graphs together with the $S(V)$ graphs enable us better to portray and understand bias. Nevertheless, as has been well known for a long time, there is still ambiguity in the single member district system as to what is fair in states that have a dominant political party. In such unbalanced states, bias and responsiveness become intertwined (McGann et al. 2016, 67). For example, if the two-party polit-

ical balance in a state is 60/40, then a plan that has each district with a 60/40 preference would likely elect all its representatives from the same party. That would be fair under the aforementioned β symmetry principle.¹² It may also be the only possibility if the political geography of the state is completely homogeneous.¹³ However, if the state's political geography were completely heterogeneous, one could draw a map that guarantees a 60/40 split in the seats,¹⁴ but one could also combine the precincts to give all seats to the dominant party or to any split between these extremes.

Whatever one deems fair in the preceding examples, they illustrate that the political geography of a state could clearly be important for assessing bias and responsiveness in unbalanced states (Rodden 2019). We examine this further in Supplementary Appendix C where we recall the ideal of proportionality in section C.1. This leads us in section C.2 to a method to measure bias in unbalanced states that takes political geography into account. However, that approach is difficult to apply to any single plan, so the main line of analysis in this article proceeds differently in section 5. Global symmetry is defined in subsection 5.1. Subsection 5.2 defines a completely new measure (γ) of bias, followed by subsections elaborating on the declination (δ)

¹⁰Technically, we estimate the voter preferences of the districts in a plan. These should be balanced such that half the seats would be won if equally attractive candidates are nominated by the parties when the overall statewide vote is split evenly. Of course, these conditions are unlikely to be met in any real district which is why actual elections are important, so our methods are therefore not designed to predict outcomes for individual districts. Nevertheless, for states with many districts, such influences tend to average out, so it becomes possible to predict overall outcomes more reliably. Even so, that is not the goal. Rather, as mentioned by Cervas and Grofman (2020, 7), it is to test the fairness and responsiveness of a plan *ab initio*, before specific candidates are chosen and before people and parties decide how to allocate resources to the various contests. It is also certainly not the goal to predict the vote fraction in a future election, only to estimate what the overall preference would be given the overall vote fraction V .

¹¹Katz et al. (2020) define $\beta(V)$ for all V and prefer to evaluate at $V = \langle V \rangle$, the most probable statewide vote or a range including $\langle V \rangle$. For succinctness, we will call this the β measure.

¹²One would have 40/60 preferences in each district if 1/3 of the dominant party's voters switch and then the other party would win all the seats and therefore $\beta = 0$.

¹³If every precinct has the same 60/40 balance, all districts must have the same 60/40 balance.

¹⁴This is most easily seen by defining complete heterogeneity as 60% of the precincts having only dominant party voters and creating districts only from like-minded precincts.

measure (Warrington 2018) and the lopsided outcomes (LO) measure (Wang 2016, 1263), and concluding in section 5.5 by considering maximal responsiveness as the primary goal. Comparisons of these and more standard measures of bias are given in section 6 which crucially demonstrates that a core subset consisting of five of these measures of bias is substantially consistent across the 11 states, including unbalanced states, and we propose that suitably averaging their values into what we call composite bias Ω provides a superior estimate of bias. Subsection 6.3 also tests measures of bias for durability as the statewide vote share changes.

A general discussion ensues in section 7. Our main result that several measures of bias agree for unbalanced states as well as for balanced states leads us to conclude that total bias can be reliably measured in all states even though the political geography varies by state.

2. ANALYSIS OF FOUR POLITICALLY BALANCED STATES

This section also describes some of our terminology.

2.1. Example of Colorado

We begin with the state of Colorado (CO). The average statewide vote $\langle V \rangle$ for 12 elections in the time period 2004–2012 was 50.6% Democratic, making this a politically balanced state in that time frame. Figure 1 shows the average two-party Republican vote fraction v_j for each congressional district (CD $_j$, $j = 1-7$) along with the standard error of the mean over 12 elections.¹⁵ The districts are rank ordered according to their placement j in a list that is ordered by the vote. We call this kind of graph the rank $r(v)$ graph because it plots the district rank versus the votes v_j in the districts.¹⁶ According to the $r(v)$ graph in Figure 1, CD1 and CD2 were strongly Democratic and CD4 and CD5 were strongly Republican in the 2011 CO map, whereas CD3, CD6, and CD7 were more competitive.¹⁷

An $r(v)$ graph displays the partisan preferences of each district in a state. It is also often used to count the partisan number of seats in the following too simple way. In Figure 1, one might say that preferences¹⁸ in CO are for three Democratic seats, in districts 1, 2 and 7, and four Republican seats, in

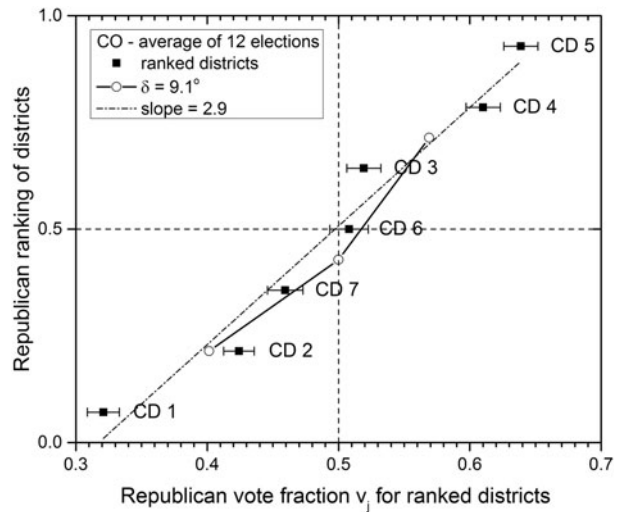


FIG. 1. Colorado $r(v)$ graph for the 2011 plan shows Republican district rank vs. Republican district vote. The slope of the linear fitted line to the districts is 2.9. The horizontal uncertainty bars are standard errors of the mean. The declination is the difference in the angles of the two solid lines (*vide infra*).

districts 3–6. However, that “all-or-nothing” assignment of the integers 1 or 0 to each district doesn’t take into account uncertainties in competitive districts like CD6. Instead of “integer assignment,”

¹⁵Following McDonald et al. (2018), we used data from Wolf (2014). Lists of the chosen elections for all states are given in Supplementary Appendix A, along with a brief description of Wolf’s data. Along with others (Backstrom et al. 1990; Gronke and Wilson 1999; McDonald 2014; Cervas and Grofman 2020; Powell et al. 2020; Abramowitz et al. 2006), for the purpose of determining political preference, we prefer statewide elections data over actual congressional elections which do not uniformly evaluate precincts in different congressional districts, partly because of incumbency and uncontested elections. Also, there are more statewide races for more precise statistics.

¹⁶It is convenient to express the rank axis r as $(j - 1/2)$ divided by the number of districts in the state, where j is the rank in the ordered list, so the rank axis uniformly spans the range 0 to 1 for all states.

¹⁷Indeed, districts 1, 2, 4, and 5 remain with the same party when any of the statewide elections are applied to them, whereas districts 3, 6, and 7 switched parties for different elections. Note that the uncertainty bars on the $r(v)$ data in Figure 1 are standard errors of the mean which indicate the precision of the partisan preferences of the districts. The standard deviations are the square root of 12 (number of elections) greater; while the standard deviation would better show the switching of parties in CD3 and CD7 from different statewide elections, the standard error of the mean better indicates the precision of the district preferences.

¹⁸Notice that we do not write that CO would be expected to elect three Democrats. By using statewide election results, we deliberately do not take incumbency in a district into account because the goal is to evaluate the partisan preferences of the 2011 map independently of political contingencies.

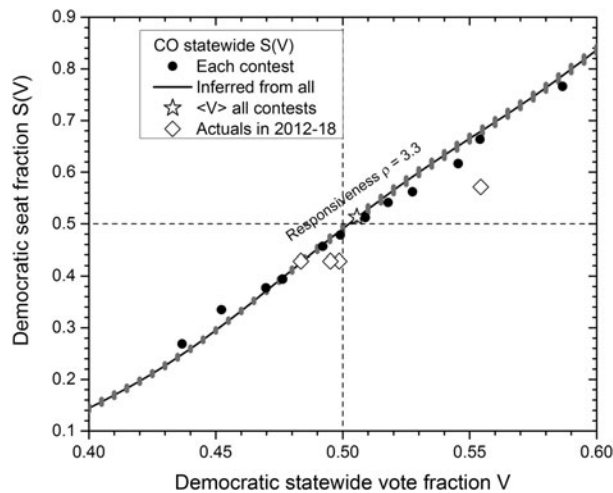


FIG. 2. $S(V)$ graph showing Democratic seat fraction preference vs. Democratic votes fraction curve for the Colorado 2011 enacted plan as a continuous solid curve. Each election ϵ generated an $S_\epsilon(V)$ curve (not shown) using proportional shift and fractional seats, and the shown solid $S(V)$ line is the average with standard error of the mean uncertainties as solid vertical bars. The solid circles are the results of each statewide election applied to the plan. The star shows the estimated seat fraction for the average statewide vote $\langle V \rangle$ for all elections. The diamonds show the fraction of seats won in the actual congressional contests carried out under the 2011 map.

we assign probabilities to districts thereby obtaining fractional seats.¹⁹ Furthermore, Figure 1 is for only one overall statewide vote share $\langle V_R \rangle = 49.4\%$ Republican, the complement of the Democratic two-party vote share. What is also needed is a different kind of figure, namely, the seats-votes $S(V)$ graph, which estimates the fraction of seats S for any two-party statewide vote fraction V . We use the proportional shift method to adjust each district vote v_i as the statewide vote shifts.²⁰ Our $S(V)$ graph for CO is shown in Figure 2.

An important result from $S(V)$ curves is the responsiveness of a map. How many seats are expected to change as the statewide vote V changes around V is just the slope dS/dV of the $S(V)$ curve at V . Of course, the slope varies with V if the curve is not linear and then the most relevant measure of responsiveness is $\rho = dS/dV$ at $V = \langle V \rangle$. Because the $S(V)$ curve for CO is quite linear, however, its responsiveness is nearly the same for all V . An important finding for the CO 2011 plan is that its responsiveness is quite high, namely, $\rho = 3.3$.²¹ This is far higher than the responsiveness $\rho = 1$ idealized by proportionality and even for the efficiency gap (EG) (McGhee 2014; Stephanopoulos and McGhee 2018) which

idealizes responsiveness of $\rho = 2$. It is closer to the classical cube “law” (Kendall and Stuart 1950) that has $\rho = 3$ at $V = 0.5$.

Each solid circle in Figure 2 shows an unshifted result from one election. Already, those points could be simply fit to provide a reasonable $S(V)$ curve. Our method is more precise, as is indicated by the standard errors of the mean shown as vertical bars on the $S(V)$ curve.²² The diamonds in Figure 2 show actual congressional seats for the 2012–2018 district elections. These agree with the $S(V)$ curve as well as could be expected, given that they must be integers in actual elections, and that they are affected by contingencies such as incumbency.

2.2. North Carolina

We turn next to North Carolina (NC), another nearly balanced state with Democratic average vote $\langle V \rangle = 51.5\%$ in the statewide elections in our data set. Figure 3 shows its $r(v)$ graph. It is very far from linear as shown by the dash-dot line, which is the best linear fit, so the corresponding slope of 1.85 is meaningless as a measure of responsiveness. Figure 3 emphasizes the well-known fact that the 2011 map for NC has three heavily packed Democratic districts and ten safe Republican districts that are not packed.

An interesting measure of bias is the declination δ , which Warrington (2018) has defined as the difference in the angles of the two straight lines joining the open circles in Figures 1 and 3.²³ Packing

¹⁹For example, a district that has a 50/50 partisan preference should count as half a seat for both parties. Details for fractional seat assignment and other technical aspects in this paragraph are given in Supplementary Appendix A. Fractional seats have been employed in various ways by Gelman and King (1994, 532) and Cottrell (2019), and the concept is nicely explained by McGann et al. (2016, 58–60) where the 5% variation that is employed in our study was suggested.

²⁰Although a proportional shift (Nagle 2015; 2019) is conceptually superior to the commonly used uniform shift, the difference between the two is not consequential for the results in this article as is shown in Supplementary Appendix B.

²¹This value of ρ means that a swing of 5% in the statewide vote V would change the estimated seat fraction by 16.5%. For CO with seven districts, that is a net swing of essentially one congressional seat.

²²Also, our method does not require choosing a fitting function, such as a bilogit (King and Browning 1987). Although a linear fit to the circles in Figure 2 would work well enough for CO, other states are considerably nonlinear.

²³Each outer circle locates the average vote for each party’s seats and its average district rank. The middle open circle in Figures 1 and 3 locates the rank that divides Democratic won seats from Republican seats. More details regarding this measure are provided in section 5.3.

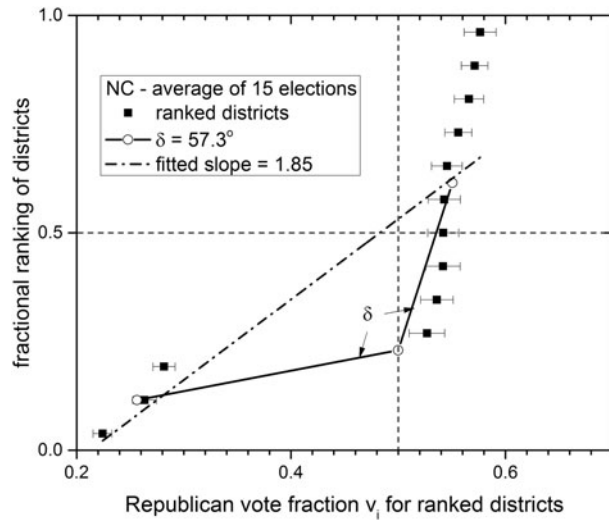


FIG. 3. Republican district rank vs. Republican district vote $r(v)$ graph for North Carolina 2011 plan. The difference in the angles of the two solid lines is the declination δ which equals 57.3° using all-or-nothing seats.

Democrats into a few districts moves the middle open circle in Figure 3 to smaller ranking of the districts won by Democrats which increases the differences in the angles of the lines connecting to the average party votes, so a positive value of δ correlates with an advantage for Republicans. This measure gives a large positive δ for NC, in contrast to a much smaller value for CO.

Figure 4 shows the $S(V)$ curve for NC.²⁴ As it is derived from the data in Figure 3, it too is highly non-linear. Again, the individual statewide election results shown by solid circles agree well with the $S(V)$ curve. So do the actual congressional election results for 2012–2018, but these also show that NC voters tended to vote more conservatively for Congress than for the statewide elections; the latter average $\langle V \rangle$ is shown by the star whereas the range of the actual congressional statewide vote average is centered near 0.48. This is important for evaluating responsiveness. The slope of the curve at $\langle V \rangle$ is quite high, $\rho = 4.0$ in Figure 4. However, the slope is close to zero near $V = 0.45$ which is the low end of the actual range of congressional elections, thereby indicating an unresponsive map favoring Republican incumbents in ten districts.

We will use Figure 4 to illustrate our definitions of several other quantities of interest. One such is the slope R of the arrow pointing from the center of the graph to the star on the curve.²⁵ A negative value of $R = -10$ signifies an anti-majoritarian result,

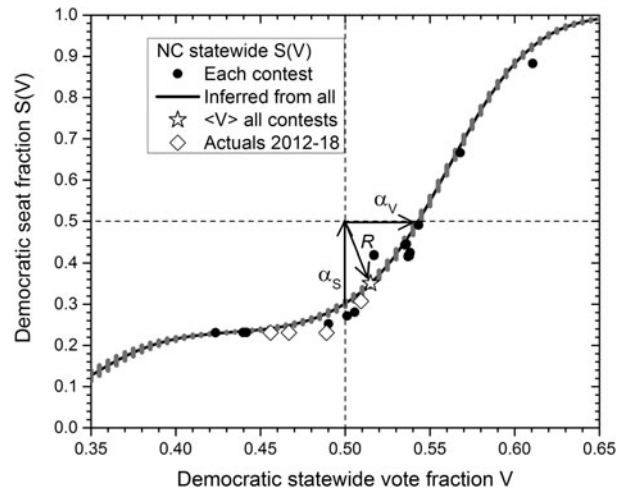


FIG. 4. Seats-Votes $S(V)$ graph for the North Carolina 2011 enacted plan. The solid $S(V)$ curve was obtained using proportional shift and fractional seats. Each contest is the result of each statewide election applied to the map. The star shows the average statewide vote for all elections. Diamonds show actual outcomes. The arrows indicate measures defined in the text.

fewer than half the seats for more than half the vote.²⁶ For unbalanced states R is positive and can be thought of as a winner’s bonus. For want of a better name, it may be thought of as an overall responsiveness.

An important seats-based measure, much employed in the literature, is the seats bias defined as the difference in the fraction of seats from 0.5 when $V = 0.5$. In this article we will designate this measure by α_S .²⁷ Its magnitude is the length of the vertical arrow in Figure 4, favoring Republicans by 20%. An important votes-based measure will be designated α_V ; it focuses on the estimated fraction of the vote for 50% of the seats. In Figure 4, 54.3% of the vote would have to be won by the Democrats in NC to win half the seats. We designate

²⁴The reader may wonder why we have chosen to use Republican (R) rank and votes axes in our $r(v)$ graphs and Democratic seats and votes in our $S(V)$ graphs. The reason is that doing so makes the $r(v)$ and $S(V)$ curves look more alike. For example, the median district with rank 0.5 has R vote greater than 0.5 in Figure 3 and the Democratic vote V to obtain half the seats is greater than 0.5 by very nearly the same amount. Indeed, using all-or-nothing seats for each district and uniform shift, the $r(v)$ curve with R axes becomes an $S(V)$ curve with D axes.

²⁵ $R = (S(\langle V \rangle) - 1/2) / (\langle V \rangle - 1/2)$. It is quite different from responsiveness ρ which is the tangent to the $S(V)$ curve given by the derivative dS/dV at $\langle V \rangle$.

²⁶The 2018 actual congressional election also was anti-majoritarian as shown by one of the diamonds in Figure 4.

²⁷ α_S is also the same as $\beta(0.5)$.

TABLE 1. A LIST OF THE STATES AND MANY OF THE QUANTITIES OF INTEREST OBTAINED IN THIS ARTICLE

State	$\langle V \rangle$	$S(\langle V \rangle)$	R	ρ	ζ	α_S	α_V	δ	GS	γ	Ω	β	LO	PR	EG
CA	59.2	73.6	2.6	2.1	0.3	-1.9	-0.6	-3	-1.8	-4	-0.6	2.9	9.2	-14	-5
IL	60.0	78.8	2.9	3.1	0.1	5.8	1.6	6	1.9	2	1.0	1.2	11.7	-19	-9
MD	59.3	85.7	3.7	1.0	0.8	-5.2	-1.0	-33	-2.6	-26	-2.4	1.4	4.1	-26	-17
MA	60.0	96.0	4.6	1.9	0.3	6.4	1.0	NA	2.0	-27	1.2	-2.9	NA	-36	-26
CO	50.6	51.3	2.2	3.9	0.2	1.1	0.3	0	0.8	1	0.2	1.0	0.3	-1	0
NC	51.5	35.1	-10.0	4.0	0.1	19.8	4.3	37	6.2	21	4.0	19.0	11.1	16	18
OH	51.3	41.4	-6.0	4.5	0.1	13.8	3.1	22	4.1	14	2.7	13.2	6.9	10	11
PA	52.9	43.6	-2.2	4.1	0.1	16.1	4.4	24	5.5	18	3.4	14.5	8.7	9	12
TN	41.6	19.2	3.6	0.8	1.1	16.2	2.9	35	6.1	24	3.7	-2.3	0.6	22	14
TX	40.4	27.8	2.3	1.1	0.7	5.0	1.2	17	5.2	12	1.9	-7.5	-3.0	13	3
SC	43.0	16.1	4.9	0.9	1.0	14.1	2.3	48	5.0	28	3.7	0.9	4.2	27	20

Values are expressed as percentages except for the δ angle and the responsiveness measures R , ρ , and ζ . Positive values for measures of bias favor Republicans.

- $\langle V \rangle$ is the average statewide two-party vote for Democrats.
- $S(\langle V \rangle)$ is fractional Democratic seats at $\langle V \rangle$.
- R is an overall measure of responsiveness or winner's bonus defined in Fig. 4.
- ρ is the slope of the $S(V)$ curve at $\langle V \rangle$.
- ζ is a measure of responsiveness defined in section 5.5.
- α_S is half the difference in party seats at $V=0.5$.
- α_V is the excess vote required for half the seats.
- δ is the value of the declination angle calculated using fractional seats and votes.
- GS is a global symmetry measure described in section 5.1.
- γ is the fair difference in party seats at $\langle V \rangle$ using ρ and defined in section 3.1.
- Ω is the average of the five previous measures, each normalized to the α_V scale.
- β gives the counterfactual symmetry difference in seats, defined in the introduction.
- LO shows values for the lopsided outcomes measure described in section 5.4.
- PR is the deviation from proportionality $S(\langle V \rangle) - \langle V \rangle$.
- EG is the efficiency gap (McGhee 2014) $S(\langle V \rangle) - 2\langle V \rangle + \frac{1}{2}$.

mM (not shown in table) is the median minus mean (McDonald and Best 2015).

the α_V value to be this vote at $S=1/2$ minus 0.5; it is the length (in %) of the horizontal arrow in Figure 4.²⁸

2.3. Other balanced states

The results in the previous section clearly identify the 2011 NC map as unfair to Democrats and unresponsive to voters. We have also drawn $r(v)$ and $S(V)$ graphs for the 2011 maps of two additional fairly evenly balanced states, Pennsylvania (PA) and Ohio (OH).²⁹ The α_S , α_V , δ , and R values for PA and OH are rather similar to NC in being unfair and anti-majoritarian, and quite unlike the results for CO which appear to be both fair and responsive.

Table 1 collects numbers for the quantities of interest mentioned in this section for balanced states. This table also shows values of other quantities that will be defined in subsequent sections. It further exhibits results for several unbalanced states to which we turn in the next two sections. Uncertainties in many of the quantities can be discerned by the reader from the error bars in the $S(V)$ graphs. For example, typical uncertainties in α_S are less than 1%. Uncertainties have not been estimated for δ ; the values in Table 1 were calculated using fractional

district seats and votes which differ from the values in the legends of our $r(v)$ graphs which used the original definition of Warrington (2018) that used all-or-nothing seat assignment.

3. REPUBLICAN MAJORITY STATES

This section examines politically unbalanced states with a substantial Republican majority.

3.1. Tennessee and the "Wall"

Tennessee (TN) has an average D vote $\langle V \rangle = 0.416$. Its $r(v)$ graph in Figure 5 shows a highly Democratically packed CD9, while CD5 leans Democratic and there are seven safe Republican districts that figuratively form a "wall." The graph is quite non-linear and this and the wall are reflected in the large positive value of the declination angle δ .

²⁸The goal of the α_V metric, which is based on the $S(V)$ graph, is similar to that of the popular median minus mean metric (mM) (McDonald and Best 2015) which is based on the data in the $r(v)$ graph.

²⁹These and graphs for other states will be available in Supplementary Appendix F.




FIG. 5. Rank-vote graph for Tennessee.

The corresponding $S(V)$ curve in Figure 6 agrees well with the statewide races (solid circles) and it agrees as well as can be expected with the actual congressional results given that those are constrained to be integers and that the incumbency effect may have helped retain the second Democratic seat. The values of the seats-based measure α_S and the votes-based measure α_V given in Table 1 are similar to those of the balanced states NC, OH, and PA. This comparison suggests that TN is also biased against Democrats. Unlike those states, R is positive, and there were no anti-majoritarian results (solid circles in Fig. 6) where more than half the votes yielded less than half the seats. The large value $R=3.6$ gives a substantial winner’s bonus of two seats compared to $R=1$ proportionality. The much smaller value of the responsiveness ρ compared to all the other states in Table 1 is consistent with the drawing of mostly safe seats for both parties as would be done in a bipartisan gerrymander.

Figure 6 introduces a new measure of bias which we will identify as γ . It uses the responsiveness ρ evaluated at $\langle V \rangle$ to calculate a seat fraction S_ρ . In Figure 6, an arrow emanates from the fair point at $S=0.5$ for $V=0.5$ with slope ρ . The end of the arrow at $V=\langle V \rangle$ locates S_ρ which is shown as a solid square in Figure 6. The seat difference between S_ρ and $S(\langle V \rangle)$ is defined as γ . Alternatively, Figure 6 shows an arrow with slope ρ emanating from the $S(\langle V \rangle)$ point to a projected $S(0.5)$ value, and then γ is the difference compared to the fair $S(0.5)=0.5$ value. Because γ uses both responsiveness and seats, it could be described as a “respon-

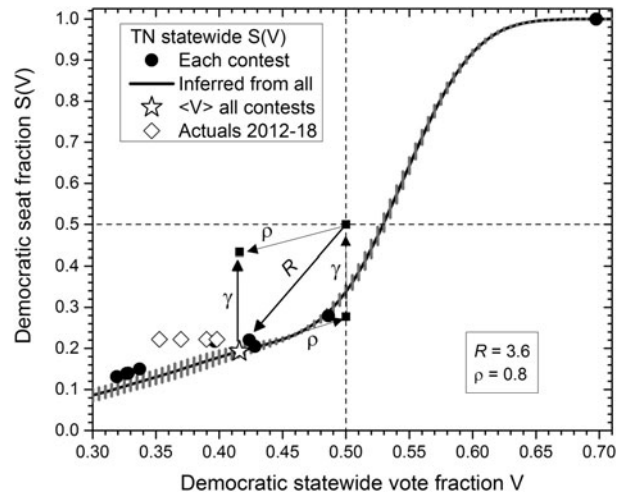


FIG. 6. Seats-votes graph for Tennessee.

siveness & seats” measure. A full discussion of the merits and characteristics of this measure will be deferred to section 5.2. Here we will only note that it is computed entirely from data at the statewide vote $\langle V \rangle$ and so it does not employ counterfactual shifting.

3.2. Flaw in the β measure

We now address the use of symmetry to estimate bias according to the measure defined as β in the introduction. The $S(V)$ curve in Figure 6 gives $S(\langle V \rangle)=0.192$ for Democrats with $\langle V \rangle=0.416$. The counterfactual symmetrically opposite vote is $1 - \langle V \rangle=0.584$, and at that value on the $S(V)$ curve the fraction of D seats is 0.857, so the counterfactual fraction of R seats is 0.143. As this is less than the factual D seat $S(\langle V \rangle)=0.192$, the symmetric measure of bias β posits that TN is biased in favor of Democrats, contrary to all the other measures. As this differs from expectations, one might question the accuracy of our $S(V)$ methodology.³⁰ However, the small uncertainties we

³⁰Of course, our value for β was obtained from a model for shifting the vote. The alternative uniform shift model essentially agrees by giving a similarly small value for β , although of the opposite sign. It may also be noted that TN is unusual compared to other Republican majority states in that there was one election, 2006 Governor, that had a strongly Democratic vote share of 0.70. This outlier election is responsible for giving the visibly large uncertainties for $0.3 < V < 0.45$ in Figure 6 and also for the rather larger standard errors of the mean for the districts in Figure 5 compared to other states. Nevertheless, our inclusion of this election demonstrates the robustness of our $S(V)$ curves over a large range of vote share.




FIG. 7. $r(v)$ graph for Texas.

obtain for our $S(V)$ curves have led us to reconsider the β measure and to conclude that it gets confounded in unbalanced states.

The key to understanding the flaw in the β measure comes from the $r(v)$ graph. If the vote were to shift strongly Democratic, the entire wall of safe Republican districts would fall. This is in contrast to the sole highly packed Democratic district that anchors $S(\langle V \rangle)$. A simple toy example illustrates this phenomenon. Consider a state with ten districts, nine of which have party A preference of 0.65 and one which has party A preference of 0.25 so the state has an overall preference $\langle V_A \rangle = 0.61$ for party A.³¹ Party A would then be expected to obtain nine seats with this set of preferences. For simplicity let us assume for the counterfactual that each district undergoes a uniform swing to party B by $\Delta V = 0.22$ which achieves the counterfactual statewide party B preference $V_B = 0.61$. Then there are nine districts with party A preference of 0.43 and one district with preference 0.03. For simplicity, let us use all-or-nothing district assignment which would then give party B ten seats in the counterfactual. The β measure would then draw the absurd conclusion that this state is biased in favor of party B because it would have obtained more seats if it had counterfactually received the same statewide vote fraction as party A actually received.³² This illustrates that by treating the two parties differently—most clearly illustrated by asymmetrical $r(v)$ graphs with a significant angle of declination—map drawers can achieve advantage for their party while the simple β measure of symmetry suggests the oppo-

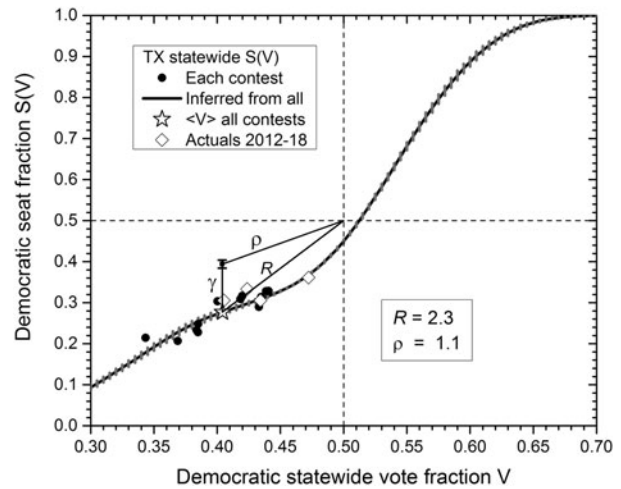


FIG. 8. $S(V)$ graph for Texas.

site. For a comparison that is relevant for the α_S measure, a uniform shift of 0.11 to $V = 0.5$ would still give nine seats to party A using all-or-nothing district seat assignment or 7.6 seats using fractional seats.

We wish to emphasize that this section only criticizes the β measure of bias, not the fundamental concept that fairness requires symmetry, as will become apparent in section 5.1.

3.3. Texas and South Carolina

We have also analyzed two other Republican majority states, Texas (TX) and South Carolina (SC). Results shown in Table 1 for the α_S and α_V measures for SC are similar to those for TN and the value of δ is even higher. The same measures for TX are somewhat smaller, but still biased in favor of Republicans. In contrast, the β value is nearly zero for South Carolina and even would assign substantial bias in favor of Democrats in TX, again indicating that the β measure is flawed.³³

³¹This assumes that all districts had the same number of voters, i.e., no turnout bias. Effective turnout bias is small in this study as is documented in Supplementary Appendix A.2.

³²Using proportional shift and fractional seats slightly alleviates the flaw in the β value by giving 9.6 instead of 10 seats to party B in the counterfactual. We also note that other examples reveal that the unfairness that is not diagnosed by the β measure comes not from packing minority party districts but from building a wall of majority party seats that are safe but not packed.

³³Our data for SC and TX also differ from that of NC in that there was no statewide election won by the Democrats. That means that the counterfactual $S(1 - \langle V \rangle)$ had to be extrapolated well beyond the vote share range of the available elections, in contrast to the less uncertain interpolation within the range of available vote share that was done for TN. Nevertheless, the general shapes of the extrapolated $S(V)$ curves for SC and TX are quite similar to the one for TN as shown in Figure 8 for TX.




FIG. 9. $r(v)$ graph for California.

Included here as Figure 7 is the $r(v)$ graph for TX. It also exhibits a wall of safe Republican districts in the $0.6 < V < 0.7$ range, although the TX wall is not as steep or high as the TN wall. The $S(V)$ graph in Figure 8 deviates from linearity, although somewhat less than the $S(V)$ graph for TN.

4. DEMOCRATIC MAJORITY STATES

4.1. California and Illinois

It is widely recognized that California (CA) has been a leading state for redistricting reform, so it is especially interesting to examine its congressional plan using our methods. The $r(v)$ graph in Figure 9 is basically linear with a small value of the declination δ favoring Democrats. The linear fit to the $r(v)$ data gives a slope of 2.1. This is the same as the responsiveness ρ shown in the $S(V)$ graph in Figure 10.³⁴ The winner's bonus $R=2.5$ for CA is considerably larger than proportionality ($R=1$) and somewhat larger than the EG ($R=2$), but still considerably smaller than the winner's bonuses in TN and SC. The α_S and α_V measures in CA favor Democrats, but less so than they favor Republicans in TX and much less than in TN and SC as shown in Table 1.

We have also analyzed the 2011 Illinois (IL) plan. Table 1 shows that the IL plan is somewhat more responsive than CA; that gives a relatively larger Democratic seat fraction in IL because both states have nearly the same average $\langle V \rangle$ vote. Its $r(v)$ graph is not as linear as for CA, but its δ is

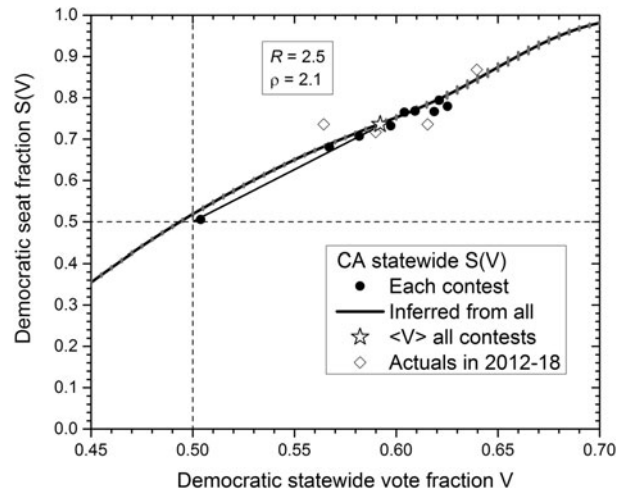


FIG. 10. $S(V)$ graph for California.

smaller and even favors Republicans. There appears to be a soft wall of six Democratic districts centered near 0.43 R vote in the $r(v)$ graph, but this wall appears to be close enough to $V=0.5$ such that, along with the extreme packing of Democrats in other Chicago districts, IL has small positive α_S and α_V values favoring Republicans. (Supplementary Appendix F contains the IL graphs.) We concur with the analysis of McDonald et al. (2018, 323) who wrote “Despite the outcries of unfairness in some quarters of the press and a court seeing a ‘blatant political move’ ... Democrats’ self-help maneuvering was largely a matter of tamping down some of the pro-Republican effects of residential patterns.” Similar comments were written by McGann et al. (2016, 105).

4.2. Maryland and Massachusetts

We now turn to two other Democratic majority states that differ substantially from CA and IL and also from each other. Maryland (MD) is widely regarded as having been intentionally gerrymandered by Democrats. Its $r(v)$ graph in Figure 11 shows that the infamous CD1 is a safe Republican district, and CD6 leans substantially Democratic. The complaint is that CD6 could be made more competitive or even leaning Republican if Democratic

³⁴Our $S(V)$ curve agrees quite well with one drawn by McGann et al. (2016, 78), although our uncertainties are smaller. They also provide values for the α_S measure for all states in their Table 3.A.2. Our value for CA is in excellent agreement and values for other states are also satisfactory considering that we use statewide elections and they used district elections.

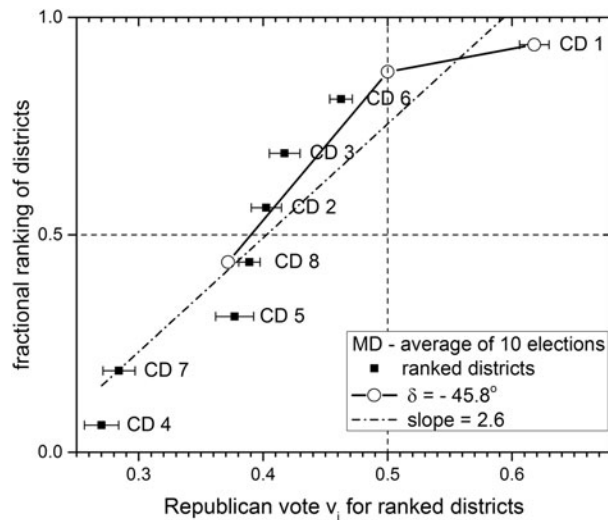


FIG. 11. $r(v)$ graph for Maryland.




FIG. 13. $r(v)$ graph for Massachusetts.

voters were exchanged with R voters in adjacent CD8. That would shift the middle declination circle down in Figure 11 which would reduce the very large negative value of δ and it would soften the already soft wall consisting of CD5, CD8, CD2, and CD3.

The $S(V)$ curve in Figure 12 shows that the winner's bonus R is quite large but the responsiveness ρ is much smaller, consistent with drawing relatively safe districts for both parties. Figure 12 also shows that the α_S and α_V measures favor Democrats, although Table 1 shows that the magnitudes of these bias values are much smaller than the corresponding values for five states that these measures purport to be biased in favor of Republicans.

Massachusetts (MA) has a similar two-party average vote share in our data compared to MD, but MA has not elected any Republicans to Congress under its 2011 plan whereas MD has always elected one. The $r(v)$ graph for MA in Figure 13 has a highly packed Democratic district CD7 and a wall of eight Democratic districts with Republican support in the range $0.35 < v_i < 0.45$. A linear fit just to this wall gives a large pseudo responsiveness of 7.3. The $S(V)$ graph in Figure 14 shows that a large winner's bonus like $R=4.6$ results in the dominant party obtaining nearly all the seats when $\langle V \rangle$ differs from 50% by as little as 10% of the vote. There were ample elections near $V=0.5$ to give credibility to the α_S and α_V measures which actually

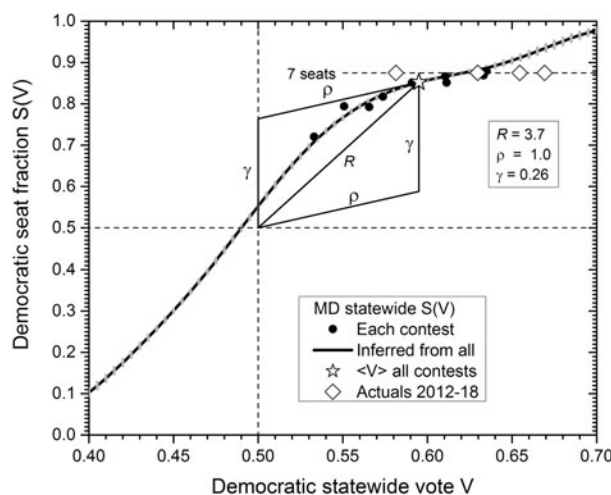


FIG. 12. $S(V)$ graph for Maryland.

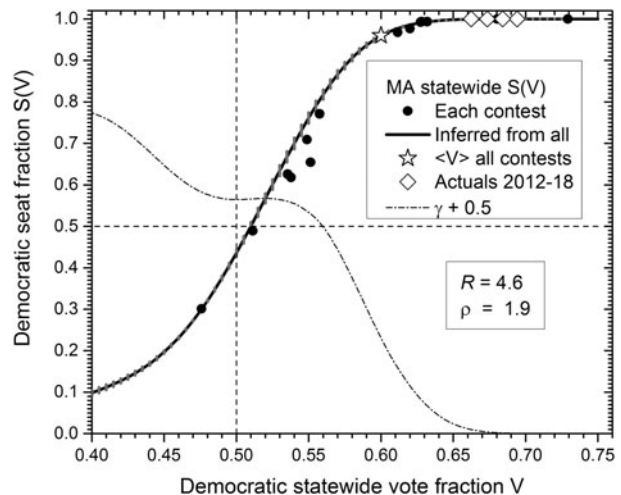


FIG. 14. $S(V)$ graph for Massachusetts.

indicate bias in favor of Republicans in the 2011 MA plan.³⁵ On the other hand, the α_S and α_V measures indicate bias for the Democrats in the MD plan, but that plan provides for one Republican seat whereas the MA plan essentially guarantees none.

5. ADDITIONAL IDEALS AND MEASURES

The MA and MD comparison just described illustrates the difficulty with assessing bias in unbalanced states. We have pursued this issue in Supplementary Appendix C by considering the impact of political geography. While we are enthusiastic about the concepts described there, the difficulty in carrying out the associated procedure motivates our consideration in this section of simpler measures in addition to the α_S , α_V , and β measures already introduced. Each of these additional measures implies an ideal mathematical form for zero bias that will be characterized.

5.1. Symmetry and the global GS measure

The concept that a district plan should treat parties symmetrically is appealing, so we have been disappointed that the β measure applied to $S(V)$ curves fails as emphasized in section 3.2. However, there are other ways to measure symmetry besides focusing just on the average statewide vote at $\langle V \rangle$ and its counterfactual at $1 - \langle V \rangle$ as is done in the β measure. In Supplementary Appendix D we describe a way to look at symmetry in the $r(v)$ graphs. However, that does not lead to a measure that we use in this article, so here we go directly to a measure of bias that uses global symmetry of the $S(V)$ graph.

The ideal for global symmetry is simply that $\beta(V)=0$ for all V . To see how global symmetry bias (GS) is measured for plans that do not have global symmetry, consider the example in Figure 15. The solid line in Figure 15 shows the $S(V)$ curve for Democrats in South Carolina and the dashed line shows the $S(V)$ curve for Republicans. The two curves differ considerably. Recall that we defined $\Delta S(V)$ as half the difference $S_D(V) - S_R(V)$ because then $\Delta S(1/2)$ is the α_S value of bias and $\Delta S(\langle V \rangle)$ is the β value. In Figure 15 β is much smaller than α_S because $\langle V \rangle$ is close to the point where the $S(V)$ curves for the two parties cross.

If it were not for our analysis of the $r(v)$ graphs in section 3.2, one might say that it is unclear from

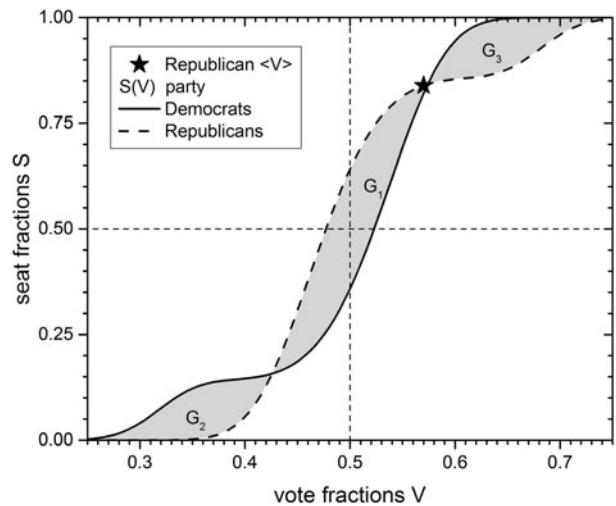


FIG. 15. Comparison of the $S(V)$ curves for the two parties in South Carolina. The solid line shows the Democratic seat fraction versus the Democratic two-party vote fraction and the dashed line is the corresponding curve for Republicans. Each curve is the inversion of the other. Half their difference at $V=0.5$ is the α_S value of bias and at $V=\langle V \rangle$ it is the β value of bias. The grey area between the two curves is the GS value of bias.

Figure 15 that the α_S measure is more valid than the β measure. However, in the case of uniform shifts the area identified by G_1 in the figure where $\Delta S(V)$ is negative has identical size to the sum of the two areas identified by G_2 and G_3 where $\Delta S(V)$ is positive.³⁶ Therefore, when the α_S value is large, there must be crossover values of $V_c \neq 1/2$ at which $\Delta S(V_c)=0$. In the cases of SC, MD, and TN, $\langle V \rangle$ is close to a crossover V_c . Similar crossover points at similar values of V_c also occur in states like NC, OH, and PA that are clearly biased; this is a further indication that α_S is a better measure of bias than β .³⁷ Nevertheless, it would be better not to have to choose a value of V at which to measure bias, especially when the chosen V is substantially different from $\langle V \rangle$.

The total shaded area $GS = G_1 + G_2 + G_3$ between the two $S(V)$ curves in Figure 15 provides such an alternative measure of bias that avoids this criticism

³⁵McGann et al. (2016, 77) and McDonald et al. (2018, 321) also noted a Republican bias in the MA plan.

³⁶When the proportional shift is used to construct the $S(V)$ curves, the equivalence of these areas is not exact, but it is still approximately true.

³⁷Likewise, the α_V votes measure of bias is better than evaluating the difference in votes required to obtain the same number of seats at the statewide $\langle V \rangle$.

of the α_S and β measures.³⁸ Because it integrates over all values of the vote V , it does not depend on making a choice of the value of V as the other seats-based measures do. Because the difference between the two curves can also be viewed as the differences in votes needed to obtain seats, GS is best described as a combination of seats-based and votes-based measures.³⁹

5.2. New γ measure of bias

A major concern with the α_S and α_V measures of bias for unbalanced states is that they do not evaluate near the average statewide vote $\langle V \rangle$. The next three subsections turn to three measures of bias that do evaluate exclusively at $\langle V \rangle$ and that therefore do not employ counterfactual V .

This subsection elaborates on our γ measure newly introduced in section 3.1 with values reported in Table 1 and shown graphically in Figures 6, 8, 12, and 14. It is rooted on the basic fairness principle of half the seats for half the vote V . It calculates an ideal fraction of seats

$$S_p = \frac{1}{2} + \rho(\langle V \rangle - \frac{1}{2}). \quad (5.2.1)$$

Equation (5.2.1) is rather like the ideal fraction of seats calculated by traditional proportional representation (PR) and by the efficiency gap (EG) except that PR imposes the value $\rho=1$ and the efficiency gap imposes $\rho=2$. In contrast, the γ measure uses the responsiveness determined at $V=\langle V \rangle$, thereby avoiding an arbitrary choice of the ρ factor, instead basing it on an empirical value at $\langle V \rangle$. Then, similarly to PR and EG, the γ measure of bias is the difference between the ideal fraction and the measured fraction

$$\begin{aligned} \gamma &= S_p - S(\langle V \rangle) \\ &= \frac{1}{2} + \rho(\langle V \rangle - \frac{1}{2}) - S(\langle V \rangle). \end{aligned} \quad (5.2.2)$$

As γ depends upon both responsiveness and seats, it could be described as a responsiveness & seats or R&S bias, although writing the simple γ symbol is more convenient. It is also useful to recall the definition of $R = (S(\langle V \rangle) - \frac{1}{2}) / (\langle V \rangle - \frac{1}{2})$ and rewrite

$$\gamma = (\rho - R)(\langle V \rangle - \frac{1}{2}). \quad (5.2.3)$$

Another favorable feature of the γ measure is that it especially penalizes bipartisan gerrymandering, which is characterized by a smaller value of ρ . In

unbalanced states R is generally larger than ρ . For Democratic states with $\langle V \rangle$ greater than $\frac{1}{2}$, a smaller ρ makes γ even more negative, appropriately reporting more bias. For Republican states with $\langle V \rangle$ less than $\frac{1}{2}$, a smaller ρ makes γ even more positive, again appropriately reporting more bias.⁴⁰

Next, we consider what kind of $S(V)$ curves give $\gamma=0$. An important class is unbiased linear curves that pass through the graphical center $(\frac{1}{2}, \frac{1}{2})$ with whatever slope $R=\rho$. CA, IL, and CO are essentially in this class. However, a general symmetric $S(V)$ curve will have values of γ that vary with V , usually with γ becoming negative as $\langle V \rangle$ increases from $\frac{1}{2}$. Of greater concern is non-symmetric curves which may have $\gamma=0$ at special values of $\langle V \rangle$.⁴¹ The most problematic of our states in this regard is MA. Figure 14 shows that $\gamma(V)=0$ at $V=0.56$ and, as V increases to $\langle V \rangle=0.6$, it indicates large negative bias in favor of Democrats. Although one might be inclined to welcome this as a common-sense result, note that the $S(V)$ curve eventually has to curve over at large V which necessarily gives a large negative γ . In the case of MA this occurs already at $\langle V \rangle$, so we are inclined to discount the γ result for MA. However, for other unbalanced, biased states like TN, $\gamma(V)$ has nearly the same value for all V near $\langle V \rangle$ because $S(V)$ has a nearly constant slope in that region. We will return to a quantitative analysis of the γ measure for other states in section 6.3.

5.3. Declination measure δ

This measure (Warrington 2018; 2019) utilizes the two solid lines that we have included in our

³⁸GS was previously named the geometric measure of bias B_G (Nagle 2015, 351) as it is the percentage of the geometric area within the total seats/votes box defined by 0 to 1 on both axes. Prof. Grofman has kindly brought to our attention that GS is a special case of one of the eight measures in Grofman (1983, 308).

³⁹Because the integral of the signed difference between the two curves is zero in the case of uniform shift and small even for proportional shift and fractional seats, GS is defined to take the absolute value of the difference which is the total shaded area in Figure 17. Which party is favored by the necessarily positive value of GS is then taken to favor the party favored by α_S and α_V . This can lead to ambiguity of the sign when the difference between the two curves is small, but such cases would be deemed acceptably fair with either sign.

⁴⁰This effect also applies in balanced states.

⁴¹This occurs when a straight line drawn from $(0.5, 0.5)$ is tangent to the $S(V)$ curve at V .

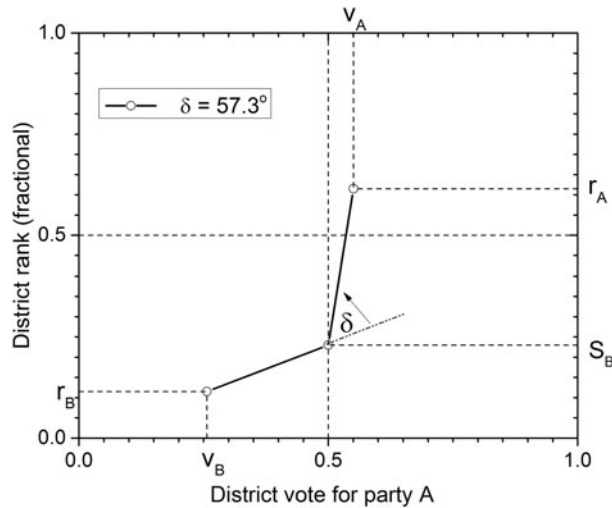


FIG. 16. Definitions of quantities r_A , v_A , r_B and v_B used to calculate ideal seats values for the δ measure of bias. The values shown are the same as those for NC in Fig. 3.

$r(v)$ graphs. Although the difference in angles δ between these two lines does not translate to the seats and votes quantities that most people would like to know, we like this construct because it visually illustrates walls of safe seats that are characteristic of unfair maps. In this subsection we review what the ideal declination ($\delta=0$) requires for $S(V)$ partly because it has interesting similarities and contrasts with the measure in the next subsection.

The basic definition of the declination depends upon just a few quantities in the $r(v)$ graph. Figure 16 illustrates these quantities; r_A and r_B are the average ranks of those districts that are won by parties A and B respectively, and v_A and v_B are the corresponding vote shares.⁴² Equating the angles of the two lines results in the ideal $\delta=0$. When the statewide vote for party A is V_A , straightforward algebra derives the following ideal $\delta=0$ relations (Campisi et al. 2019, 375; Katz et al. 2020, 175)

$$(\frac{1}{2}-v_B)-(v_A-\frac{1}{2}) = \frac{1}{2}-V_A \quad (5.3.1)$$

and

$$S_A-\frac{1}{2} = (V_A-\frac{1}{2})/[(4v_A-2)-(2V_A-1)]. \quad (5.3.2)$$

Equation (5.3.2) emphasizes that the seat fraction required for unbiased plans by the δ measure depends not only upon the vote V_A but also upon an additional characteristic, the average vote v_A .⁴³ Figure 17 shows the ideal $S(V)$ curves for two values of v_A , 0.6 and 0.7.

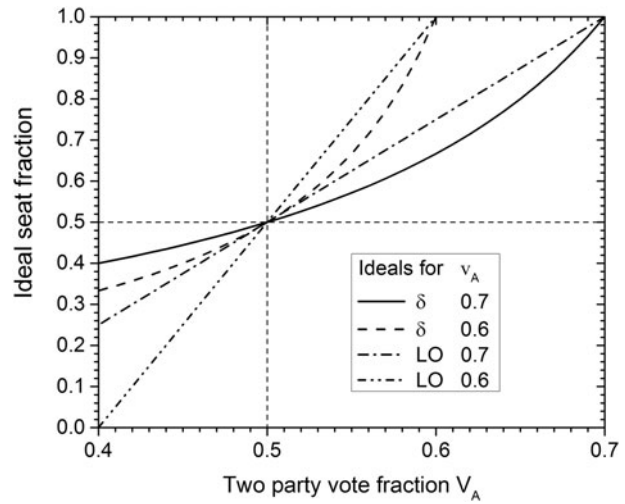


FIG. 17. Party A $S(V)$ curves for ideal declination ($\delta=0$) and ideal lopsided outcomes (LO=0, next subsection) for two values of v_A which is the average voter preference in districts won by party A.

Figure 17 shows that party A wins all the districts if $v_A=0.6$ (dashed curve) compared to 2/3 of the districts if $v_A=0.7$ (solid curve).⁴⁴ Any measure that allows this kind of variation means that measure violates a primary principle for measuring bias (McGhee 2017; Cover 2018), namely, that a partisan map drawer not be allowed to increase the seat fraction for a given vote fraction without incurring a change in the value given by the measure.⁴⁵ In the case of the declination measure, when party A is in power and expects $V_A > \frac{1}{2}$, it can increase S_A for the same V_A by decreasing v_A with no change in δ . Fortunately, decreasing v_A

⁴²The rank of the middle circle is twice the average rank r_B of seats won by party B, so this is S_B . Note that the district votes won by a party can have many distributions, such as (1) all having the same value or (2) all being quite different. The lines themselves are suggestive that the district votes are linear with their rank, but this is not generally required, although such linear $r(v)$ graphs do result in the declination not varying with the statewide vote when the uniform shift is employed. As noted earlier, we prefer to use fractional seats and vote shares in our calculations for Table 1, but it is convenient to use the simpler all-or-nothing seat assignment method (Warrington 2018) in Figure 19 and our $r(v)$ graphs.

⁴³Or v_B , but that is determined from Eq. (5.3.1) once V_A and v_A are given.

⁴⁴Using Eq. (5.3.1) the corresponding packing fractions $1-v_B$ for party B are 0.5 and 0.4 for $v_A=0.6$ and 0.7, respectively. These numbers emphasize that packing of party B is not allowed by the declination ideal.

⁴⁵Katz et al. (2020, 175) have similarly criticized the δ measure.

also increases the competitiveness of the districts and a partisan party A map drawer would be unlikely to reduce v_A low enough to create competitive districts when party A is the majority party. Unfortunately, when $V_A < 1/2$, Figure 17 shows that party A obtains more seats by having all districts more packed, in which case party A would encourage a bipartisan gerrymander while satisfying $\delta = 0$.

5.4. Lopsided outcomes measure of bias

This measure focuses on discriminatory packing. The ideal for this measure is that the excess vote share for districts won by party A averaged over those districts (namely, $v_A - 1/2$ in the previous section) equals the excess vote share for districts won by party B averaged over those districts (namely $1/2 - v_B$). Then, the measure of packing bias is⁴⁶

$$LO = (1/2 - v_B) - (v_A - 1/2). \quad (5.4.1)$$

A positive value of LO would indicate greater packing of party B voters and would therefore indicate a bias in favor of party A.

The ideal $LO = 0$ and Eq. (5.4.1) require that $v_B = 1 - v_A$. Together with the relation $V_A = v_A S_A + v_B S_B$, simple algebra gives the seats-votes relation,

$$S_A - 1/2 = (V_A - 1/2) / 2(v_A - 1/2). \quad (5.4.2)$$

Figure 17 shows the $S(V)$ curves for two values of v_A . Just as for the δ measure in the previous subsection, given statewide vote V_A in an unbalanced state with $V_A > 1/2$, a partisan map drawer could increase party A seat fraction by decreasing v_A , so the LO measure, like the δ measure, violates the principle that a map drawer should not be able to change the seat outcome without incurring a change in the value given by the measure of bias.⁴⁷ Although this LO measure is similar to the δ measure in the sense that both of their ideal curves depend upon v_A , Figure 17 shows that they have different ideal curves.⁴⁸ The ideal $S(V)$ curve obtained from the LO measure is similar to the proportionality and EG measures in that it is linear. Furthermore, its ideal responsiveness in Eq. (5.4.2) is $\rho = R = 1/2(v_A - 1/2)$; this relates to accommodating political geography which would be expected to produce smaller values of $v_A - 1/2$ for more homogeneous states. However, the actual value of LO is given by Eq. (5.4.1) and that depends upon $1/2 - v_B$ which is a separate variable for states with $LO \neq 0$.

5.5. Minimal inverse responsiveness ζ

We turn in this subsection from fairness to responsiveness because many reformers believe that districts should be competitive and that requires a statewide plan to be more responsive. Responsiveness makes it easier for voters to change the party in power, and it provides for more robust elections. It is therefore pertinent to consider maximal responsiveness as an ideal and to devise a measure of the extent to which a map does not satisfy the ideal. Such a formula has to take into account a maximum ideal value ρ_{\max} . Clearly the maximum occurs when all the districts j have partisan preferences $v_j = 0.5$. Of course, this is only possible when the statewide vote is $V = 0.5$. If one uses all-or-nothing assignment of a district's seat, then a small shift in statewide vote flips all the districts; that means $\rho_{\max} = \infty$. Infinity is not a subtractable number and that makes it awkward to quantify how a real map differs from the most responsive. To accommodate this, we define an inverse responsiveness measure ζ as:

$$\zeta = (1/\rho) - (1/\rho_{\max}). \quad (5.5.1)$$

Ideal responsiveness is then indicated by $\zeta = 0$, and more responsive maps are identified as having smaller values of ζ . With all-or-nothing district seat assignment $1/\rho_{\max}$ is simply zero. However, when the vote swings by a small amount, it is quite unlikely that all the fully competitive districts will swing to the same party, so it is preferable to use fractional seats, and from its slope we obtain $\rho_{\max} = 10$. We consider this value of ρ_{\max} to be a reasonable maximum responsiveness for balanced states. This reduces ζ by 0.1 in Eq. (5.5.1) compared to using the assignment of just 0 or 1 to a district's seat. For unbalanced states a derivation of ρ_{\max} is deferred to Supplementary Appendix E.

⁴⁶The usual calculation for each party simply adds the excess vote fractions just in those districts won the party and divides by the number of those districts. We use fractional seats and fractional excess votes for responsive districts, although differences are usually small.

⁴⁷Similarly to the ideal δ measure, party A could increase v_A in order to increase S_A when $V_A < 1/2$, Katz et al. (2020, 174) have also published Eq. (5.4.2) and criticized the LO measure.

⁴⁸This is also easily seen from Eq. (5.3.1) which is required for $\delta = 0$ whereas $LO = 0$ requires the left-hand side of Eq. (5.3.1) to be zero, and these two requirements differ except when $V_A = 1/2$.

6. COMPARISONS

Now that data have been shown and various measures have been described, we come to the most important part of this article in subsection 6.1 where we synthesize our results to measure total partisan bias in unbalanced states. Comparison of measures that don't work well is given in subsection 6.2, and the durability of measures is shown in section 6.3.

6.1. Best measures of partisan bias in unbalanced states

Different measures of bias measure different quantities. The α_S , β , and γ measures focus on seat bias in the S(V) graph, the α_V measure focuses on vote bias, and GS measures a combination of the two. Regarding the r(v) graphs, the LO measure focuses on packing districts, the declination δ measure is a geometric quantity that relates to both packing and cracking, and like the α_V measure, the more traditional median minus mean (mM) (McDonald and Best 2015), focuses on vote bias. Questions naturally arise: Which of these measures is best? And is there any agreement among them? These questions cannot be answered by examining a single map for a single state. In this subsection we show that they can, however, be answered by comparing the performance of measures across multiple states.⁴⁹

The comparison is facilitated by normalizing the values obtained by the measures to a common scale. This gives the graphical comparison shown in Figure 18. In this figure the α_S , δ , γ , and GS values were each normalized to the α_V values, so the meaning of the scale in the vertical axis is the percentage of the Democratic vote in excess of 50% for half the seats.⁵⁰ The average of these normalized measures is shown by the stars. We call this composite bias and assign it the symbol Ω . Values of Ω are given in Table 1.

The normalized measures shown in Figure 18 are visually well correlated as is confirmed by the r values in the legend.⁵¹ All the signs agree except for the MA γ outlier which has been omitted from the statistics for the reasons given in section 5.2. We hypothesize that the composite average Ω of these measures provides a better estimate of bias than any one of them. The t-test obtains the 99% confidence ranges for Ω shown by the uncertainty bars in Figure 18. These indicate substantial partisan

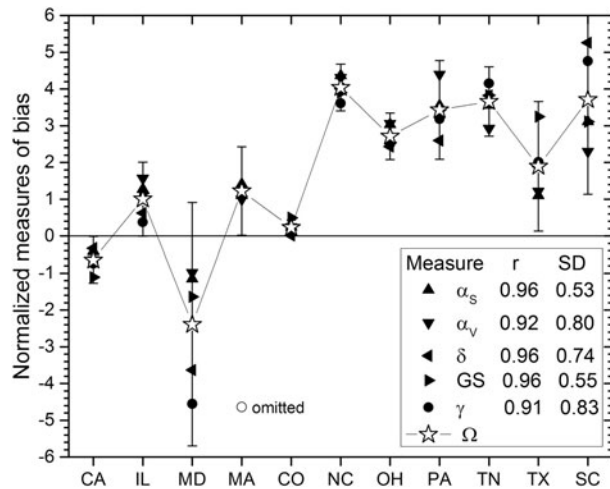


FIG. 18. Comparison of bias across states for five measures normalized to the α_V measure. Mean values are shown by stars with 99% confidence ranges. The legend shows the values for each measure of the Pearson r correlation and the standard deviation relative to the state mean averaged over states.

bias for Republicans in PA, NC, OH, SC, TN, and TX. At the 95% confidence level (not shown in Fig. 18), MD and CA are biased in favor of Democrats and MA and IL are biased in favor of Republicans.⁵²

For each measure the legend in Figure 18 also shows the average value over all states of its standard deviation (SD) relative to the stars in Figure 18. This individual SD of each measure provides an estimate of its uncertainty when it is applied separately from the other measures. At the 95% confidence level of two standard deviations, these SD values would indicate bias for all five of these measures in PA, NC, OH, SC, and TN and none of them indicates bias in CO and CA. In MD only the α_V measure falls short of the 95% confidence level. For the three states MA, IL, and TX, some

⁴⁹Although there have been many studies comparing states, we are unaware of this particular application to assess measures of partisan bias. We also note that one could also address these questions by comparing many maps for the same state.

⁵⁰The normalization factor for each measure was obtained by regression of its values to the α_V values. Standardizing disparate measures to obtain an average has also been accomplished by Stephanopoulos and Warshaw (2020).

⁵¹The pairwise comparisons of the five measures are also well correlated.

⁵²When comparing two plans for the same state, the uncertainties are likely to be correlated for each measure which would reduce the uncertainty for discriminating their relative bias ranges.

of the measures fall on either side of but usually close to the 95% confidence level. These results suggest that any one of these measures could often be an adequate measure of bias whenever it is inconvenient to calculate all of them. When all five can be calculated, we suggest using the composite value Ω . This combines the different indicators of bias that are associated with each of the individual measures, and it provides a way to obtain a statistical estimate of uncertainty as shown in Figure 18.⁵³

6.2. Measures that don't work in unbalanced states

Measures that disagree numerically from the averages shown in Figure 18 are shown in Figure 19. The β value agrees well for balanced states with the starred average because $\langle V \rangle$ is nearly the 50% that is used for the α measures. Figure 19 also shows values of β close to zero for unbalanced states which simply reiterates the point made in section 3.2 that there has to be a crossover point between the $S(V)$ curves for D and R in the unbalanced range as indicated in Figure 15.

The PR values derived from proportionality (see Supplementary Appendix C.1 for definition) in Figure 19 also agree with the other measures for balanced states, but they differ substantially in unbalanced states which we attribute to political geography. Such deviations would be smaller if accurate estimates would be made of PR_{SMD} in Eq. (C.2.1). Of course, PR is very strongly correlated with the winner's bonus R in unbalanced states.⁵⁴

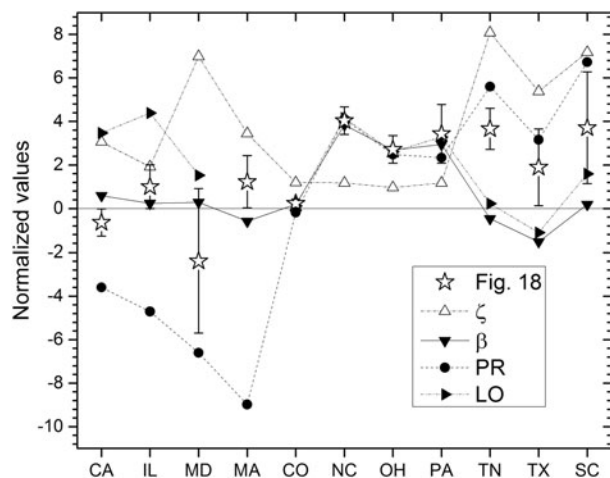


FIG. 19. Comparison across states for four measures and the average from Fig. 18, normalized to the α_V measure.

The LO values shown in Figure 19 are consistent with those of other measures in balanced states by indicating much more bias in favor of Republicans in NC, OH, and PA compared to CO. It has previously been noted (Wang et al. 2018, 313) that LO is not suited for detecting bias in unbalanced states. Indeed, the value of LO indicates a bias in favor of Republicans in MD and an even larger Republican bias in IL and CA. The Republican LO bias for CA is easily seen from its $r(v)$ graph in Figure 9. Because CA has essentially symmetrical $r(v)$ and $S(V)$ graphs, average packing is necessarily smaller in the R won districts than in the D won districts.⁵⁵ In TX, LO gives a bias in favor of Democrats. Both these unlikely values suggest that LO contains a spurious piece for unbalanced states which makes it report too much bias in favor of the minority party.⁵⁶

Values of ζ would not be expected to correlate well with any of the other values because ζ is a measure of responsiveness, not of partisan bias.⁵⁷ Surprisingly, plans in the heavily biased balanced states such as PA, NC, and OH have smaller, more favorable values of ζ than unbalanced states. This is due to their statewide Democratic vote occurring for $\langle V_D \rangle$ slightly greater than $1/2$ where the $S(V)$ curve rapidly rises. Evaluating ρ at smaller $\langle V_D \rangle$ where the actual congressional votes occurred would make ζ larger for these states. Even so, ζ would not be correlated with measures of bias in the unbalanced states.

6.3. Durability

It is desirable that the numerical values given by a measure of bias do not change much for typical maps when the vote swings by plausible percentages around the best estimate for the statewide average $\langle V \rangle$. Many measures of bias change dramatically with small changes in the vote when all-or-nothing assignment of a district's seat is used

⁵³Notice that these uncertainties are different than the generally smaller statistical uncertainties in the five basic measures which arise from the different statewide elections.

⁵⁴In the Democratic majority states, the appropriate correlation is between R values and negative PR values.

⁵⁵This follows from the definition of LO with respect to the difference in separation of the outer open circles from the inner open circle in the $r(v)$ graphs.

⁵⁶By comparison of Eq. (5.4.1) with Eq. (5.3.1), $\langle V_D \rangle - 1/2$ should be subtracted from LO.

⁵⁷Bipartisan gerrymandering reduces responsiveness but does not affect partisan fairness in balanced states.




FIG. 20. Non-durability of five measures of partisan bias and the ζ measure of responsiveness applied to the 2011 enacted map for NC, normalized to 1 at $V = \langle V \rangle$. Actual values at $\langle V \rangle$ are shown in Table 1.




FIG. 21. Non-durability of five measures of partisan bias and the ζ measure of responsiveness applied to the 2011 enacted map for TN, normalized to 1 at $V = \langle V \rangle$ except for the LO measure (see FN 62). Actual values at $\langle V \rangle$ are shown in Table 1.

for states with a small number of districts. We will call this instability. It does not reflect the true character of a plan or of a measure of bias. Instability is appropriately avoided by calculating fractional seats.

This subsection discusses a more fundamental type of change which is often described as durability (Grofman 2019; Wang et al. 2018) and as sensitivity (Warrington 2019). We define a measure to be absolutely durable if there is no change in the bias as the vote swings.⁵⁸ An example of an absolutely durable measure of bias is the α_S measure. Although α_S is evaluated at $V = 1/2$, that evaluation is performed on the $S(V)$ curve which doesn't change after it is drawn for a map using many past elections.⁵⁹ Likewise, the α_V and GS measures do not depend on choosing a statewide vote. Durability has also been a strong argument in favor of the median minus mean (mM) measure which uses the $r(v)$ graph.⁶⁰

Other measures of bias are not absolutely durable. Non-durability of the δ , PR, EG, ζ , LO, and γ measures is examined for NC in Figure 20 and for TN in Figure 21. These two figures illustrate that none of these measures is absolutely durable for all states.⁶¹ While the LO measure is most durable for NC, it is least durable for TN. The PR and γ measures are most durable for TN, but PR is least durable for NC while Figure 14 shows durability concerns for γ in MA. The δ measure is fairly durable for TN, although it varies more for NC. Of course, one way to superficially make these mea-

asures durable would be to average over a range of V centered on $\langle V \rangle$. However, such averaging tends to bring the value of bias back to what is obtained just using the central value obtained at $\langle V \rangle$, so the intrinsic non-durability remains. This argues in favor of using the durable measures of bias mentioned in the previous paragraph. The ζ measure of the departure from ideal responsiveness defined by Eq. (5.5.1) is not constant in Figures 20 and 21, simply because the $S(V)$ curves are not linear. For CA, which has nearly linear $S(V)$ and $r(v)$ graphs, the γ , ζ , and δ measures are quite durable. The Ω measure is the most durable measure in Figures 20 and 21.

⁵⁸This definition neglects systematic temporal variation differentially by precinct and district. Although the used data do not allow us to address precinct level changes, differential temporal variation of districts is contraindicated by the relative stability of their rankings over different elections.

⁵⁹The $S(V)$ curves we calculate are quite insensitive to large variations in V as seen by the small uncertainties in our $S(V)$ figures.

⁶⁰The mM measure is durable insofar as the median district vote swings equally with the statewide vote when the uniform shift is employed, and that is only marginally changed when proportional shift is employed for typical swings. The mM values do depend delicately on a single median district compared to the conceptually similar α_V measure, so we include only the latter in our set of core measures.

⁶¹The numerical values of the different measures are quite different which leads us to normalize all the values to unity at $V = \langle V \rangle$. An exception is made for LO in TN because its value is nearly zero at $\langle V \rangle$; the graph shows $1 + 10LO$, which emphasizes that LO changes sign in this range of V .

7. DISCUSSION

We preface this discussion by reminding the reader that our measures of bias are for total partisan bias, whatever its provenance. Overtly intended gerrymandering is just one part. Another part is unremediated bias due to a state's political geography and its map-drawing rules.

7.1. *Balanced states*

Although the primary focus of this article is on unbalanced states that lean strongly towards one party, it is important to emphasize that the thorny issues that arise for such states do not much affect balanced states. It does not matter whether one prefers proportionality (PR) or the efficiency gap (EG) (McGhee 2014) or the cubic law (Kendall and Stuart 1950) or bilogits (King 1989); when the vote share is close to 50%, they all converge on the unassailable criterion of 50% seat share for 50% vote share. While vote share is never precisely 50%, we have shown that many measures of bias agree rather well with each other for the 2011 plans of the four nearly balanced states CO, PA, NC, and OH.

Figure 18 shows this agreement for the seats-based measure α_S , the votes-based measure α_V , the geometric declination measure δ , a global symmetry measure GS, and a new responsiveness & seats-based measure γ . Figure 19 extends this agreement to the LO packing measure, the proportionality measure PR, and the extended symmetry measure β .⁶² This agreement is revealed in these figures by normalizing the values obtained in Table 1. Different measures provide values for different quantities, such as angles for the δ measure, seats for the α_S measure, and votes for the α_V measure. These quantities provide different perspectives on the underlying bias. These quantities naturally have different scales, nearly 60 degrees for δ and only 5% for α_V . Normalization provides a common scale that allows one more easily to compare different measures.

The agreement in this article of many measures of bias for balanced states confirms the unsurprising conclusion that the enacted 2011 plans in NC, OH, and PA are highly biased in favor of Republicans. The same analysis also reveals that the CO plan is quite fair. All this reassures map drawers that measuring bias in balanced states is quite achievable with a wide range of measures.

7.2. *Measures of bias for unbalanced states*

The preceding fine agreement of the normalized values obtained for all measures of bias when applied to balanced states only holds for some of those measures when applied to unbalanced states. Figure 18 in section 6.1 shows agreement for the seats-based measure α_S , the votes-based measure α_V , the geometric declination measure δ , a global symmetry measure GS, and a new measure γ that uses both seats and responsiveness at the average vote $\langle V \rangle$. We will call these the core measures and we propose that any one of these could be used to measure bias in unbalanced states.⁶³ This is a key finding in this article.

While the α_S , α_V , GS, and γ measures all use the $S(V)$ curves that shift the primary $r(v)$ data and employ fractional seats, the δ measure relies only on the primary $r(v)$ data.⁶⁴ The δ measure ideal is subject to manipulation as shown by Figure 17, but it is best at quantifying walls of safe, but not packed, districts favoring one party that are visually apparent in the $r(v)$ graphs. As shown in Figures 20 and 21 in section 6.3, values of bias obtained from the δ and γ measure depend upon the average vote $\langle V \rangle$, so they are free of the vote shifting counterfactual. However, they are not as durable as the other three $S(V)$ based measures. The α_S and the α_V values are absolutely durable but only because they are evaluated near the center of the $S(V)$ curve, which is displaced from the statewide vote for unbalanced states. However, the $S(V)$ curves for unbalanced states like SC, TN, and TX that have strongly biased values of α_S and α_V have quite similar shape to those of NC, OH, and PA whose values of α_S and α_V clearly show bias according to the fundamental principle that half the seats should be obtained for half the vote. The α_S and γ measures give a relatable quantity, seat bias, and the α_V measure gives vote bias, whereas the δ and GS quantities are more abstract. The normative principle for the α_S , α_V , and γ

⁶²One can add the efficiency gap measure EG to this list as its bias values are very near the same as the proportionality PR values for these balanced states.

⁶³It is, of course, logically possible that all five measures could be mutually wrong, but this appears unlikely in view of overall agreement of the results with a priori intuition.

⁶⁴Nagle (2019) has criticized the δ measure for being unstable, but our use of fractional seats makes it stable. Warrington (2019) has instead proposed adding fictitious buffer districts to achieve stability.

measures is the fundamental half the seats for half the vote, whereas the normative principal for the δ and GS measures is symmetry. GS measures absolute asymmetry over the whole range of the $S(V)$ curve, but which party is biased against requires appealing to one of the other measures. Clearly, none of these five measures is the silver bullet for evaluating bias in unbalanced states.

A natural question is: Why do these measures generally agree with each other? We suggest that each, in its own way and albeit imperfectly, is getting at the core underlying bias in a plan. We propose that together, like buckshot, they are sufficient for the job at hand, namely, to show that measuring bias is not limited to balanced states. This then leads us to consider a suitably normalized average, the composite measure Ω , which we suggest is better than any of the individual measures from which it is composed because it incorporates the various strengths of those measures.⁶⁵

7.3. *Two rejected measures of bias for unbalanced states*

The β measure of bias is based on a symmetry principle that seems so plausible that it came as a surprise that it does not work in states with a dominant party. Figure 19 shows that β has small values for all the unbalanced states. In section 3.2 we show how and why this occurs based on examination of the $r(v)$ graphs. Those graphs are highly asymmetric as seen by their values of δ , but they lead to $S(V)$ curves that are necessarily rather symmetric, but only near the value $\langle V \rangle$ for unbalanced states. For other values of V , especially those for which values of α_S , α_V and GS are obtained, the $S(V)$ curves are highly asymmetric. That is why these are better measures of bias than β . Withdrawing confidence in the β measure does not repudiate the concept that symmetry between parties is still an ideal. It simply means that a broader notion of symmetry is required. Supplementary Appendix D defines district symmetry (DS) in the $r(v)$ graph, and the GS measure uses global deviation from symmetry in the $S(V)$ graph. Of course, the δ measure directly illustrates deviations from symmetry in the $r(v)$ data. The α_S and α_V measures zoom in on the central portion of the $S(V)$ curve where asymmetry is robustly apparent. The γ measure only uses the symmetry principle at the 50/50 point along with counterfactual-free data.

The LO motivation of focusing on packing looks attractive a priori. However, values of LO bias in Figure 19 would claim that CA and IL are strongly biased in favor of Republicans while TN is fair and TX is biased in favor of Democrats, disagreeing strongly with our core measures and common perceptions. Section 5.4 argues that this indicates a systematic artifact that assigns too much bias in favor of the minority party in unbalanced states.

7.4. *Proportional ideal*

In Supplementary Appendix C.1 an argument is reiterated that proportionality is the ideal fair representation for partisan groups of voters and that the PR measure defined as $S - V$ is the appropriate measure of bias. But then in Supplementary Appendix C.2 the well-known incompatibility of proportionality and the single member district system due to partisan geography is reiterated. In particular, MA is an example of a state where relative geographic uniformity of partisan voters makes it essentially impossible to draw congressional districts that would durably give Republicans a proportional share.

A redistricting procedure based on the proportionality ideal would be to find plans that come closest to satisfying proportionality for the most likely vote share $\langle V \rangle$.⁶⁶ As emphasized by the example of PA (Nagle 2019), a serious drawback is a state imposing constraints on the set of acceptable plans that makes it impossible to achieve a fair map. Furthermore, the ensuing PR_p estimate of bias defined in Supplementary Appendix C.2 can't be obtained just from a plan itself but requires drawing many other possible maps. That is impracticable and certainly limits the number of people who could evaluate maps they have drawn.

Of course, researchers are now capable of generating ensembles of many maps by using various computer simulation techniques. As described in the introduction, often that is done to establish

⁶⁵The estimates of uncertainty shown in Figure 18 may also be thought of as providing an estimate of the degree of imperfection inherent in each measure relative to the underlying bias that we estimate by Ω .

⁶⁶Since there would be other maps that would grant the majority party a larger winner's bonus R , we are aware that it may be deemed naïve even to suggest such a criterion. However, a criterion that would guarantee minority party seats, by packing minority party voters if necessary, has recently been favorably discussed by Katz et al. (2020, Appendix B).

benchmarks or baselines that give average values of bias and then measuring intentional bias as differences from the baseline average. This has been used for court cases to evaluate intentional gerrymandering, but that is not what PR_p in Supplementary Appendix C.2 attempts to measure. It measures bias as the difference from the plan that comes closest to proportionality, and that putatively fairest plan may even be an extreme outlier in the set of computer generated plans when a state's political geography favors one party.⁶⁷ This is consistent with the quote at the end of section 4.1 from McDonald et al. (2018, 323) to the effect that the Illinois legislature drew an outlier map in 2011 that turned out to be rather fair. The challenge with implementing this approach is in determining what is the fairest possible plan with respect to proportionality. For that reason, in the body of this article we have retreated to our five core measures α_S , α_V , δ , γ , and GS and their composite Ω that appear to measure total bias reliably without this complication and the related complication of having to consider political geography and its interaction with the state's rules for map drawing.

7.5. Responsiveness

We have considered several quantities related to responsiveness. The basic one is ρ ; it measures how responsive the plan is at the expected value of the statewide preference $\langle V \rangle$. Bipartisan gerrymanders that lock in safe seats for either party have small values of ρ , so it is an important measure for reformers. It is most easily measured from the $S(V)$ curve. Table 1 shows values of ρ that are suspiciously small for MD, SC, TN, and TX. We have also defined a reciprocal responsiveness measure ζ in section 5.5 that allows comparison of a plan to a realistic maximum responsiveness that takes into account the difference between unbalanced and balanced states. This responsiveness measure is not well correlated with bias, reaffirming that fairness and responsiveness are two separate quantities (King and Browning 1987). We suggest that $\zeta=0$ defines an appropriate ideal responsiveness.

An overall responsiveness R , by taking on a negative value, identifies an anti-majoritarian outcome in balanced states, and it assesses the winner's bonus in unbalanced states. In the latter case a large value of R is associated with unfairness to the minority party as in MA, MD, SC, and TN, but it likely

has a minimum value that depends on the political geography of a state. This measure of responsiveness is therefore entangled with bias in both balanced and unbalanced states.

The ratio ρ/R is also of interest; a ratio near 1 in unbalanced states, such as CA and IL, follows from their linear district symmetry (DS) as seen in Supplementary Appendix D. Bipartisan gerrymandering in unbalanced states reduces ρ/R by reducing ρ . This ratio is also small in balanced states that are biased because R is generally large and even negative for anti-majoritarian states like PA, NC, and OH.

Another measure of responsiveness that we have not focused on is the fraction of competitive districts. A rough way to estimate this is to multiply ρ by a competitive range of votes, typically 10%.⁶⁸ It is somewhat discouraging that this means that there would only be 20% competitive districts even with $\rho=2$ super proportional responsiveness. In any case, the $S(V)$ curve is the most appropriate vehicle to assess responsiveness/competitiveness.

7.6. Future application to redistricting

In order to focus on the thorny issue of measuring bias in unbalanced states, this article has performed an analysis across states using the single 2011 plan for each state. Moving forward to redistricting in 2021, we propose using five core measures in all states for the purpose of comparing the bias in different proposed plans. Although any one of these measures would provide an estimate of the total bias, a better estimate would average their normalized values as in our composite Ω measure. The normalization factors we obtain from our analysis across states are, for normalization to the α_V measure, 0.21 for α_S , 0.083 for δ , 0.66 for GS, and 0.17 for γ . These would provide a first approximation for new maps in a state. However, it is quite likely that these normalization factors will be different for different states. For example, 12 maps were analyzed for α_S and α_V using the same 7s election data in Table 1 in Nagle (2019, 69). Regression on

⁶⁷Cf. the introduction and footnote 3. Also, note that this section applies equally to other measures like the efficiency gap that has an ideal $S(V)$ curve.

⁶⁸Better than a fixed range with a sharp cutoff would be a gradual fractional district measure. While competitiveness at the statewide level and responsiveness have distinct measures, they are clearly closely related.

these data gives 0.31 for the normalization of α_S to α_V , rather larger than the above value of 0.21. Therefore, as more maps are drawn, more precise values of the normalization factors can be obtained.⁶⁹ The main result of this article remains the identification of the core measures of bias for all states and the analysis that devises a composite measure Ω .

8. CONCLUSION

Based on the results in this article, we propose that total partisan bias can be measured reliably for the problematic unbalanced states as well as for balanced states. Any or all of the five core measures or their composite Ω measure can be used with past election results to assess partisan bias in maps drawn by redistricting commissions, in maps drawn by individuals using various map drawing software, and in maps drawn by computer algorithms. Furthermore, responsiveness can also be obtained immediately from S(V) graphs. Measures of fairness and responsiveness that are perceived to be reliable would encourage the enactment of election law that would include these fundamental concepts, and in a form that could have justiciable bite.

SUPPLEMENTARY MATERIALS

Supplementary Appendix A
 Supplementary Appendix B
 Supplementary Appendix C
 Supplementary Appendix D
 Supplementary Appendix E
 Supplementary Appendix F

REFERENCES

- Abramowitz, A. I., B. Alexander, and M. Gunning. 2006. "Incumbency, Redistricting, and the Decline of Competition in U.S. House Elections." *Journal of Politics* 68: 75–88.
- Altman, M., and M. D. McDonald. 2018. "Redistricting by Formula: An Ohio Reform Experiment." *American Politics Research* 46: 103–31.
- Backstrom, C., L. Robins, and S. Eller. 1990. "Establishing a Statewide Electoral Effects Baseline." In *Political Gerrymandering and the Courts*, ed. B. Grofman, 145–70. New York, NY: Agathon Press.
- Best, R., J. S. Krasno, D. B. Magleby, and M. D. McDonald. 2020. "Detecting Florida's Gerrymander: A Lesson in Putting First Things First." *Social Science Quarterly* 101: 37–52.
- Bradlee, Dave. *Dave's Redistricting*. <<https://davesredistricting.org>>.
- Campisi, M., A. Padilla, T. Ratliff, and E. Veomett. 2019. "Declination as a Metric to Detect Partisan Gerrymandering." *Election Law Journal* 18: 371–87.
- Cervas, J. R., and B. Grofman. 2020. "Tools for Identifying Partisan Gerrymandering with an Application to Congressional Districting in Pennsylvania." *Political Geography* 76: 102069(1–12).
- Chen, J. W., and J. Rodden. 2013. "Unintentional Gerrymandering: Political Geography and Electoral Bias in Legislatures." *Quarterly Journal of Political Science* 8: 239–69.
- Cottrell, D. 2019. "Using Computer Simulations to Measure the Effect of Gerrymandering on Electoral Competition in the Us Congress." *Legislative Studies Quarterly* 44: 487–514.
- Cover, B. P. 2018. "Quantifying Partisan Gerrymandering: An Evaluation of the Efficiency Gap Proposal." *Stanford Law Review* 70: 1131–233.
- Duchin, M., T. Gladkova, E. Henninger-Voss, B. Klingensmith, H. Newman, and H. Wheelen. 2019. "Locating the Representational Baseline: Republicans in Massachusetts." *Election Law Journal* 18: 388–401.
- Gelman, A., and G. King. 1994. "A Unified Method of Evaluating Electoral Systems and Redistricting Plans." *American Journal of Political Science* 38: 514–54.
- Gelman, A., G. King, and A. C. Thomas. 2012. "Judgeit Ii: A Program for Evaluating Electoral Systems and Redistricting Plans." *Gary King*. <<http://gking.harvard.edu/judgeit>>.
- Grofman, B. 2019. "Tests for Unconstitutional Partisan Gerrymandering in a Post-Gill World." *Election Law Journal* 18: 93–115.
- 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: 2–35.
- Grofman, Bernard. 1983. "Measures of Bias and Proportionality in Seats-Votes Relationships." *Political Methodology* 9: 295–327.
- Gronke, P., and J. M. Wilson. 1999. "Competing Redistricting Plans as Evidence of Political Motives—the North Carolina Case." *American Politics Quarterly* 27: 147–76.
- Katz, J. N., G. King, and E. Rosenblatt. 2020. "Theoretical Foundations and Empirical Evaluations of Partisan Fairness in District-Based Democracies." *American Political Science Review* 114: 164–78.
- Kendall, M. G., and A. Stuart. 1950. "The Law of Cubic Proportion in Election Results." *British Journal of Sociology* 1: 183–97.
- King, G. 1989. "Representation through Legislative Redistricting—a Stochastic-Model." *American Journal of Political Science* 33: 787–824.
- King, G., and R. X. Browning. 1987. "Democratic Representation and Partisan Bias in Congressional Elections." *American Political Science Review* 81: 1251–73.

⁶⁹Furthermore, following up on footnote 49, estimates in the difference of the bias between two plans should be addressed using Student's t-test on the differences in the values of the five normalized measures.

- McDonald, M. D., D. B. Magleby, J. Krasno, S. J. Donahue, and R. Best. 2018. "Making a Case for Two Paths Forward in Light of *Gill v. Whitford*." *Election Law Journal* 17: 315–27.
- McDonald, M. P. 2014. "Presidential Vote within State Legislative Districts." *State Politics & Policy Quarterly* 14: 196–204.
- McDonald, Michael D. 2009. "The Arithmetic of Electoral Bias, with Applications to U.S. House Elections." *APSA 2009 Toronto Meeting Paper*. Available at SSRN: <<https://ssrn.com/abstract=1451302>>.
- McDonald, Michael D., and Robin E. Best. 2015. "Unfair Partisan Gerrymanders in Politics and Law: A Diagnostic Applied to Six Cases." *Election Law Journal* 14: 312–30.
- McCann, A. J., A. S. Smith, M. Latner, and J. A. Keena. 2016. *Gerrymandering in America: The House of Representatives, the Supreme Court, and the Future of Popular Sovereignty*. 1st ed. New York, NY: Cambridge University Press
- McGhee, E. 2017. "Measuring Efficiency in Redistricting." *Election Law Journal* 16: 417–42.
- McGhee, E. 2014. "Measuring Partisan Bias in Single-Member District Electoral Systems." *Legislative Studies Quarterly* 39: 55–85.
- Metric Geometry and Gerrymandering Group. *Districtr*. <<https://districtr.org/>>.
- Nagle, J. F. 2017. "How Competitive Should a Fair Single Member Districting Plan Be?" *Election Law Journal* 16: 196–209.
- Nagle, J. F. 2015. "Measures of Partisan Bias for Legislating Fair Elections." *Election Law Journal* 14: 346–60.
- Nagle, J. F. 2019. "What Criteria Should Be Used for Redistricting Reform?" *Election Law Journal* 18: 63–77.
- Powell, R. J., J. T. Clark, and Dube. M. P. 2020. "Partisan Gerrymandering, Clustering, or Both? A New Approach to a Persistent Question." *Election Law Journal* 19: 79–100.
- Rodden, J. A. 2019. *Why Cities Lose*. New York, NY: Hachette Book Group.
- Stephanopoulos, N. O. 2013. "The Consequences of Consequentialist Criteria." *UC Irvine Law Review* 3: 669–715.
- Stephanopoulos, N. O., and E. M. McGhee. 2018. "The Measure of a Metric: The Debate over Quantifying Partisan Gerrymandering." *Stanford Law Review* 70: 1503–68.
- Stephanopoulos, N. O., and C. Warshaw. 2020. "The Impact of Partisan Gerrymandering on Political Parties." *Legislative Studies Quarterly* 45: 609–643.
- Thornburgh, D. *Draw the Lines*. <<https://drawthelinespa.org/>>.
- Wang, S. S.-H. 2016. "Three Tests for Practical Evaluation of Partisan Gerrymandering." *Stanford Law Review* 68: 1263–321.
- Wang, S. S.-H., B. A. Remlinger, and B. Williams. 2018. "An Antidote for Gobbledygook: Organizing the Judge's Partisan Gerrymandering Toolkit into Tests of Opportunity and Outcome." *Election Law Journal* 17: 302–14.
- Warrington, G. S. 2019. "A Comparison of Partisan-Gerrymandering Measures." *Election Law Journal* 18: 262–81.
- Warrington, G. S. 2018. "Quantifying Gerrymandering Using the Vote Distribution." *Election Law Journal* 17: 39–57.
- Wolf, Stephen. 2014. "DRA Update: 2012 President & Downballot Election Results Estimates + Display Table Template." *Daily Kos*. August 4. <<http://www.dailykos.com/stories/2014/8/4/1318876/-DRA-Update-2012-President-Downballot-Election-Results-Estimates-Display-Table-Template>>.

Address correspondence to:
John F. Nagle
Department of Physics
Carnegie Mellon University
Wean Hall
Pittsburgh, PA 15213
 USA

E-mail: nagle@cmu.edu

Received for publication June 29, 2020; received in revised form November 14, 2020; accepted December 6, 2020; published online February 16, 2021.

On measuring two-party partisan bias in unbalanced states

John F Nagle and Alec Ramsay

Supplementary Appendix A: Data

A.1 Source.

The data come from statewide election returns compiled and disaggregated to voter tabulation districts (VTDs) by Stephen Wolf at Daily Kos (Wolf 2014). Since U.S. elections at all levels are administered by county or local governments, Wolf disaggregated county level returns to assign votes to VTDs. While this is not ideal, first order discrepancies were avoided by using available votes cast in the VTD in the 2008 presidential election and the proportion of the county’s population living in a VTD. McDonald, et al. 2018 have reported that Wolf’s disaggregation method was sufficiently accurate by comparison to VTD level data that were available in several states. Even if Wolf’s data were deemed unsuitable to determine bias for the actual states in question, these data are sufficient to test measures of bias as this can be done for any consistent set of data.

State	2004 President	2006 Senate	2006 Governor	2006 Attorney General	2006 Secretary of State	2006 Comptroller	2006 Treasurer	2006 Auditor	2008 President	2008 Senate	2008 Attorney General	2008 Secretary of State	2008 Treasurer	2010 Senate	2010 Governor	2010 Attorney General	2010 Secretary of State	2010 Comptroller	2010 Treasurer	2010 Auditor	2012 President	2012 Senate	2012 Governor	2012 Attorney General	2012 Secretary of State	2012 Treasurer	2012 Auditor	2013 Senate (special)	
CA	✓								✓												✓	✓							
IL	✓													✓	✓	✓	✓	✓	✓		✓								
MD	✓	✓	✓	✓		✓			✓					✓	✓						✓	✓							
MA	✓	✓	✓	✓					✓	✓				✓	✓	✓	✓	✓		✓	✓	✓	✓					✓	
CO	✓		✓	✓	✓									✓	✓	✓	✓				✓								
NC	✓								✓	✓		✓	✓	✓	✓	✓	✓				✓	✓		✓	✓	✓	✓	✓	
OH	✓	✓	✓	✓	✓		✓	✓	✓		✓			✓	✓	✓	✓		✓	✓	✓	✓							
PA	✓	✓	✓	✓					✓					✓	✓						✓	✓		✓					
TN	✓	✓	✓						✓	✓					✓						✓	✓							
TX	✓	✓	✓	✓					✓	✓					✓	✓					✓	✓							
SC	✓		✓						✓	✓				✓	✓	✓	✓	✓			✓								

Table A.1 Used statewide elections for each state as well as a downballot composite for each state.

The states that were used are given in Table A.1 Our initial choice of states to analyze was to follow McDonald, et al. 2018 who chose the four states, IL, MA, MD, all unbalanced D states, along with balanced OH. We then wanted to have a comparable number of unbalanced R states, choosing TX as the largest one to match CA which we thought would be valuable as the widely touted most fairly districted state, as well as two smaller unbalanced R states, TN and SC, to match MA and MD. We also added notoriously biased balanced states PA and NC as well as CO which looked *a priori* to be the least biased balanced state. The used elections and results are in the research/elections directory of this GitHub repository:redacted for review. Elections that had significant third-party votes were not used.

A.2. Calculations for $r(v)$ and $S(V)$ graphs

For each congressional district j in each election we calculated the two-party Democratic vote fraction $v_j = D \text{ vote} / (D \text{ vote} + R \text{ vote})$. For each district, the average vote fraction and the standard error of the mean was calculated and plotted in the $r(v)$ graphs. The statewide two-party vote fraction in each election was calculated in two ways: 1) the average of the district vote fractions and 2) the actual two-party vote fractions in the state. The difference of 1) minus 2) is the turnout bias (McDonald 2009). Table A.2 shows values for the CA elections that we used. Table A.3 shows average values of turnout bias for all analyzed states.

Election	08P	Down	10AG	10C	10G	10SS	10T	12P	12S	Mean	STD
$\langle V \rangle$ %	62.1	59.7	50.4	60.4	56.7	58.2	60.9	61.9	62.5	59.2	3.8
t-bias %	-0.2	-0.6	-0.7	-0.8	-0.8	-0.9	-0.8	-0.6	-0.5	-0.7	0.2

Table A.2 Average statewide vote $\langle V \rangle$ and turnout bias for CA elections

state	CA	IL	MD	MA	CO	NC	OH	PA	TN	TX	SC
% t-bias	0.7	0.6	0.4	0.7	-0.3	0.2	0.3	-0.2	0.3	2.2	0.1
% std	0.2	0.3	0.2	0.4	0.1	0.3	0.2	0.2	0.2	0.3	0.2

Table A.3 Average turnout bias for all states.

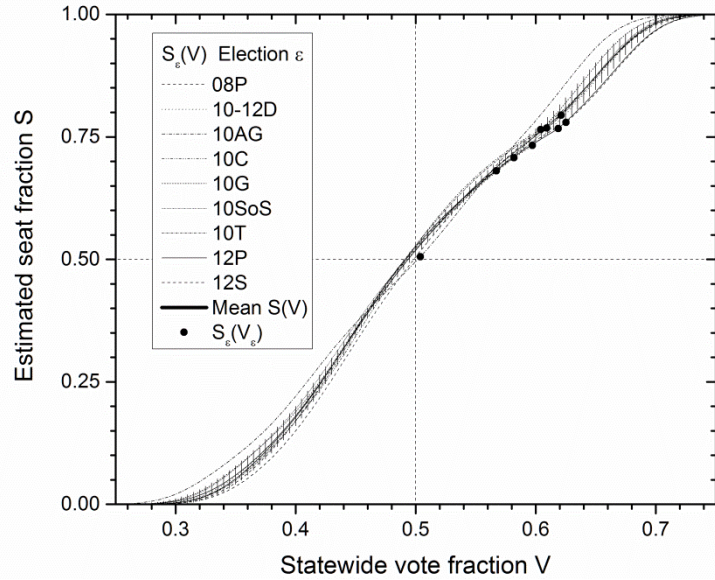
Because turnout bias is small, it has been ignored in the main text where the tables and figures show the actual two-party vote fractions.

For each election ε a seats/vote $S_\varepsilon(V)$ graph was calculated for statewide vote V_k from 25% to 75% in intervals of 0.5%. For each V_k a voter preference $v_j(V_k)$ was calculated for each district j using a proportional shift from the factual $v_j(\langle V \rangle)$ as defined by Nagle 2015, 2019. When the shifted statewide party preference V_k for party A is smaller than $\langle V \rangle$, party A voters in each district are shifted to party B. The proportional shift shifts the same fraction of party A voters in each district. In contrast, the uniform shift shifts the same fraction of total voters from party A to party B regardless of how few party A voters there may be.¹ The fraction of a seat estimated for each shifted district was then calculated using party seat probability $P(V) = 1 - \frac{1}{2}(1 + \text{prob}((V - \frac{1}{2})/0.04))$ where prob is the usual probit function, here with variance 0.04.² The sum of seat fractions over all districts was then $S_\varepsilon(V)$ for that election. Figure A.1 shows $S_\varepsilon(V)$ curves for several elections in CA. Then, the $S(V)$ curve is the average of the $S_\varepsilon(V)$ curves and the error bars in the $S(V)$ graphs are the standard error of the means at each V . Finally, solid circles in the $S(V)$ figures show $S_\varepsilon(V_\varepsilon)$ where, for each statewide election ε , V_ε is the statewide vote and S_ε is the average sum of district fractional seats. Differences between the individual $S_\varepsilon(V_\varepsilon)$ and $S(V_\varepsilon)$ are another measure of the uncertainty of our estimates.

¹ The uniform shift can obviously result in districts with less than 0% or greater than 100% preferences, but that can't happen with the proportional shift.

² For example, this function estimates that a district with 55% preference for party A has a seat likelihood of 10.5% for party B as shown in Fig. 1 in Nagle 2019.

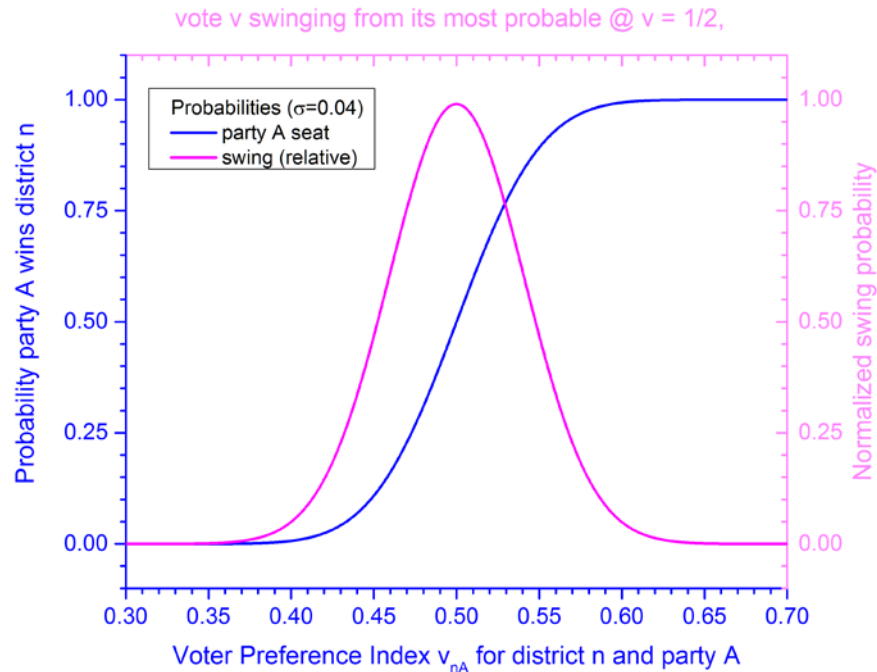
Fig. A.1 Solid circles show seats $S_\varepsilon(V_\varepsilon)$ for nine CA elections ε from whose district votes are obtained the shifted Democratic $S_\varepsilon(V)$ curves. Averaging gives the mean $S(V)$ curve shown with standard deviations.



A.3. Supporting information for use of fractional seats. (Material in this subsection has been added after publication.)

Fig. 1 is here redrawn with the same $P(V)$ as in the paper with $\sigma = 0.04$. Fig. 1 also shows its derivative $p(v) = (dP(v)/dv)/N$ where N was chosen to normalize $p(0.5) = 1$. This $p(v)$ derivative is the underlying Gaussian $\exp(-(v - 1/2)^2/2\sigma^2)$ that is the basis of the normal distribution and the probit function and much of the field of statistics.

Fig. 1. This is what is assumed in the paper.



This Gaussian is the probability $p(v)$ that the vote in an election be v votes if the mean value V equals $1/2$. In words, $p(v)$ shows the probability of a swing $v - 1/2$ in an election. When the mean vote is 0.5, then the magenta curve shows the probability for swings in the actual vote. Since the sum of the probabilities for v exceeding 0.5 is also 50%, this means that the probability


of party A winning is 0.5 and this provides the $(\frac{1}{2}, \frac{1}{2})$ point on the blue probability curve. For any other average vote v_{nA} on the blue horizontal axis, the magenta curve is shifted to $v = v_{nA}$ and the sum of all probabilities exceeding $\frac{1}{2}$ is calculated to provide the probability that party A wins the district.

There are two main assumptions here. The first is choosing the probability distribution function to be a Gaussian. The classic argument for this is that Gaussian probability is what ensues from lots of random causes that can't be predicted and that are uncorrelated with each other. This is the standard assumption in statistics and probability since Gauss 200 years ago. However, it is always possible that the distribution follows a different distribution. For example, given an average partisan preference, a district might swing further towards one party less than half the time, but the swing would be larger on average; this is skewness in the distribution of swings. Another example is that the swings are large in both directions with few near zero swings; this is kurtosis in the distribution.

Of the many standard statistical tests for rejecting Gaussian normality, the Shapiro-Wilk p-test was applied using Origin software to the 194 districts in the 11 states in this paper using all the elections in Appendix A. It was found that 25 districts were rejected for normality at the usual 5% level. However, a single outlier election can cause districts to be rejected; this was most obvious in TN where the swing for 2006 Governor was 28% above the mean. Omitting this election reduced the number of TN rejected districts from 2 to 0. The p-test also identified this election as an outlier when it was applied to all the statewide elections. This second use of the p-test was then employed in other states and the 2010 Attorney General election in CA and the 2010 Senate election in SC were also found to be outliers. Removing these elections reduced the number of rejected districts to 5 in CA, 1 in NC and 2 in TX for a total of 8 districts (4.1%) rejected for normality. Of course, 5% of random draws from a Gaussian distribution should be rejected for normality at the 5% level, so the assumption of normality is empirically supported.

The second assumption is the numerical value of the width $\sigma = 0.04$ of the Gaussian distribution. Of course, σ is not expected to be the same for each district or for each state. We can estimate σ from the second moment of the empirical distribution of the swings.³ The resulting average district σ for each state is shown in the next figure.

Fig. 2. Empirical values of average district σ by state. Average over all states is $\sigma = 0.055 \pm 0.017$. Average weighted by number of districts is $\sigma = 0.051 \pm 0.010$.



³ This is just the square root of the average squared mean deviation from the average vote.

Choosing different values of σ makes a significant difference in obtaining a fractional seat value for one district. However, when we obtain the sum of fractional seats of a party in a state, the differences in the sum is smaller because there will usually be districts leaning in both directions, although not the same number which is one reason why fractional seats should be calculated.⁴ Of course, DRA could be programmed to compute σ for each state and even for each district, although the accuracy will generally be poorer than what is shown in Fig. 2 because there are generally fewer elections in the DRA database than the 9 to 16 in Table A.1.⁵

In summary, using $\sigma = 0$ is all-or-nothing which unrealistically assumes that a district that slightly leans to a party by 1% is as likely to be won by that party as a district that favors that party by 10%.⁶ We assert that the use of $\sigma = 0$ is an assumption, and a much worse assumption than assuming normal statistics with a value of σ that is consistent with the literature and strongly supported by data.⁷

Supplementary Appendix B: Comparison of proportional and uniform methods for shifting vote share

It has long been recognized that reliable extrapolations can be made by shifting the vote that is obtained in actual elections. The simplest way to do this, much used over the years, is the uniform shift method which shifts every precinct and every district by the same percentage as the shift $V - \langle V \rangle$ in the statewide vote. While this time-honored method is likely to be reliable for competitive districts j with v_j close to 0.5, it has the obvious flaw for unbalanced districts that v_j may become greater than 1 or less than 0 (King 1989), so we employ the proportional shift method. This is based on the more plausible model that it is equally likely that any voter anywhere in the state is equally likely to shift his/her vote in the same proportion as the statewide shift. Of course, this is a model. It is quite likely that some precincts have such strong partisan preferences that hardly anyone will vote differently when the statewide vote changes and other precincts will be more responsive than average. We have done a cursory analysis that suggests that such differences tend to average out at the level of congressional districts when many elections are considered, but that is beyond the scope of this paper. The purpose of this appendix is simply to show how the $S(V)$ curves compare for the two different models.⁸

⁴ In the past results have been compared using $\sigma = 0.02$ with those from using $\sigma = 0.04$ with acceptable differences.

⁵ A suitable compromise for DRA would be to calculate σ from the statewide elections. That σ has negligibly smaller values than those shown in Fig.2.

⁶ It might be noted that using a very large value of σ gives $P(V) = \frac{1}{2}$ for all V , which unrealistically says that partisan preference makes no difference in future elections.

⁷ A criticism of our use of normality tests is that they tend to be more accurate with more than 30 data points, but we only have, depending upon the state, between 9 and 16 elections. A superficially appealing alternative to exceed 30 data points would combine the swings of all districts in a state. However, such a data combination should not be used to test district normality because there is strong correlation in the district swings for each statewide election; analyzing such a data set for normality would require a more complicated analysis with a reduced effective number of independent data points.

⁸ Katz, et al. 2020,172 have recently noted that the proportional swing has three times the error of uniform swing in their data base. The errors are smaller and more nearly equal for the two methods in our analysis. We also note that their paper used alternative (1) in Sec. A.2 for the vote fraction.

Figure B.1 shows the comparison for PA. As expected, both methods agree with each other and with the actual election results in the actual vote range. The difference for large shifts is also expected. The uniformly shifted $S(V)$ only goes to 0 when the statewide shift away from Democrats is large enough to shift the most Democratic districts to become safe Republican (this also shifts some Republican districts to have more than 100% Republican voters). In contrast, as the statewide vote shifts against Democrats using a proportional shift, the most Democratic districts lose Democratic voters more rapidly, so the Democratic proportional shift $S(V)$ curve lies below the uniformly shifted $S(V)$ curve for small V . Similarly, as the vote shifts toward Democrats it takes a larger statewide shift for the most Republican districts to become Democratic, so the uniform shift $S(V)$ lies below the proportional shift $S(V)$ for large V .

Fig. B.2 for CA shows the same differences between the uniform shift and the proportional shift as Fig. B.1. Fig. B.2 also shows $\beta(V)$ where 0.5 has been added to β in order to centralize the curves on the same figure. The proportionally shifted $\beta(V)$ reverses sign as V increases from 0.5 where it is equal to α_S and slightly favors Democrats to $\langle V \rangle = 0.6$ where it would indicate bias in favor of Republicans, contrary to all our core measures in Section 7 as shown in Table 1 in the main text. The uniformly shifted $\beta(V)$ follows the same course with V but its magnitude is much smaller at $\langle V \rangle$ than the proportional shift because $S(V)$ is larger at the counterfactual $V = 1 - \langle V \rangle$. For other states the sign is even different as shown in Table B.1. Such large differences in the results using two plausible models for shifting votes is further reason not to trust the β measure of bias. The α_S measure is clearly better as Table B.1 shows that its values have the same sign in each state for both shifting methods, although the magnitudes do differ, most in TX where votes in the actual elections require the greatest extrapolation to $V = 50\%$.

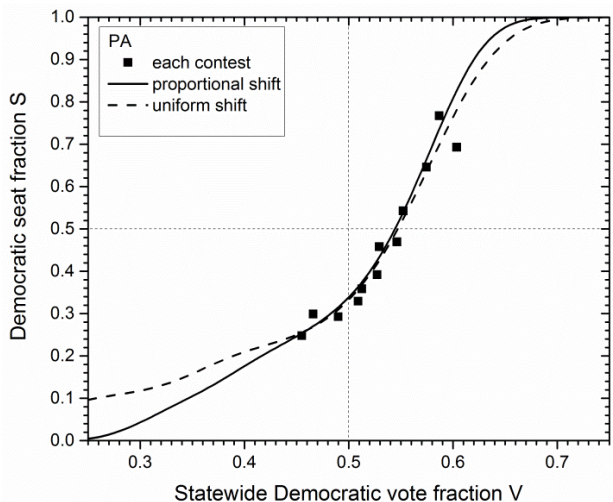


Fig. B.1 $S(V)$ for two shift models in PA

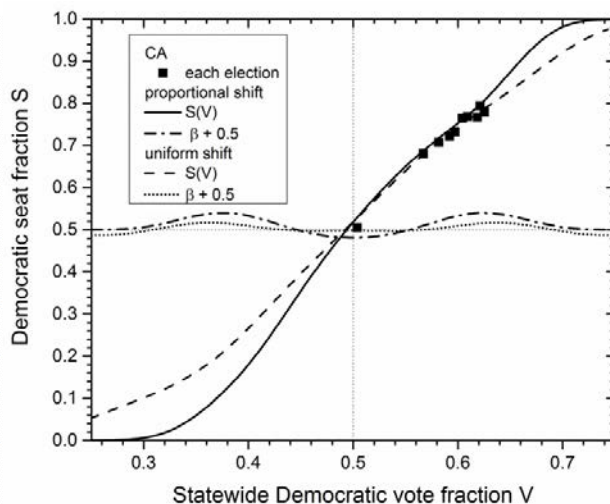


Fig. B.2. $S(V)$ for two shift models in CA

Table B.1. Percentage values for α and β for unbalanced states as determined using proportional shift and uniform shift.

state	CA	IL	MD	MA	TN	TX	SC
beta values	--	--	--	--	--	--	--
prop shift	2.9	1.2	1.4	-2.9	-2.3	-7.5	0.9
uniform shift	0.6	1.6	-2.1	-3.0	2.2	1.1	4.9
alpha values	--	--	--	--	--	--	--
prop shift	-1.9	5.8	-5.2	6.4	16.2	5.0	14.1
uniform shift	-0.2	9.3	-3.7	12.7	21.2	12.7	16.9

Appendix C: Proportionality vs. Political Geography and the SMD System

C.1 The proportionality ideal

Proportionality is widely recognized as the ideal in countries that use a list or mixed member proportional system to determine the fraction of seats awarded to different parties, but courts have not recognized it in the United States. Indeed, SCOTUS opinions have specifically denied that proportionality has constitutional support. A more cogent argument against proportionality is that it is not generally achievable over a wide range of the vote in the SMD system. A counterargument is that proportionality only has to be achievable over the expected range of the vote in each state. We will come back to achievability in the next subsection.

Let us we review why proportionality can be considered the ideal, even in an SMD system. We begin with the assertion that each voter should be represented equally as any other voter. Two voters of like mind in the same district with the same representative clearly have equal representation, but that is different from the representation of a voter of opposite mind in the same district. That is an unfortunate feature of the SMD system of representation. Nevertheless, groups of voters of like mind can still be represented equally on average in the SMD system and that is how one can argue for proportionality as the ideal.⁹ In this view representatives are shared among voters, so the empowerment of each voter of like-mind is the number of representatives of the like-minded party divided by the number of voters who voted for that party. Let us designate S and V as the fraction of seats and voters for party A and $1-S$ and $1-V$ as the fractions for the other party B. Then, the share of representation per voter of the two parties is just S/V and $(1-S)/(1-V)$ respectively. Requiring equal shares for voters in both groups and trivial algebra then requires $S = V$, which is just proportionality.¹⁰ In terms of measures of responsiveness, this is $R = \rho = 1$.

If proportionality is the democratic ideal, that suggests a simple measure of bias, namely,

$$PR = S(\langle V \rangle) - \langle V \rangle . \quad (C.1.1)$$

This evaluation is only performed at the average vote share $\langle V \rangle$ in an attempt to avoid the achievability problem of demanding proportionality over a wide V range. The PR values for the states we have analyzed are shown in Table 1¹¹.

The PR measure is related to overall responsiveness R by

$$PR = (R - 1)(\langle V \rangle - 1/2) , \quad (C.1.2)$$

⁹ Brief Amici Curiae of Common Cause and the Campaign Legal Center, Inc, *Shapiro v. McManus*, 14-990 (U.S. Aug. 13, 2015).

¹⁰ Note that one can also generalize this in two ways. First, even representatives of the opposite party provide representation for a voter of the opposite party, just not as much; although the algebra is more complicated, proportionality again ensues. Second, one can derive proportionality for any number of parties.

¹¹ Regarding the sign of PR, when Republican values are inserted for S and $\langle V \rangle$ as is done in Table 1, then positive values of PR mean bias in favor of Republicans.

which shows that PR depends both on the winner's bonus R and the state imbalance $\langle V \rangle - 1/2$.¹²

Another theoretical ideal is provided by the efficiency gap (EG) principle that the wasted votes should be equal for both parties. This results in $R = \rho = 2$. Other variations have also been discussed (Nagle 2017). Historical averages give values of R near 2. This has suggested that idealizing an $R = 2$ might take into account the effect of the SMD system. Then one has the well-known EG measure of bias (McGhee 2014),¹³

$$EG = (S(\langle V \rangle) - 1/2) - 2(\langle V \rangle - 1/2) . \quad (C.1.3)$$

Like the PR measure, EG also depends upon R and $\langle V \rangle - 1/2$ according to $EG = (R - 2)(\langle V \rangle - 1/2)$.

C.2 Political geography and a modified proportionality measure

The end of Section 4.2 essentially asks a question: Why should MD be generally acknowledged to have been gerrymandered whereas MA is not so considered when the estimated 4% fraction of Republican seats in MA is so much smaller than the 14.3% fraction in MD for the nearly identical 60% partisan imbalance in both states? This question can be reframed in terms of the quantitative measures in the previous subsection. Why should the MD plan not be considered fairer than the MA plan when MD has smaller PR bias¹⁴ than MA if the value $PR = 0$ is to be considered the ideal based on proportionality? One answer could be simply that PR may not be a good measure of bias. We next consider a more nuanced answer to this question that takes into account the contribution of political geography in the SMD system to the PR measure.

If the political complexion of an unbalanced state is homogeneous, then any map must lead to the dominant party winning all the seats. We believe that this is close to being the case for MA because we have been unable to draw even one reasonable looking Republican leaning district in MA.¹⁵ In contrast, the 2011 MD plan clearly showed that it was possible to draw a map with a strongly Republican district. While the MD plan is fairer to its minority party voters by conforming better to proportionality than MA, one clearly can't blame the makers of the MA plan for an unfair map if it was absurd to draw even one Republican district. In the case of MA, we suggest that its large value of PR bias can mostly be attributed to its relatively homogeneous political geography interacting with the SMD system.

In other unbalanced states it is certainly possible to draw some districts that are favorable to the minority party, but it still may not be possible to achieve $PR = 0$ due to their political geography. In that case, the fairest achievable plan would be the one that minimizes the absolute

¹² Eq. (C.1.2) emphasizes how PR varies with R for a state's relative vote share. However, the factor $(\langle V \rangle - 1/2)$ is crucial when considering balanced states as its small value drives R into negative, anti-majoritarian values when the state is biased as seen for NC, OH, and PA in Table 1.

¹³ Although the EG was originally expressed in terms of wasted votes, this form is equivalent when turnout bias is zero, and this is the form that properly accommodates turnout bias Cover 2018, McGhee 2017, Katz, et al. 2020.

¹⁴ MD also has smaller γ_2 bias.

¹⁵ Duchin, et al. 2019 have shown rigorously that it is impossible to draw any district that would have elected a Republican based on 5 statewide elections in the 2002-2010 cycle. They also reported that MA has become more heterogeneous, and we have used the districtr software program to draw a contiguous, although not 'reasonable looking', district that had substantial Republican preference in all of the 8 statewide election results after 2008.

value of PR given the political geography of the state. We suggest that the value of PR that is closest to zero be called PR_{SMD} because it is a bias due, not to intentional gerrymandering, but to the incompatibility of the SMD system with proportionality. For maps that do not achieve the PR_{SMD} value, that value could then be subtracted to obtain an effective plan bias

$$PR_P = PR - PR_{SMD}. \quad (C.2.1)$$

PR_P can then be compared for plans within each state and also to compare plans enacted by different states.

Let us illustrate how this proposal works for MA. For MA, Table 1 shows $PR = -36\%$ which is implausibly even larger than the Republican PR advantage in balanced, clearly biased states NC, PA and OH. Let us consider as a rough estimate made from drawing a few maps and informed by the results of Duchin, et al. 2019 that about one durable and acceptable looking Republican majority district can be drawn in MA. If so, then the maximum fraction of Republican seats at $V_R = 40\%$ is $1/9$ and $PR_{SMD} = (1/9) - 0.40 = -29\%$. Then, the modified bias is $PR_P = -36\% - (-29\%) = -7\%$. This is considerably smaller in magnitude and more realistic looking than the unmodified $PR = -36\%$.

This MA example indicates that the absolute values of PR for other unbalanced states in Table C.1 may also be too large and require similar modification to take account of political geography. Consistency then suggests that a similar modification should also apply to balanced states, so we turn next to PA where we uncover cautionary concerns. Substantial numbers of maps for PA were drawn for the recent court case which overturned the 2011 Congressional map. Those maps were drawn under the tight constraint that the number of county splits was to be minimized.¹⁶ Under those constraints, Nagle 2019 estimated that the maximum number of Democratic seats would be 7.5 for 50% Democratic two-party vote.¹⁷ One might then suggest that $PR_{SMD} = (18-7.5)/18 - 0.5 = 8\%$ using 50% vote. The problem with accepting this in Eq. (C.2.1) is that any PA plan with 58% R seats for 50% vote would be assigned zero bias ($PR_P = PR - PR_{SMD} = 0$) while strongly violating the most fundamental fairness criterion of 50% seats for 50% vote. The root cause for not being able to draw a fair PA map with more than 7.5 seats for 50% vote is simply the tight county split constraint imposed by the PA Supreme Court on acceptable maps. It was shown that loosening the county split criterion allowed a fair map to be drawn (Nagle 2019). That having been done, then the minimum absolute value of PR_{SMD} is 0 and only PA plans that satisfy the fundamental 50/50 fairness criterion have $PR_P = 0$. While this shows that consistency can be established in PA in the context of this PR_P modified measure, it emphasizes the importance of the constraints imposed on acceptable maps and concern over how that choice is made.

¹⁶ PA practice requires population deviation not to exceed *one person*; then the minimal number of county splits has to be precisely the number of districts minus one, barring some highly improbable set of county populations. (A proof is provided on the web site <http://lipid.phys.cmu.edu/nagle/Technical>.) Note also that the number of county splits is not the same as the number of split counties, *e.g.* one county split between three districts counts as two splits.

¹⁷ The estimate for the 2018 adopted PA map was 7 D seats for 50% of the vote ($\alpha = 11\%$) and 9 seats for 53.6% of the vote ($\alpha_2 = 3.6\%$) which indeed was the 2018 outcome. Cervas and Grofman 2020,9 have estimated 7.8 D seats for 50% of the vote for the adopted map.

Recapitulating what we learn from the MA and PA examples, the underlying goal of the proportionality ideal is to assign bias to a plan whenever proportionality is not achieved without regard for the source of the bias. The PR_p modification attempts to remove that part of the bias that can be attributed to the SMD system which usually brings about a winner's bonus R greater than the proportionality ideal $R = 1$; Eq. (C.1.2) shows this effect. The primary goal of the PR_p modification is to enable proportionality values for unbalanced states to appear consistent with values obtained for balanced states which do not need such a modification.¹⁸

As indicated above, determination of PR_{SMD} has to be done by drawing many maps for each state. This is beyond the scope of this paper which has instead developed the approach in the main text.

Supplementary Appendix D: Symmetry in the $r(v)$ graphs

Let us here consider a very basic definition of symmetry at the district level. When there is a district that has a stronger preference for party A than the average statewide preference, then we define district-level symmetry to require that there be another district that has the same stronger preference for party B. Let us label two such districts m and n and designate the preferences of these districts by their expected vote share for party A as v_m and v_n . Then, if the statewide vote share is $\langle V \rangle$, this pair of districts is defined to be symmetrical when $v_m - \langle V \rangle = \langle V \rangle - v_n$. Symmetry for the entire state is then achieved if every district is so paired.¹⁹ We will call this district symmetry, abbreviated DS. For our $r(v)$ graphs, $\langle V \rangle$ is the Republican vote $\langle V_R \rangle$.

The nearly linear $r(v)$ graph for CA in Fig. 9 comes quite close to exhibiting DS, so it is not an unattainable ideal. Indeed, any $r(v)$ graph whose districts fall on a straight line automatically exhibits DS. However, linearity is only a sufficient condition for DS, but it is not necessary. District symmetry only requires what is called inversion symmetry. Inversion about the point $(r=0.5, v=\langle V \rangle)$ transforms each point located at r and v on the $r(v)$ graph into a point at $r' = 1 - r$ and $v' = 2\langle V \rangle - v$.²⁰ If a graph is transformed by inversion and remains identical to what it was before inversion, then we say it has DS. Fig. D.1 shows an $r(v)$ graph that has three examples that have DS. In addition to the linear one, there are two non-linear examples that are relevant for considering responsiveness. All are centered at $r = 0.5$ and the statewide average $v = \langle V \rangle$ which we chose to be 0.6 for the examples in Fig. D.1. However, the slope is quite different for these three examples at this central point; this slope is named ρ_{sym} in the figure legend. The actual responsiveness ρ is essentially the slope at $v = 0.5$ because that is the midrange where

¹⁸ For a balanced state, Eq. (C.1.2) shows that any plausible value of the winner's bonus does not contribute significantly to PR because the $\langle V \rangle - 1/2$ factor is small. For balanced states NC, OH and PA, the actual PR bias is due to the large negative anti-majoritarian values of R . We also note that a redefinition of R to be the slope of the vector from the $V=50\%$ point on the $S(V)$ curve instead of from the (50,50) center leads one back to the α_s measure.

¹⁹ For a state with an odd number of districts, at least one district would have $v_j = \langle V \rangle$.

²⁰ The reason for choosing the inversion center at $\langle V_R \rangle$ is that this becomes the center of the $S(V)$ graph.

districts most likely change party for small statewide vote swings. The legend of Fig. D.1 shows values of ρ_{sym} and ρ , as well as the winner's bonus R .²¹

Fig. D.1. Three examples of $r(v)$ graphs that have district symmetry (DS) with the same average statewide vote $\langle V_A \rangle = 0.6$ for party A.


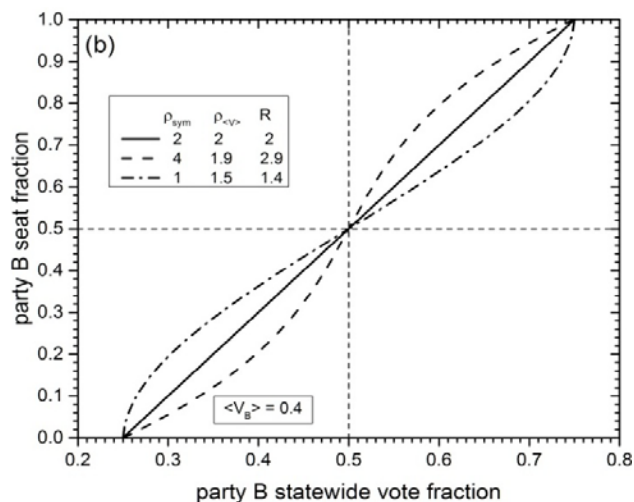


Figure D.2 shows the corresponding $S(V)$ curves which, for simplicity, are obtained by uniformly shifting the curves in Fig. D.1 by 0.1 to the left. Regarding responsiveness, it is interesting to compare ρ for the three examples in Fig. D.2. When $\langle V_B \rangle = 0.3$, the plan represented by the dot-dash line has the largest slope and the most seats for minority party B. But when $\langle V_B \rangle = 0.5$, the most responsive plan is the one represented by the dashed line in Fig. D.2 and by open circles in Fig. D.1.²²

Fig. D.2. The corresponding $S(V_B)$ curves assuming uniform swing of the $r(v)$ graphs and all-or-nothing seat assignment.



²¹ The slope at $v = 0.5$ is exactly the responsiveness if one uses the uniform swing and all-or-nothing district seat assignment. Likewise, the value of the winner's bonus R is the slope of the straight line from the intersections of the curves with the $v = 0.5$ vertical line to the central point at $r = 0.5$.

²² One can also obtain a rough measure of the number of competitive districts when $\langle V_A \rangle = 0.6$ by counting the number of districts that lie within the window 0.45 to 0.55. Then, the third plan shown by triangular symbols is least responsive, although it gives the minority party the most seats. However, when we consider $\langle V \rangle = 0.7$, that same plan is the most responsive because all the curves shift to the right by 0.1 which is the same as counting the number of responsive districts in a window which shifts to the left by 0.1 in the $r(v)$ graph. One may also define the competitiveness of individual districts using fractional seats probabilities, such as was done by Nagle 2019.

Appendix E: Estimated ideal responsiveness for unbalanced states

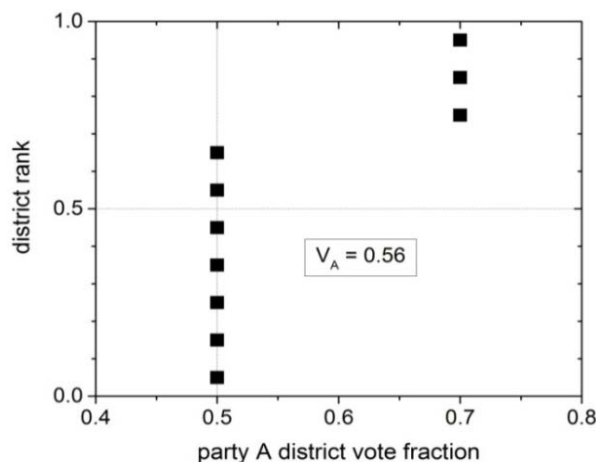
The ideal ρ_{\max} in Eq. 5.5.1 in the main text should be modified when a state is unbalanced.²³ One might then retain as many fully competitive $v_j = 1/2$ districts as possible and pack the remaining districts with as large a fraction of the majority party voters as possible. Figure E.1 shows an example of a corresponding $r(v)$ graph. Let us suppose that the maximum average packing of majority voters is v_M and that this is large enough for those districts to be completely safe. The responsiveness is then the product of the fraction S_C of competitive districts and the responsiveness of the fully competitive district with partisan preference $1/2$. When the party A vote is V_A , one finds $S_C = (v_M - V_A)/(v_M - 1/2)$.²⁴ For example, for an unbalanced state that has $V_A = 0.6$ and $v_m = 0.7$, S_C decreases to $1/2$ and the $1/\rho_{\max}$ term in Eq. (5.5.1) increases to 0.2.

It is also interesting that this model for maximal responsiveness gives the relation

$$S_A - 1/2 = (V_A - 1/2)/(2v_M - 1) \quad (\text{E5.5.2})$$

when v_M is large enough for completely safe districts. This gives winners bonus $R = 1/(2v_M - 1)$. In the unrealistic case of being able to pack districts with all of the majority party voters (*i.e.* $v_M = 1$), one has $R = 1$, like the proportional representation ideal. For a more realistic value of v_M like 0.75, then $R = 2$, like the EG ideal.

Fig. E.1 Suggested rank graph to maximize responsiveness ρ for a state with 10 districts, maximal average majority party packing $v_M = 0.7$, and average statewide vote $V_A = 0.56$.



²³ In this general case, it becomes more complicated when the fractional district seat is determined as a continuous function of v . For simplicity in this paragraph, we will simply use all-or-nothing assignment to determine the seats when the statewide vote is V_A , with $1/2$ assigned to fully competitive districts with partisan preference $1/2$.

²⁴ This follows directly from $V_A = 1/2 S_C + v_M(1 - S_C)$.

Appendix F: Rank and Seats Graphs for OH, PA, SC and IL




Fig. SAF1. $r(v)$ for Ohio

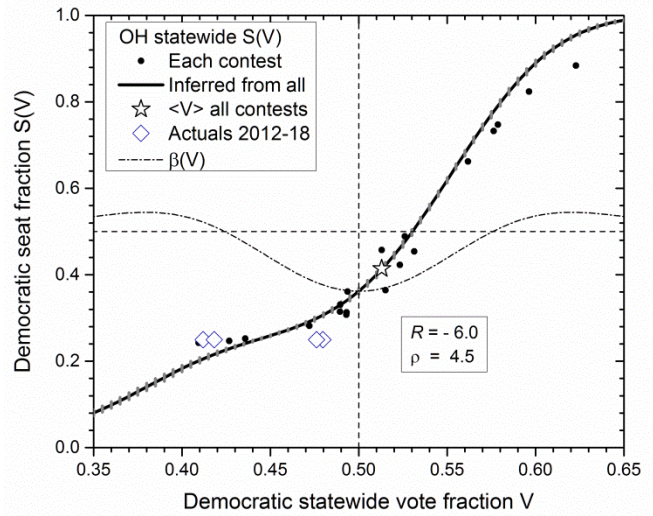


Fig. SAF2. $S(V)$ for Ohio

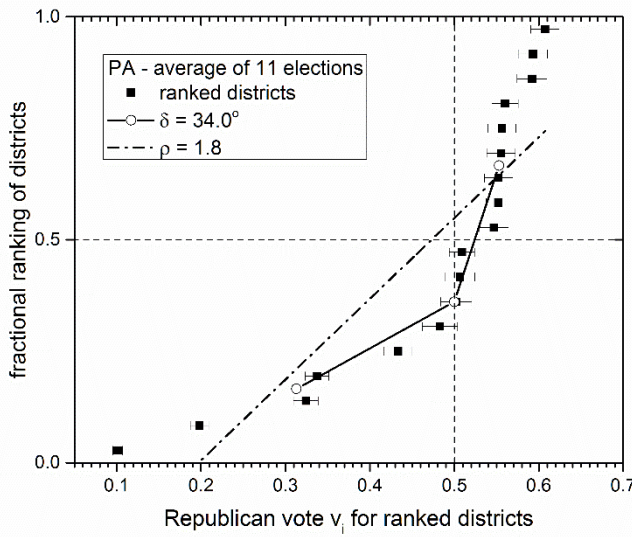


Fig. SAF3. $r(v)$ for Pennsylvania

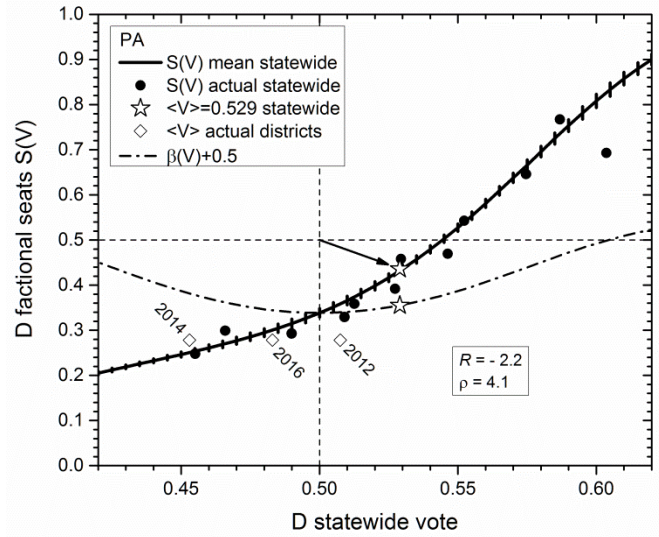


Fig. SAF4. $S(V)$ for Pennsylvania

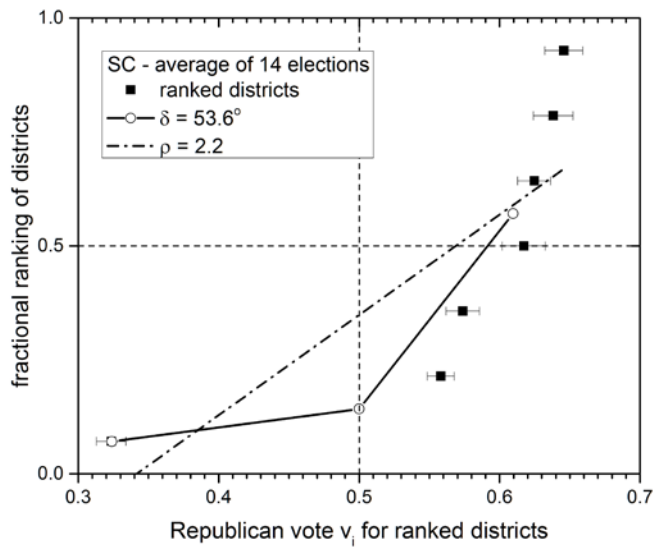


Fig. SAF5. $r(v)$ for South Carolina




Fig. SAF6. $S(V)$ for South Carolina

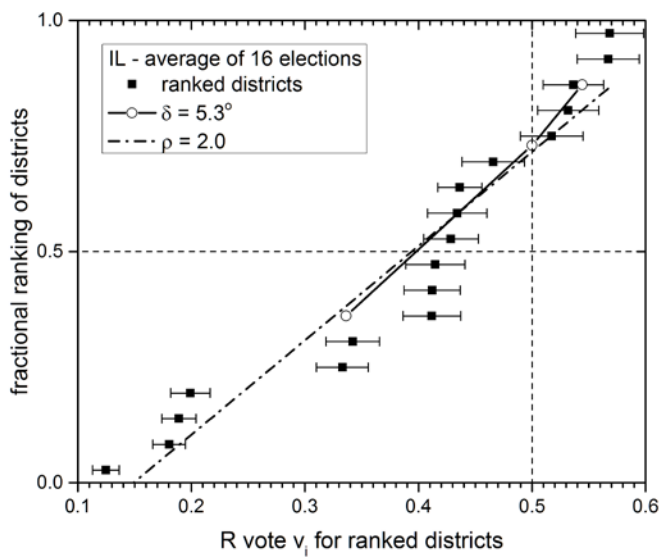


Fig. SAF7. $r(v)$ for Illinois

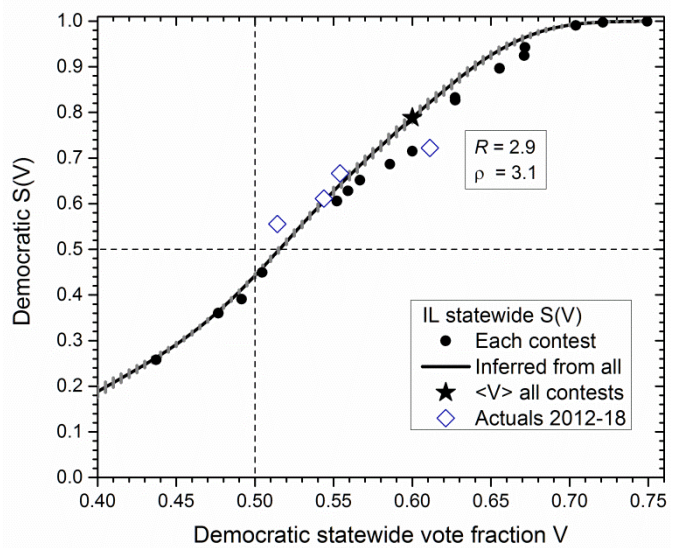


Fig. SAF8. $S(V)$ for Illinois

Appendix Bibliography

- Cervas, J. R., and B. Grofman. 2020. "Tools for Identifying Partisan Gerrymandering with an Application to Congressional Districting in Pennsylvania." *Political Geography* 76: 102069(1-12).
- Cover, B. P. 2018. "Quantifying Partisan Gerrymandering: An Evaluation of the Efficiency Gap Proposal." *Stanford Law Review* 70: 1131-233.
- Duchin, M., T. Gladkova, E. Henninger-Voss, B. Klingensmith, H. Newman, and H. Wheelen. 2019. "Locating the Representational Baseline: Republicans in Massachusetts." *Election Law Journal* 18: 388-401.
- Katz, J. N., G. King, and E. Rosenblatt. 2020. "Theoretical Foundations and Empirical Evaluations of Partisan Fairness in District-Based Democracies." *American Political Science Review* 114: 164-78.
- King, G. 1989. "Representation through Legislative Redistricting - a Stochastic-Model." *American Journal of Political Science* 33: 787-824.
- McDonald, M. D., D. B. Magleby, J. Krasno, S. J. Donahue, and R. Best. 2018. "Making a Case for Two Paths Forward in Light of Gill v. Whitford." *Election Law Journal* 17: 315-27.
- McDonald, Michael D. 2009. "The Arithmetic of Electoral Bias, with Applications to U.S. House Elections." *APSA 2009 Toronto Meeting Paper*. Available at SSRN.
- McGhee, E. 2017. "Measuring Efficiency in Redistricting." *Election Law Journal* 417: 426-31.
- . 2014. "Measuring Partisan Bias in Single-Member District Electoral Systems." *Legislative Studies Quarterly* 39: 55-85.
- Nagle, J. F. 2017. "How Competitive Should a Fair Single Member Districting Plan Be?." *Election Law Journal* 16: 196-209.
- . 2015. "Measures of Partisan Bias for Legislating Fair Elections." *Election Law Journal* 14: 346-60.
- . 2019. "What Criteria Should Be Used for Redistricting Reform?." *Election Law Journal* 18: 63-77.
- Wolf, Stephen. 2014. "Dra Update: 2012 President & Downballot Election Results Estimates + Display Table Template ". <<http://www.dailykos.com/stories/2014/8/4/1318876/-DRA-Update-2012-President-Downballot-Election-Results-Estimates-Display-Table-Template>>: Daily Kos.