Implementing partisan symmetry: Problems and paradoxes

DeFord, Dhamankar, Duchin, Gupta, McPike, Schoenbach, and Sim

December 9, 2019

Abstract

We apply elementary mathematics to the partisan symmetry standards proposed in the political science literature and derive surprising properties. We then consider prospects for implementation of partisan symmetry standards with respect to recent voting patterns in Utah, Texas, and North Carolina, finding problems and paradoxes in each case. This should raise major concerns about the suitability of partisan symmetry scores for use in redistricting reform efforts.

1 Introduction

In the political science literature, there is a long legacy of work on gerrymandering, or the act of drawing political boundary lines with an ulterior motive. One of the questions attracting the most attention has been to measure the degree of partisan advantage secured by a particular choice of redistricting lines. To counteract that requires a baseline notion of partisan fairness. The family of fairness metrics with perhaps the longest pedigree is called *partisan symmetry* scores [1, 2, 3, 6, 7, 8, 9]. These are premised on the intuitively appealing fairness notion that the share of representation awarded to one party with its share of the vote *should* also have been secured by the other party, had the vote shares been exchanged. For instance, if Republicans achieve 40% of the seats with 30% of the vote, then it would be fair for the Democrats to also achieve 40% of the seats with 30% of the vote.

At the heart of the symmetry ideal is a commitment to the principle that half of the votes should secure half of the seats. There are several particular fairness metrics or scores that derive their logic from this core axiom. The *mean-median metric* is vote-denominated: it produces a signed number that is often described as measuring how far short of half of the votes a party can fall while still securing half the seats. A similar metric, *partisan bias*, is seat-denominated. Given the same input, it is said to measure how much more than half of the seats will be secured with half of the votes. The ideal value of both of these metrics is zero. These are two in a large family of partisan metrics that can be geometrically described in terms of symmetry of the seats-votes curve.

The focus in the current note is to show that there are serious obstructions to the practical implementation of symmetry standards. This is of pressing current interest because, as we write, states are racing to adopt redistricting reform measures. In 2018 alone, four states passed constitutional amendments as voter referenda (CO, MI, MO, OH), and another wrote reforms into state law (UT). In Utah, partisan symmetry has now been adopted as a criterion to be considered by the new independent redistricting commission before plans can be approved.[1] We sound a note of caution here, showing that the versions of these scores that are realistically useable are eminently gameable by partisan actors and do not have reliable interpretations.

Utah itself gives strong evidence of the interpretation problems: with respect to recent voting patterns, a good "symmetry" score logically entails a Republican congressional sweep; what's more, the popular symmetry scores described above actually make a sign error and flag all plans with any Democratic representation as major *Republican* gerrymanders.

[1] "The Legislature and the Commission shall use judicial standards and the best available data and scientific and statistical methods, including measures of partisan symmetry, to assess whether a proposed redistricting plan abides by and conforms to the redistricting standards" that bar party favoritism [13].

1.1 Literature review

1.1.1 Building the seats-votes curve with available data

We consider an election in a state with k districts and two major parties, Party A and Party B. A standard rhetorical device in the political science literature is the "seats–votes curve," a plot representing the relationship of the vote share for Party A to the seat share for the same party. Observed outcomes are single points in V-S space—for instance, $(.3, .4)$ represents an election where Party A got 30% of the votes and 40% of the seats—but various methods have been used to extend the plot, such as fitting a curve from a given class to observed data points. We will focus on a second construction of seats-votes curves that is emphasized in [8]: beginning with a single observation, a partisan swing assumption is employed to vary the vote share. When that swing is linear, it generates a step function spanning from $(0, 0)$ to $(1, 1)$. (See Figures 1-2 below for examples.) This linear uniform partisan swing is the leading method proposed for use in evaluation. Katz–King–Rosenblatt note that curve-fitting is more suited "for academic study... than for practical use" in evaluation of plans. Grofman noted in 1983 that linear swing is preferred in practice to more sophisticated models [6, n.14], and it has been invoked even this year (2019) in expert reports and testimony [11].

1.1.2 Deriving symmetry scores from the seats-votes curve

Given a seats-votes curve, many scores have been discussed to evaluate its symmetry; here, we will focus on the mean-median score MM, the partisan bias score PB, and the partisan Gini score PG, all of these which have been considered for at least 35 years. (Definitions are found in the next section.) Grofman's 1983 survey paper [6] lays out eight possible scores of asymmetry once a seats-votes curve has been set. His Measure 3 is vote-denominated bias, which would equal MM under linear uniform partisan swing; similarly his Measure 4 corresponds to PB, and Measure 7 introduces PG. Because the partisan Gini is defined as the area between the seats-votes curve and its reflection over the center (seen in the shaded regions in Figs 1-2), it is easily seen to "control" all the other possible symmetry scores: when PG = 0, its ideal value, all partisan symmetry metrics also take their ideal values, including MM and PB. This agrees with Katz–King–Rosenblatt [8, Def 1], where the coincidence of the curve and its reflection, i.e., PG = 0, is called the "partisan symmetry standard." In this note, our Theorem 3 gives precise necessary and sufficient conditions for the partisan symmetry standard to obtain.

The literature invoking MM and PB as measures of bias is too large to survey comprehensively. We note that the particular interpretation of median minus mean as the Party A advantage is fairly standard in the journal literature, such as: "The median is 53 and the mean is 55; thus, the bias runs two points against Party A (i.e., $53 - 55 = -2$)" [10]. Also standard is the connection to the seats-votes curve: MM "essentially slices the S/V graph horizontally at the S = 50% level and obtains the deviation of the vote from 50%" [12, p351].[2]

1.1.3 Applying symmetry scores in practice

Numerous scholars have advanced these scores and standards for courts as practical tests of partisan gerrymandering, in amicus briefs spanning from LULAC v. Perry (2007) [1] to Whitford v. Gill (2018) [2] to the past term's Rucho v. Common Cause (2019) [3]. They have been claimed to be "reliable and difficult to manipulate" and authors have argued that while "Symmetry tests should deploy actual election outcomes" (as we do here), they will nonetheless "measure opportunity," i.e., give information about future performance [2, p17,24]. That brief explicitly proposes mean-median as a concrete choice of score for this task.

As laid out in the influential LULAC brief, "Models applying the symmetry standard are by their nature predictive, just as the legislators themselves are predicting the potential impact of the map on partisan representation. The symmetry standard and the resulting measures of partisan bias, however, do not require forecasts of a particular voting outcome. Rather, by examining all the relevant data and the potential seat divisions that would occur for particular vote divisions, it compares the potential scenarios and determines the partisan bias of a map, separating out other potentially confounding factors. Importantly, those drawing the map have access to the same data used to evaluate it, and no data is required other than what is in the public domain" [1, p11]. This paper takes up precisely this modeling task in the manner explicitly proposed by its authors.

[2] There is even more work centered on PB (notably [9]), but it is more rarely used in conjunction with linear swing, since that assumption makes its values move in large jumps.

2 A mathematical characterization of the Partisan Symmetry Standard

We begin with definitions and notation needed to state the results in this paper, and particularly Theorem 3, which characterizes when $\mathsf{PG} = 0$. We describe the vote outcome in the election using a vector of vote shares, an ordered vector composed of Party A share of the total votes for the Parties A and B in each of the k districts as follows: $\mathsf{v} = (v_1, \ldots, v_k)$, where $0 \le v_1 \le \ldots \le v_k \le 1$. The number of districts in which Party A has more votes than Party B in an election with vote shares v is $\#\{i : v_i > \frac{1}{2}\}$. This induces a function $\gamma = \gamma_{\mathsf{v}}$ by $\gamma(\overline{v} + t) = \#\{i : v_i + t > \frac{1}{2}\}/k$, which we can interpret as the share of districts won by Party A in the counterfactual that an amount t was added to A's observed vote share in every district. Varying t to range over the one-parameter family of vote vectors

$$(v_1 + t, \ldots, v_k + t)$$

is known as (linear) *uniform partisan swing*. The curve γ, treated as a function $[0, 1] \to [0, 1]$, has been regarded as measuring how a fixed districting plan would behave if the level of vote for Party A were to swing up or down. Below, we will refer to function and its graph interchangeably, and we will call it the *seats-votes curve* associated to the vote share vector v.

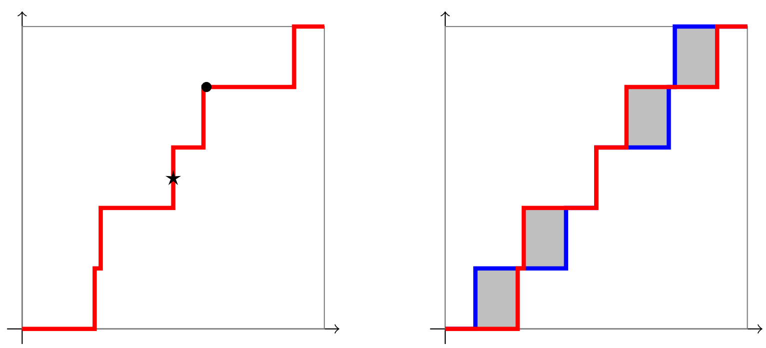

Figure 1: Red: The seats-votes curve γ generated by the vote share vector $\mathsf{v} = (.21, .51, .61, .85, .87)$, which gives $\overline{v} = .61$. Blue: its reflection about the center point $\star$. Since MM is the horizontal displacement from $\star$ to a point on γ, this hypothetical election has a perfect $\mathsf{MM} = 0$ score, but it is not very symmetric overall, with $\mathsf{PG} = .112$, seen as the area of the shaded region between γ and its reflection.

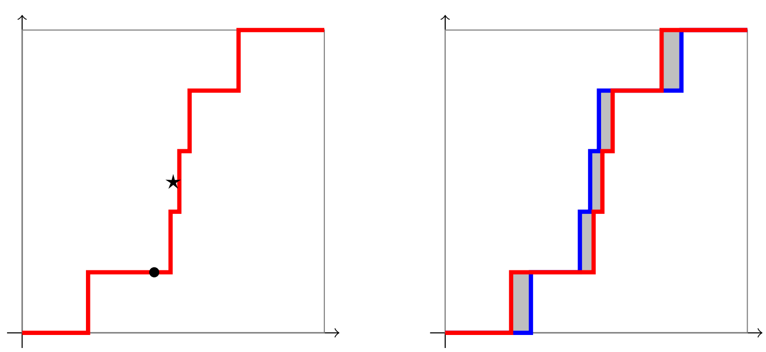

Figure 2: The seats-votes curve γ generated by the vote share vector $\mathsf{v} = (.221, .383, .417, .446, .719)$, which was the observed outcome in the 2016 Congressional races in Oregon from the Republican point of view. This gives a mean of $\overline{v} = 0.4372$, and earned Republicans 1 seat out of 5. The red curve shows Republican seats at each vote share under linear uniform partisan swing. The blue curve is the reflection about the center, so it shows seats at each vote share from the Democratic point of view. This could be regarded as a situation with reasonably good symmetry, since the red and blue curves are close. Its scores are $\mathsf{PG} = .05248$, $\mathsf{MM} = -.0202$, and $\mathsf{PB} = -.1$. The sign of the latter two scores is thought to indicate a Democratic advantage.

Let the mean district vote share for Party A be denoted $\overline{v} = \frac{1}{k}\sum v_i$ and the median district vote share, v_{med}, be the median of the $\{v_i\}$, which equals $v_{(k+1)/2}$ if k is odd and $\frac{1}{2}(v_{k/2} + v_{(k/2)+1})$ if k is even because of the convention that we write them in non-decreasing order. We note that $\overline{v}$ is not necessarily the same as the statewide share for Party A except in the idealized scenario that the districts have equal numbers of votes cast. We begin with several scores based on v and γ.

Definition 1. *The* **partisan Gini** *score* $\mathsf{PG}(\mathsf{v})$ *is the area between the seats-votes curve* γ_{v} *and its reflection over the center point* $\star = (\frac{1}{2}, \frac{1}{2})$*.*

$$\mathsf{PG}(\mathsf{v}) = \int_0^1 \left|\gamma(x) - \gamma(1-x) + 1\right| dx.$$

The **mean-median** *score is* $\mathsf{MM}(\mathsf{v}) = v_{\mathrm{med}} - \overline{v}$*. The* **partisan bias** *score is* $\mathsf{PB}(\mathsf{v}) = \gamma(\frac{1}{2}) - \frac{1}{2}$*.*

These scores can all be related to the shape of the seats-votes curve γ. Partisan Gini measures the failure of γ to be symmetric about the center point $\star = (\frac{1}{2}, \frac{1}{2})$, in the sense that it is zero if and only if γ equals its reflection, and non-negative by construction. Mean-median score is the horizontal displacement from $\star$ to a point on γ,[3] which is why it is votes-denominated (vote-share being the variable on the x-axis). Similarly, partisan bias is the vertical displacement from $\star$ to a point on γ, and is therefore seats-denominated. (We note that $(\frac{1}{2}, \gamma(\frac{1}{2}))$ is a well-defined point unless there is a jump precisely at $1/2$, which occurs if some $v_i = \overline{v}$ on the nose—this should not happen with real-world data.) Below, we will focus on MM instead of PB, but we note that $\mathsf{MM} > 0 \implies \mathsf{PB} \geq 0$ because of the geometric interpretation: if γ passes to the left of $\star$ and is nondecreasing, then it must pass through or above $\star$.

We can see that the curve γ, and consequently the partisan Gini score, is exactly characterized by the points at which Party A has added enough vote share to secure the majority in an additional district. For the following analysis, it will be useful to characterize this curve in terms of the v data.

Definition 2. *The* **gaps** *in a vote share vector* v *can be written in a gap vector*

$$\updelta = (\delta_1, \delta_2, \ldots, \delta_{k-1}) = (v_2 - v_1,\ \ v_3 - v_2,\ \ \ldots,\ \ v_k - v_{k-1}).$$

The **jump points** *for vote share vector* v *are the values of* $\overline{v} + t$ *such that some* $v_i + t = \frac{1}{2}$*. We have*

$$t_1 := \frac{1}{2} - v_k\ , \qquad t_2 := \frac{1}{2} - v_{k-1}\ , \qquad \ldots \qquad t_k := \frac{1}{2} - v_1$$

as the times corresponding to these jumps, so we can denote the jumps as $j_i = \frac{1}{2} + \overline{v} - v_{k+1-i}$*, and the jump vector as* $\mathsf{j} = (j_1, \ldots, j_k)$*.*

By definition of γ, these jump points j_i are exactly the x-axis values at which γ jumps from $(i-1)/k$ to i/k.

With this notation, we can manipulate the various partisan symmetry scores. For instance, the center-most rectangle(s) formed between γ and its reflection have height $2\mathsf{PB}$ and width $2\mathsf{MM}$, allowing the derivation of inequalities relating these scores. For small k, these reduce to extremely simple expressions: $\mathsf{PG} = \frac{4}{3}|\mathsf{MM}|$ when $k = 3$, and $\mathsf{PG} = 2|\mathsf{MM}|$ when $k = 4$. (See Supplement B.)

For any number of districts, we obtain a very clean characterization of precisely which vote shares satisfy the Partisan Symmetry Standard [8, Def 1].

Theorem 3 (Partisan Symmetry Characterization)**.** *Given k districts with vote shares* v*, jump vector* j*, and gap vector* $\updelta$*, the following are equivalent:*

$$\mathsf{PG}(\mathsf{v}) = 0 \tag{1}$$

$$j_i + j_{k+1-i} - 1 = 0 \qquad \forall i \tag{2}$$

$$\frac{1}{2}\left(v_i + v_{k+1-i}\right) = \overline{v} \qquad \forall i \tag{3}$$

$$\frac{1}{2}\left(v_i + v_{k+1-i}\right) = v_{\mathrm{med}} \qquad \forall i \tag{4}$$

$$\delta_i = \delta_{k-i} \qquad \forall i \tag{5}$$

That is, the partisan symmetry standard is nothing but the requirement that the vote shares by district are distributed on the number line in a symmetric way. In particular, this tells you at a glance that an election with vote shares $(.37, .47, .57, .67)$ in its districts rates as perfectly partisan-symmetric, while one with vote shares $(.37, .47, .57, .60)$ falls short. The proof is included in Supplement A.

[3] To see this, plug in $t = 1/2 - \mathsf{MM} - \overline{v} = 1/2 - v_{\mathrm{med}}$ to deduce that $(1/2 - \mathsf{MM}, 1/2)$ is on γ.

3 Paradoxes with signed symmetry scores

Recall that mean-median and partisan bias are *signed* scores that are supposed to identify which party has an advantage and by what amount. A positive score is said to indicate an advantage for Party A (the point-of-view party whose vote shares are reported in v). Let us say that a *paradox* occurs when the score indicates an advantage for one party even though it has a very low number of seats—the fewest seats it can possibly earn for a given vote pattern, say. In other words, a paradox means that the score makes an apparent sign error.

To illustrate this, we will begin with the case of $k = 4$ districts, where the algebra is simpler. The issues are not limited to small k, however: in the empirical section we will find paradoxes of this kind in $k = 13$ and $k = 36$ cases as well.

Example 4 (Paradoxes forced by arithmetic)**.** *Suppose we have $k = 4$ districts and an extremely skewed election in favor of Party A, achieving $75\% \le \overline{v} < 87.5\%$. With equal turnout, Party B can get at most one seat. However, every partition of votes achieving this value yields $\mathsf{MM} \ge \overline{v} - \frac{3}{4}$. In particular, districting plans with any representation for Party B all have positive MM and PB, paradoxically indicating an advantage for Party A.*

The demonstration is extremely simple as a matter of arithmetic. Since $\frac{1}{2}(v_2 + v_3) = v_{\text{med}}$, we have

$$\overline{v} = \frac{1}{4}(v_1 + v_2 + v_3 + v_4) = \frac{v_1 + v_4}{4} + \frac{v_2 + v_3}{4} \implies v_{\text{med}} - \overline{v} = \overline{v} - \frac{v_1 + v_4}{2}.$$

Since $v_1 \le \frac{1}{2}$ (for B to win a seat) and $v_4 \le 1$, we get $\mathsf{MM} = v_{\text{med}} - \overline{v} \ge \overline{v} - 3/4$, as needed.

A stronger statement can be made if one takes political geography into account. It was shown by Duchin et al in a study of Massachusetts [5] that, given a set of atoms that are not to be split in redistricting, the uniformity of vote distributions can prevent the minority party from achieving representation. By the same token, uniformity can impose upper bounds on the level of packing that is possible in a district.

Example 5 (Paradoxes forced by geography)**.** *Suppose we have $k = 4$ districts and a skewed election in favor of Party A, with $62.5\% \le \overline{v} < 75\%$. Suppose the geography of the election has Party A support arranged uniformly enough that Party A vote share in a district can not exceed Q, for some $Q < 2\overline{v} - \frac{1}{2}$. Then with equal turnout, Party B can get at most one seat. However, every partition of votes achieving this value yields $\mathsf{MM} > \overline{v} - \frac{1}{4} - \frac{Q}{2}$. In particular, districting plans with any representation for Party B all have positive MM and PB, paradoxically indicating an advantage for Party A.*

That is, when there is an extremely skewed outcome (with a vote share for one party exceeding 75%), paradoxes always occur, just for arithmetic reasons. But even for less skewed elections with a vote share between 62.5 and 75% for the leading party—which frequently occurs in practice—realistic constraints imposed by political geography can force these sign paradoxes.

Proof of Q bound. First, it is easy to see that Party B can't achieve two seats: in that case, we would have $v_1, v_2 \le \frac{1}{2}$. Since we also have $v_3, v_4 \le Q < 2\overline{v} - \frac{1}{2}$, we can average the v_i to get the contradiction $\overline{v} < \overline{v}$.

To see that $\mathsf{MM} > 0$, we write

$$\mathsf{MM} = v_{\text{med}} - \overline{v} = \frac{v_2 + v_3}{2} - \frac{v_1 + v_2 + v_3 + v_4}{4} = \frac{v_1 + v_2 + v_3 + v_4}{4} - \frac{v_1 + v_4}{2} = \overline{v} - \frac{v_1 + v_4}{2}.$$

Since $v_1 < \frac{1}{2}$ and $v_4 \le Q < 2\overline{v} - \frac{1}{2}$, we have $\mathsf{MM} > \overline{v} - \frac{2Q+1}{4} > \overline{v} - \overline{v} = 0$. □

4 Investigations with observed vote data

In this section we illustrate the issues that were raised on a theoretical level above, using naturalistically observed election data together with a Markov chain technique that produces large ensembles of districting plans. All data and code are public and freely available for inspection and replication [14].

The algorithmically generated plans are not offered as a statistical experiment and come with no probabilistic claims about plan selection, but merely provide *existence proofs* to illustrate how readily gameable partisan symmetry standards will be for those engaged in redistricting. The methods also produce many thousands of examples of plans that are paradoxical in the senses described above, where partisan symmetry metrics identify the wrong party as the gerrymanderers, relative to the common understanding of that term.

In each case, we have run a recombination (ReCom) Markov chain for 100,000 steps—long enough to comfortably achieve heuristic convergence benchmarks in all scores that we measured—while enforcing population balance, contiguity, and compactness.[4] We have run trials on every election in our dataset and all results are available for comparison [14]. Below, we have highlighted the most recent available Senate race from a Presidential election year in each state.

4.1 Utah and the "Utah Paradox"

We begin with Utah, where the elections that were available in our dataset all come from 2016.[5] Utah has only four congressional districts and has a heavily skewed partisan vote share, with a statewide Republican vote share of 71.55% in the 2016 Senate race.[6] Figure 3 shows outcomes from our 100,000-step ensemble.

The vast majority (94.266%) of Utah plans found in our ensemble have all four R seats, with the remaining plans giving 3-1 splits. The chain found 2772 plans with 3 Republican and 1 Democratic seats, and we see that all of these have PG scores above 0.06. Below, we rigorously establish these bounds on seats and scores.

When looking at the full PG histogram, we see a large bulk of plans with nearly-ideal PG scores, all giving a Republican sweep (four out of four R seats). This is surprising enough to deserve a name.

The Utah Paradox

- Partisan symmetry scores near zero are supposed to indicate fairness, and signed symmetry scores are supposed to indicate which party is advantaged.
- There are many trillions of valid Congressional plans in Utah, and every single one of them with PG close to zero has a Republican sweep of the seats. In particular, even constraining symmetry scores to be in the better half of the observed values would deterministically impose a partisan outcome: the one in which Democrats are locked out of representation.
- Furthermore, the signed scores make an apparent sign error: they report all plans with Democratic representation to be significant pro-Republican gerrymanders.

Applying sorting bounds to UT-SEN16, we see that the best possible Republican quasi-district is $Q = 86\%$ Republican, achieved by picking precincts from all over the state with no regard to contiguity. Certainly any legally valid district built out of precincts is therefore at most 86% Republican.

Example 6 (The Utah paradox, quantified)**.** *The UT-SEN16 vote pattern can be divided into 4R-0D seats or 3R-1D seats. However, even though* MM, PB, *and* PG *scores can all get arbitrarily close to zero, there are no reasonably symmetric plans that secure a Democratic seat. In our algorithmic search, every plan with Democratic representation has* PG $> .069$, MM $> .034$, *and* PB $\geq .25$, *which is in the worst half of the scores observed. Thus any reasonable constraint on partisan symmetry would lock Democrats out of representation, and all plans with D representation are paradoxically reported as egregious R-favoring gerrymanders.*

Deriving quantified paradox. The UT-SEN16 election has $\overline{v} = .7155$. If we can upper-bound the possible Republican share of a district by any $Q < .931$, then the arguments of the last section show that Democrats can secure at most one seat, and that every plan with Democratic representation has MM > 0. Indeed, even a greedy assemblage of the 608 precincts with the highest Republican share in that race (which is the number needed to reach the ideal population of a district) only produces a district with R share of .888, which ensures MM $> \overline{v} + \frac{2Q+1}{4} \approx .0215$. And this is even

[4] The population balance imposed here is 2% deviation from ideal district size. Such plans are easily tuneable to 1-person deviation by refinement at the block level without significant impact to any other scores discussed here. Contiguity is enforced by recording adjacency of precincts. Compactness, at levels comparable to those observed in human-made plans, is an automatic consequence of the spanning-tree-based recombination step. For more information about the Markov chain used here, see [4].

[5] Out of GOV16, SEN16, and PRES16, none gives a stellar partisanship signal, because the Democratic candidates for Governor and Senate were quite weak, while the partisanship in the Presidential race was complicated by the presence of a very strong third-party candidate in Evan McMullen. This is a frequent issue with this method, though, since most elections have idiosyncratic stories.

[6] We note in passing that the amount of linear partisan swing needed to reverse the partisan advantage could be viewed as unreasonably large. With respect to SEN16, fully 199 out of 2123 precincts in the state have Republican vote share pushed to less than zero by the swing.

without the requirement of contiguous districts, which certainly limits the possibilities further. In our Markov chain run of contiguous plans, the highest Republican share ever encountered in a single district is .8595, which ensures $\mathsf{MM} > .035$. From Supplement B, we know that $\mathsf{PG} = 2|\mathsf{MM}|$, so this gives $\mathsf{PG} > .07$. And finally, since $j_2 = \frac{1}{2} + \overline{v} - v_3$ and $v_3 > \overline{v}$, we have $j_2 < \frac{1}{2}$, which means that $\gamma(1/2) > 1/2$. Since the seats values rise in increments of $1/4$, we have $\mathsf{PG} \geq 1/4$. □

As described in the introduction, Utah recently became the first state to encode partisan symmetry as a districting criterion in statute. This makes the Utah Paradox quite a striking example of the worries raised by using partisan symmetry scores in practice.

4.2 Texas

Next, we turn to Texas, creating a chain of 100,000 steps to explore the ways to divide up the 2012 Senate vote distribution. With 36 Congressional districts, Texas has one of the highest k values of any state (only California has more seats). The 2012 Senate race was won by a Republican with $\sim 58\%$ of the vote.

Figure 4 shows the partisan properties in the ensemble of plans, allowing us to compare extreme symmetry scores (the ostensible indicator of partisan gerrymandering) to extreme seat shares (the explicit goal of partisan gerrymandering). We find no evidence of a useful correspondence.

Over 98% of all plans give Republicans 22 to 27 seats out of 36, seen in gray in the histogram. The red bars mark the outlying plans with the most Republican seats (28 or more), while the blue bars mark the outlying plans with the most Democratic seats (21 or fewer). We can then study the histograms formed by the winnowed subsets of the ensemble with the best PG and MM scores, always in the top 5%. Note that these severely winnowed subsets not only have a shape similar to the full ensemble, but that the partisan outlier plans are not proscribed by strict symmetry thresholds. In fact, plans with the vanishingly extreme outcome of 30/36 Republican seats occur with higher frequency among the plans with $\mathsf{MM} \approx 0$ than among the full set of sampled plans.

Next we present the reverse perspective, considering how the plans with the most extreme seat shares score on partisan symmetry. Below those are histograms made only from the seat outliers: blue for plans with ≤ 21 Republican seats and red for plans with ≥ 28 Republican seats. The mean-median score does not succeed in signaling extreme seat totals, and in fact a significant number of extreme D-favoring plans paradoxically register as major Republican gerrymanders. In terms of the overall symmetry measured by PG, extreme plans for both parties can be found with nearly-optimal scores.

4.3 North Carolina

Finally, we move to a state with a much closer to even partisan split: North Carolina ($k = 13$ seats), with respect to the 2016 Senate vote ($\sim 53\%$ Republican share).

In this case, mean-median does much better than in Texas in terms of distinguishing the seat extremes: we see consistently higher scores for the maps with the most Republican seat share than the ones with the most Democratic outcomes. However, partisan symmetry still imposes no constraint on partisan gerrymandering: hundreds of maps have 10-3 outcomes (which was clearly reported in the Rucho case to be the most extreme that the legislature thought was possible) while securing nearly perfect symmetry scores. Indeed, the ensemble finds many maps that return an 11-2 outcome.

5 Conclusion

In this note, we have characterized the partisan symmetry standard as a prescription for the arrangement of vote totals across districts (Theorem 3). We follow this with examples of realistic conditions in which the adoption of strict symmetry standards (a) does not constrain extreme partisan gerrymandering and (b) can enforce unforeseen consequences on partisan outcomes. Finally, again in realistic conditions, standard partisan symmetry metrics (c) can give answers that depend unpredictably on which vote pattern is used to assess them, and even (d) can plainly mis-identify which party is advantaged by a plan.

None of these findings gives a theoretical reason for rejecting partisan symmetry as a definition of fairness. A believer in symmetry-as-fairness can certainly coherently hold that symmetry standards do not *aim* to constrain partisan outcomes, but merely to reinforce the legitimacy of district-based democracy by reassuring the voting public that the tables can yet turn in the future.

Nevertheless, these investigations should serve as a strong caution regarding the use of partisan symmetry metrics, whether in redistricting criteria before maps are adopted or in evaluation after subsequent elections have been conducted.

References

[1] Amicus brief of Gary King, Bernard Grofman, Andrew Gelman, and Jonathan N. Katz in *LULAC* (Nos. 05-204, 05-254, 05-276, 05-439)

[2] Amicus brief of Heather K. Gerken, Jonathan N. Katz, Gary King, Larry J. Sabato, and Samuel S.-H. Wang in *Whitford* (No. 16-1161)

[3] Amicus brief of Bernard Grofman and Ronald Keith Gaddie in *Rucho* (Nos. 18-422, 18-726)

[4] Daryl DeFord, Moon Duchin, and Justin Solomon, *Recombination: A family of Markov chains for redistricting*. https://arxiv.org/abs/1911.05725

[5] Moon Duchin, Taissa Gladkova, Eugene Henninger-Voss, Ben Klingensmith, Heather Newman, and Hannah Wheelen, *Locating the Representational Baseline: Republicans in Massachusetts*. Election Law Journal, to appear December 2019.

[6] Bernard Grofman, *Measures of Bias and Proportionality in Seats-Votes Relationships*, Political Methodology, Vol. 9, No. 3 (1983), 295–327.

[7] Bernard Grofman and Gary King, *The Future of Partisan Symmetry as a Judicial Test for Partisan Gerrymandering after LULAC v. Perry*, 6 Election L.J. 2 (2007) .

[8] Jonathan N Katz, Gary King, and Elizabeth Rosenblatt. *Theoretical foundations and empirical evaluations of partisan fairness in district-based democracies*. American Political Science Review, to appear 2019.

[9] Gary King and Robert X. Browning, *Democratic Representation and Partisan Bias in Congressional Elections*. The American Political Science Review, Vol. 81, No. 4 (Dec., 1987), 1251–1273.

[10] Michael D. McDonald and Robin E. Best, *Unfair Partisan Gerrymanders in Politics and Law: A Diagnostic Applied to Six Cases*. Election Law Journal, Volume 14, Number 4 (2015), 312–330.

[11] Jonathan Mattingly, Expert Report in Common Cause v. Lewis. https://sites.duke.edu/quantifyinggerrymandering/files/2019/09/Report.pdf See also https://sites.duke.edu/quantifyinggerrymandering/2019/06/26/the-animated-firewall/

[12] John F. Nagle, *Measures of Partisan Bias for Legislating Fair Elections*. Election Law Journal, Volume 14, Number 4 (2015), 346–360.

[13] Utah Code. Chapter 7 Title 20A Chapter 19 Part 1 §103. https://le.utah.gov/xcode/Title20A/Chapter19/C20A-19_2018110620181201.pdf

[14] Voting Rights Data Institute, *Partisan symmetry data repository*. https://github.com/vrdi/Partisan_Symmetry

A Equivalent conditions for $\mathsf{PG} = 0$

We briefly recall the needed notation from §2 above: vote share vector v with ith coordinate v_i and $\overline{v}$ defined as the average of the v_i; gap vector δ with $\delta_i = v_{i+1} - v_i$ and jump vector j with $j_i = \frac{1}{2} + \overline{v} - v_{k+1-i}$. These expressions defined j, δ in terms of v; neither j nor δ completely determines v, because they are invariant under translation of the entries of v, but one additional datum (such as v_1 or $\overline{v}$) suffices, with j or δ, to fix v. In this section, we begin by expressing PG in terms of the jumps j, then give equivalent conditions for $\mathsf{PG} = 0$ in terms of j, δ, and v.

As outlined above, PG measures the area between the seats-votes curve γ and its reflection. The shape of the region between those curves depends directly on the points $\mathsf{j} = (j_1, j_2, \ldots, j_k)$, since each j_i is the x value of a vertical jump in the curve and the $1 - j_i$ are the values of the jumps in the reflection. But looking at Figure 1 makes it clear that it is complicated to decompose the integral into vertical rectangles in the style of a Riemann sum, because the $\{j_i\}$ and the $\{1 - j_{k-i}\}$ do not always alternate. Fortunately, it is always easy to decompose the picture into horizontal rectangles (analogous to a Lebesgue integral), where it is now clear which red and blue corners to pair as the seat share changes from i/k to $(i+1)/k$. The curve contains the points $(j_i, \frac{i-1}{k})$, $(j_i, \frac{i}{k})$ as well as $(j_{k+1-i}, \frac{k-i}{k})$, $(j_{k+1-i}, \frac{k-i+1}{k})$. The rotated curve therefore contains the points $(1 - j_{k+1-i}, \frac{i-1}{k})$ and $(1 - j_{k+1-i}, \frac{i}{k})$, which means that the ith rectangle has height $1/k$ and width $\left|j_i + j_{k+1-i} - 1\right|$. Summing over these rectangles gives us the expression

$$\mathsf{PG} = \frac{1}{k} \sum_{i=1}^{k} \left|j_i + j_{k+1-i} - 1\right|.$$

Recall that the set of vote share vectors $\mathcal{V}$ is the cone in the vector space $\mathbb{R}^k$ given by the condition that the v_i are in non-decreasing order in $[0,1]$. The j vector is simply the v vector reversed and re-centered at $1/2$ rather than $\overline{v}$. The only condition on the gap vector δ is that its entries sum to less than one. Putting these observations together we may define the set of achievable $\mathsf{v}, \delta, \mathsf{j}$ respectively as

$$\mathcal{V} = \{(v_1, \ldots, v_k) : 0 \le v_1 \le \cdots \le v_k \le 1\},$$
$$\mathcal{D} = \left\{(\delta_1, \ldots, \delta_{k-1}) : \delta_i \ge 0 \; \forall i, \quad \sum \delta_i \le 1\right\},$$
$$\mathcal{J} = \left\{(j_1, \ldots, j_k) : 0 \le j_1 \le \cdots \le j_k \le 1, \quad \sum j_i = \frac{k}{2}\right\}.$$

The condition on j is of interest because it directly parametrizes the set of all seats-votes curves $\Gamma = \{\gamma_{\mathsf{v}} : \mathsf{v} \in \mathcal{V}\}$.

Theorem 7. *Given k districts with vote shares v, jump vector j, and gap vector δ, the following are equivalent:*

$$\mathsf{PG}(\mathsf{v}) = 0 \tag{1}$$
$$j_i + j_{k+1-i} - 1 = 0 \qquad \forall i \tag{2}$$
$$\frac{1}{2}\left(v_i + v_{k+1-i}\right) = \overline{v} \qquad \forall i \tag{3}$$
$$\frac{1}{2}\left(v_i + v_{k+1-i}\right) = v_{\mathrm{med}} \qquad \forall i \tag{4}$$
$$\delta_i = \delta_{k-i} \qquad \forall i \tag{5}$$

Proof. The condition that $\mathsf{PG}(\mathsf{v}) = 0$ has been rewritten in terms of j, and converting back to the v_i we get

$$\frac{1}{k} \sum_{i=1}^{k} |j_i + j_{k+1-i} - 1| = \frac{2}{k} \sum_{i=1}^{k} \left|\frac{v_i + v_{k+1-i}}{2} - \overline{v}\right| = 0,$$

which immediately gives (1) $\iff$ (2) $\iff$ (3) since a sum of nonnegative terms is zero if and only if each term is zero. To see (3) $\iff$ (4), just consider $i = \lceil \frac{k}{2} \rceil$ in (3) to obtain $v_{\mathrm{med}} = \overline{v}$; in the other direction, average both sides over i in (4) to obtain $\overline{v} = v_{\mathrm{med}}$. Finally, the symmetric gaps condition (5) is clearly equivalent to the symmetry of the values of v about the center v_{med}, which is (4). This completes the proof. $\square$

B Bounding PG in terms of MM

Recall that the mean-median score MM is a signed score that is supposed to identify which party has a structural advantage, and by what amount. On the other hand, the partisan Gini PG is a non-negative score that simply quantifies the failure of symmetry, interpreted as a magnitude of unfairness. We easily see that $\mathsf{PG} = 0 \implies \mathsf{MM} = 0$ by comparing (3) and (4) in Theorem 3. In this section we strengthen that to a bound from below that is sharp in low dimension.

Let us define $\operatorname{discrep}(i) = \frac{v_i + v_{k+1-i}}{2} - \overline{v}$, measuring the difference between the average of a pair of vote shares from the average of all the vote shares. This gives

$$\mathsf{PG} = \frac{1}{k}\sum_{i=1}^{k} |2\overline{v} - v_i - v_{k+1-i}| = \frac{2}{k}\sum_{i=1}^{k} |\operatorname{discrep}(i)|.$$

Note that $\operatorname{discrep}(\lceil \frac{k}{2} \rceil) = \mathsf{MM}$, as observed above, and that $\sum_{i=1}^{k} \operatorname{discrep}(i) = 0$ by definition of $\overline{v}$.

Theorem 8. *The partisan Gini score satisfies* $\begin{cases} \mathsf{PG} \geq \frac{4}{k}|\mathsf{MM}|, & k \text{ odd} \\ \mathsf{PG} \geq \frac{8}{k}|\mathsf{MM}|, & k \text{ even,} \end{cases}$ *with equality when* $k = 3, 4$.

Proof. First suppose k is odd, say $k = 2m+1$. Then $\operatorname{discrep}(m) = v_m - \overline{v} = \mathsf{MM}$, so $\sum_{i \neq m} \operatorname{discrep}(i) = -\mathsf{MM}$. We have

$$\mathsf{PG} = \frac{2}{k}\sum_{i=1}^{k} \operatorname{discrep}(i) = \frac{2}{k}\left(|\operatorname{discrep}(m)| + \sum_{i \neq m} |\operatorname{discrep}(i)|\right) \geq \frac{2}{k}\left(|\operatorname{discrep}(m)| + \Big|\sum_{i \neq m} \operatorname{discrep}(i)\Big|\right)$$
$$= \frac{2}{k}\left(|\mathsf{MM}| + |-\mathsf{MM}|\right) = \frac{4}{k}|\mathsf{MM}|.$$

The argument for even $k = 2m$ is very similar, except that $\operatorname{discrep}(m) = \operatorname{discrep}(m+1) = \frac{v_m + v_{m+1}}{2} - \overline{v} = \mathsf{MM}$. So now those terms contribute $2|\mathsf{MM}|$ to the sum and the remaining terms contribute at least $2|-\mathsf{MM}|$, for a bound of $\mathsf{PG} \geq \frac{8}{k}|\mathsf{MM}|$. That completes the proof of the inequalities.

For $k = 3$ or $k = 4$, the term $\sum_{i \neq m} |\operatorname{discrep}(i)|$ is just $2|\operatorname{discrep}(1)|$, making the inequality into an equality. □

By a dimension count argument, it is easy to see that PG can not simply be a function of MM for $k \geq 5$. But in fact, we can also show that MM does not simply depend on PG. For instance, we can write

$$\mathsf{v} = (.2, .3, .4, .5, .7), \qquad \mathsf{v}' = (.2, .31, .39, .5, .7), \qquad \mathsf{v}'' = (.19, .31, .4, .5, .7),$$

giving $\mathsf{PG}(\mathsf{v}) = \mathsf{PG}(\mathsf{v}')$ while $\mathsf{MM}(\mathsf{v}) \neq \mathsf{MM}(\mathsf{v}')$. On the other hand, $\mathsf{MM}(\mathsf{v}) = \mathsf{MM}(\mathsf{v}'')$ while $\mathsf{PG}(\mathsf{v}) \neq \mathsf{PG}(\mathsf{v}'')$.

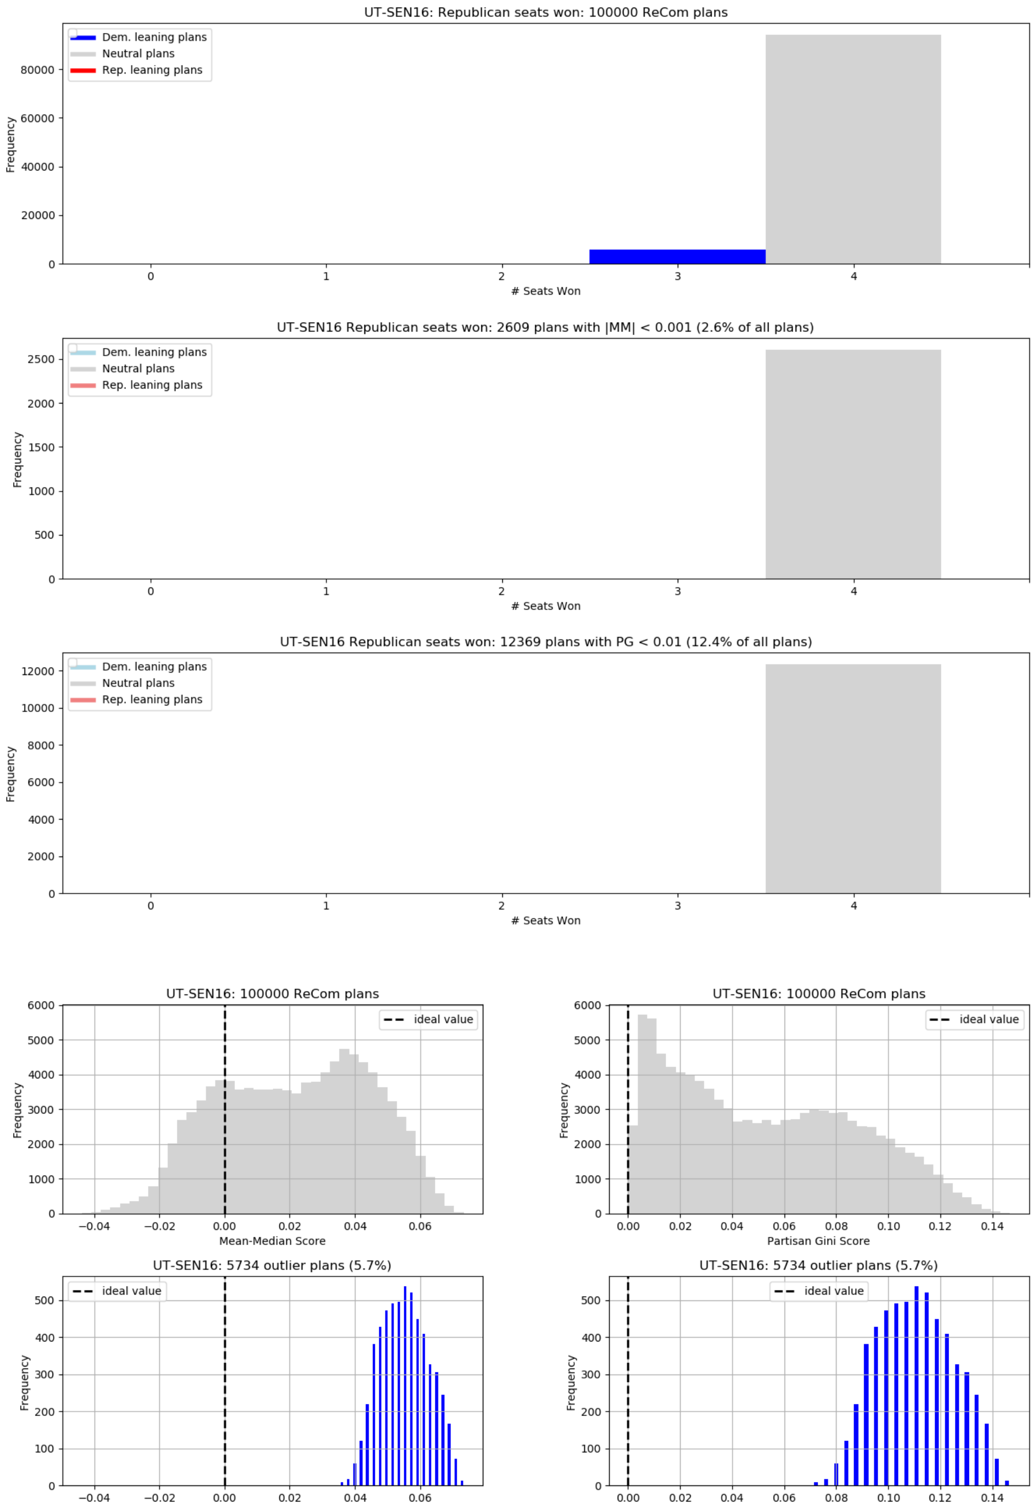

Figure 3: Ensemble outputs for Utah congressional plans with respect to SEN16 votes ($\bar{v}$ = .7155). Recall that the mean-median score reports a Republican advantage when MM > 0. The PG score is unsigned, but larger magnitude indicates greater asymmetry.

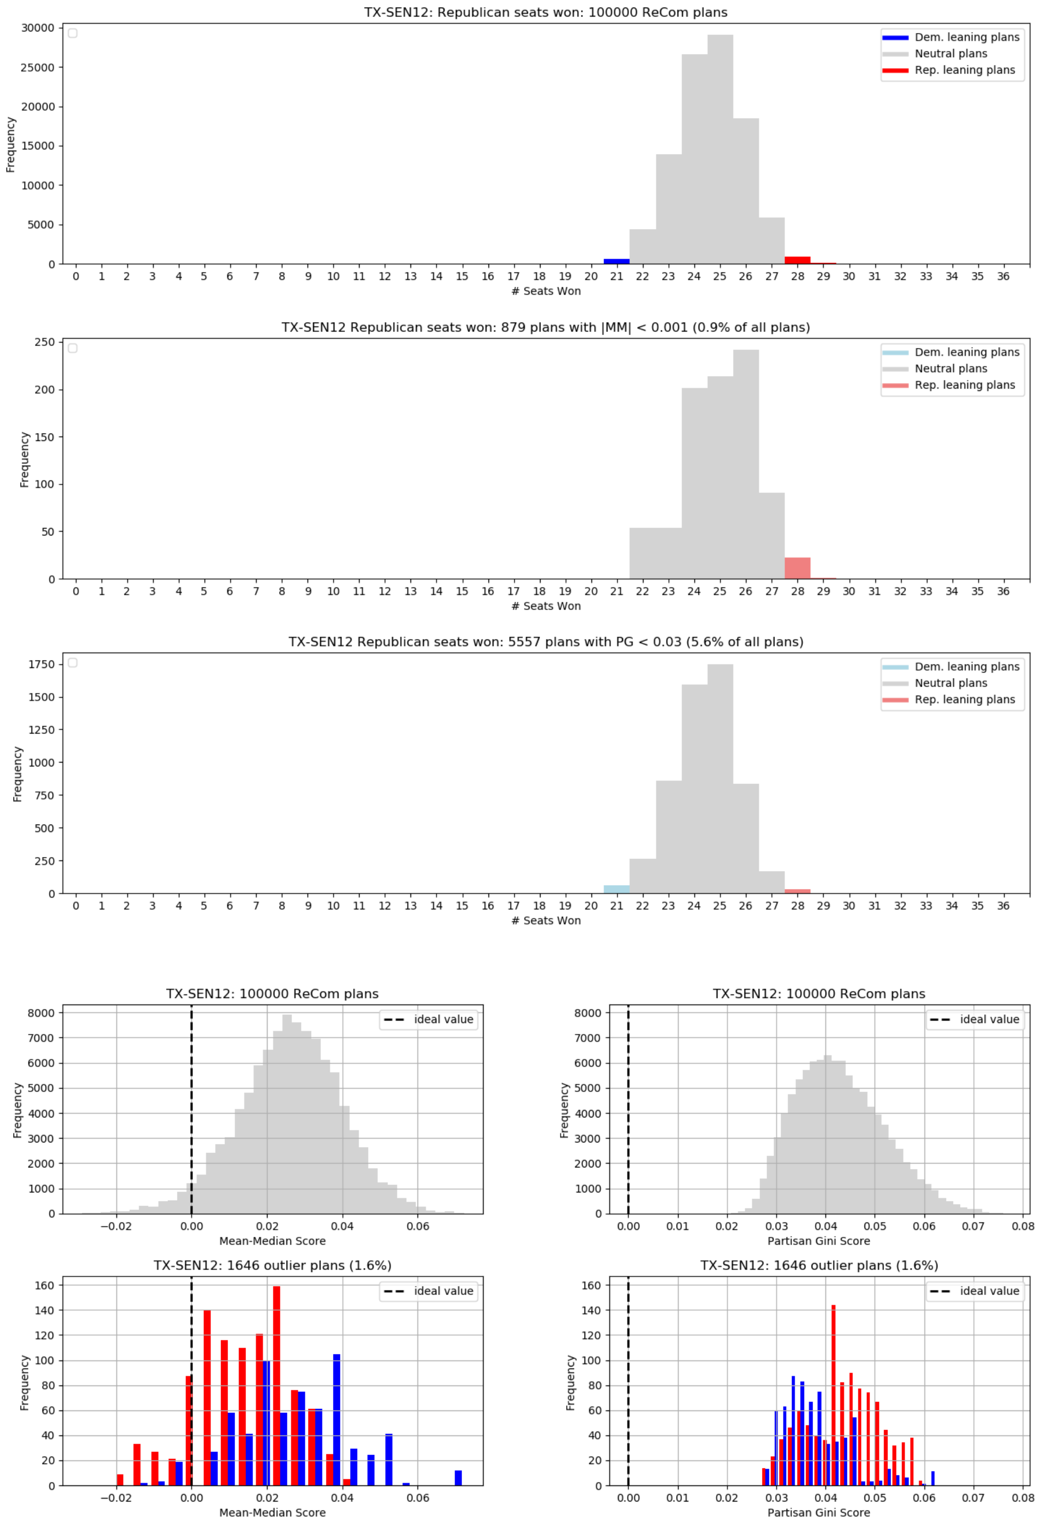

Figure 4: Ensemble outputs for Texas congressional plans with respect to SEN12 votes ($\overline{v}$ = .5815). Recall that the mean-median score reports a Republican advantage when MM > 0. The PG score is unsigned, but larger magnitude indicates greater asymmetry.

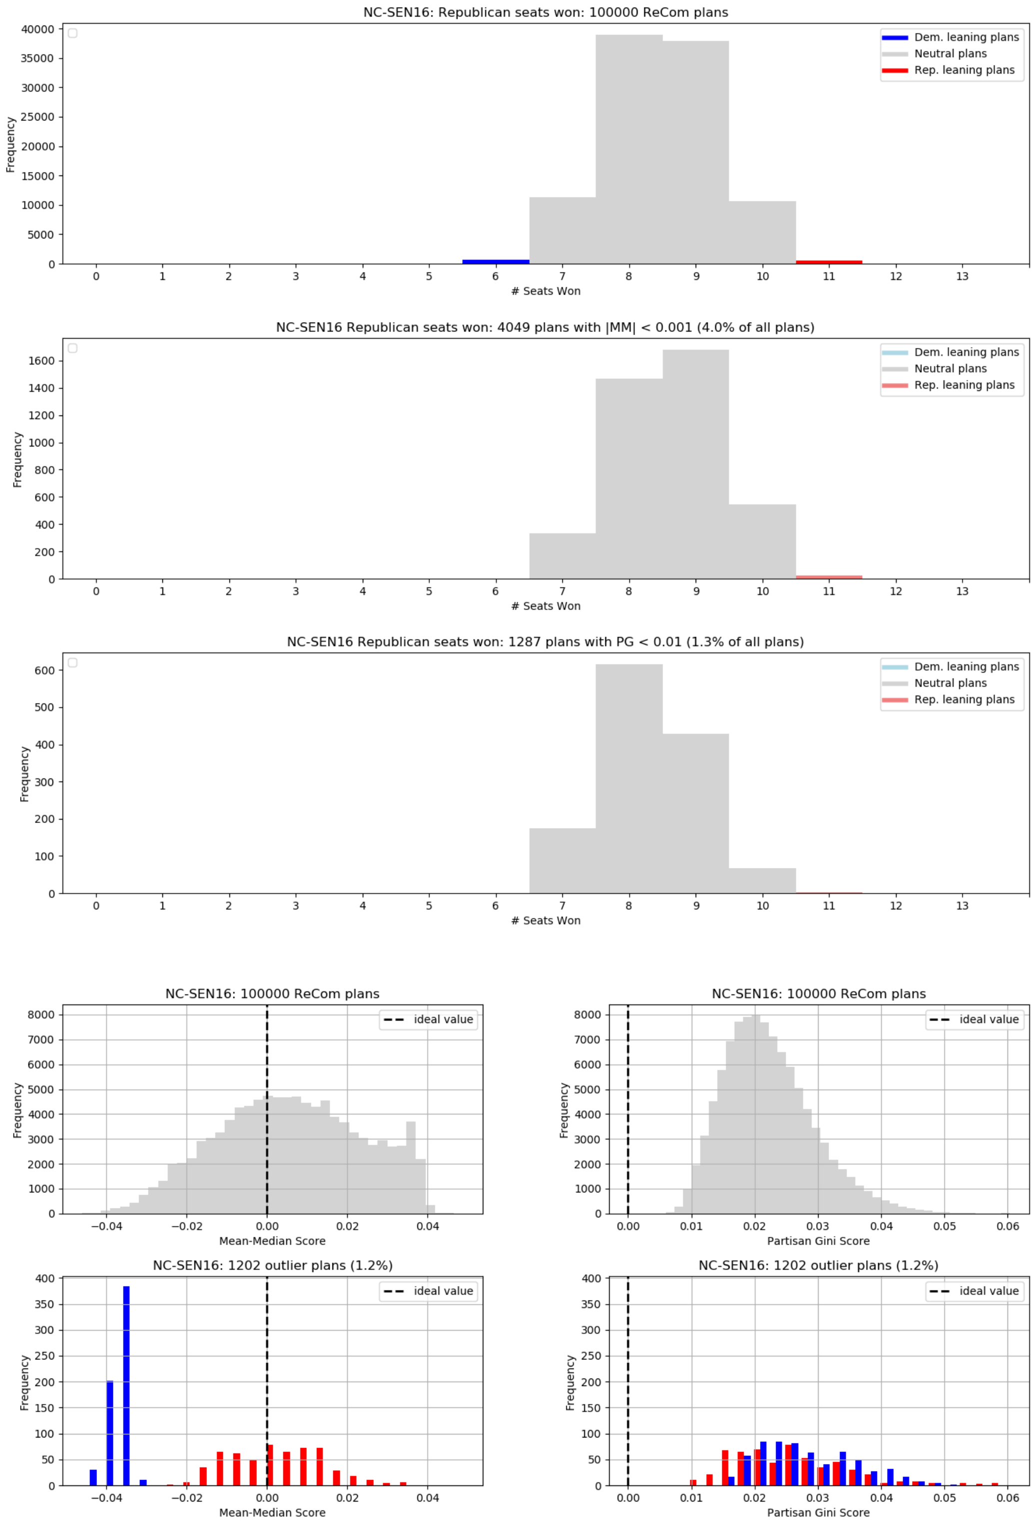

Figure 5: Ensemble outputs for North Carolina congressional plans with respect to SEN16 votes ($\bar{v}$ = .5302). Recall that the mean-median score reports a Republican advantage when MM > 0. The PG score is unsigned, but larger magnitude indicates greater asymmetry.

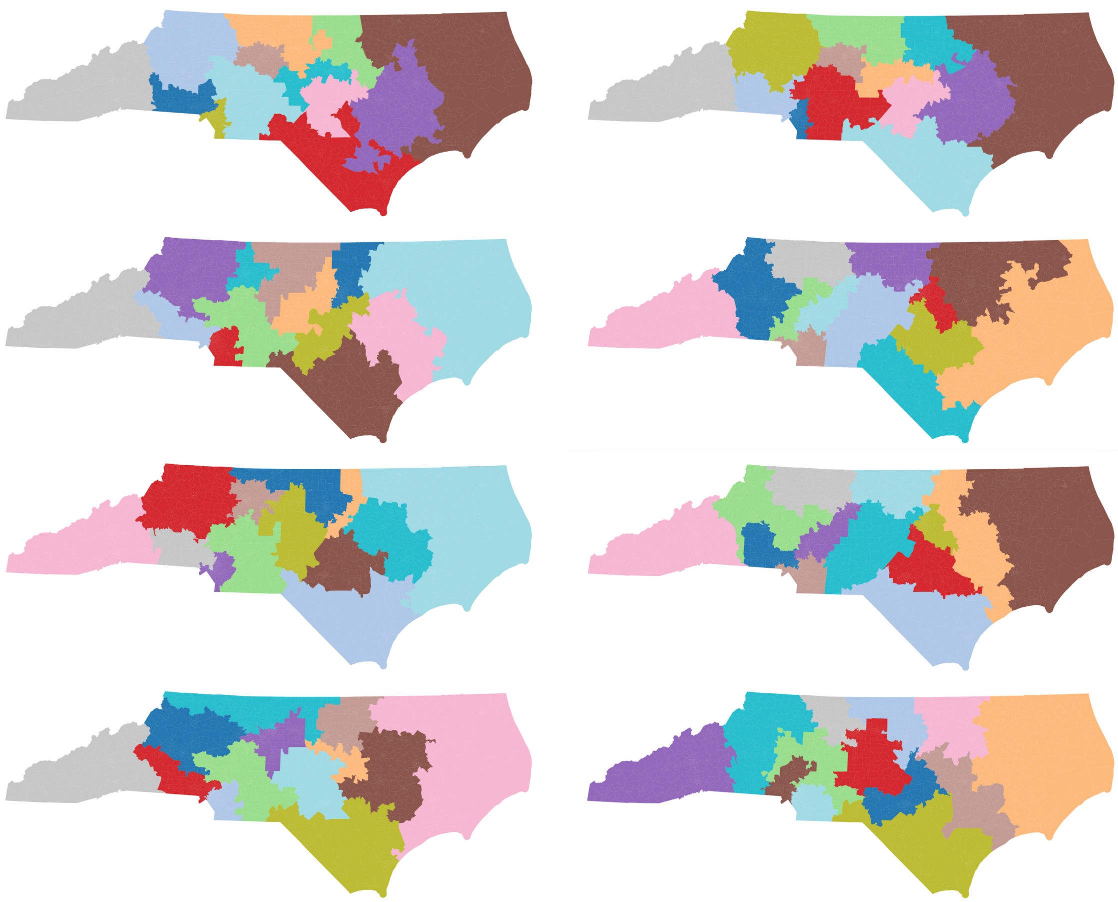

Figure 6: Each of these 13-district plans has 11 Republican-majority seats with respect to NC-SEN16 data, while having PG < .01 (in the best 3% of partisan symmetry scores). From top to bottom, the maps in the left column have PG scores of 0.0069, 0.0073, 0.0078, and 0.0098 respectively. From top to bottom, the maps in the right column have PG scores of 0.0072, 0.0075, 0.009, and 0.0099, respectively