To: Special Masters, North Carolina Superior Court, Wake County

From: Eric McGhee

Re: Measures of partisan fairness

Date: February 20, 2022

The Special Masters appointed by the North Carolina Superior Court of Wake County have asked me to provide my assessment of the partisan fairness of the remedial maps submitted by the parties to *NCLCV and Common Cause* v *Hall*, and *Harper* v *Hall*. I am a political scientist who studies elections, election administration, redistricting, public opinion, and legislative behavior. I am the creator of a popular measure of partisan gerrymandering called the efficiency gap, and co-creator with Nicholas Stephanopoulos of Harvard University of a legal test using the same. I am also a member of the Board of Directors of PlanScore, a nonpartisan website that scores redistricting plans on measures of partisan advantage. I have numerous published articles on the subject of partisan advantage in redistricting and am frequently consulted on the topic by policymakers and the media. I have a PhD in political science from the University of California, Berkeley.

The Special Masters have asked me to evaluate the remedial plans on four measures of partisan advantage: partisan symmetry, the mean-median difference, the efficiency gap, and the declination. In this memo I will explain the basis for each of these metrics to serve as background for the memos on the plans themselves.

All of these metrics measure some form of *partisan advantage*: an advantage for one party beyond what a normative concept of partisan balance would otherwise dictate. This normative baseline differs for each measure employed, meaning each captures a slightly different sense of fairness. As it happens, in a competitive state like North Carolina all the metrics tend to produce very similar results. I will describe the metrics and explain the particular sense of fairness each is meant to represent in the plainest language I can.

For those less inclined to learn the details, below is a quick summary of each measure:

- **Partisan symmetry** measures the excess seat share a party receives when it wins 50% of the votes. In a fair plan it should win 50% of the seats as well. It is measured in seat share.
- **Mean-median difference** measures the excess vote share required for a party to claim 50% of the seats. In a fair plan is should only take 50% of the vote to win half the seats. It is measured in vote share.

- **Efficiency gap** is the partisan balance of inefficient votes that do not contribute directly to a victory. In a fair plan the parties should have equal inefficient votes, reflecting no "packing" or "cracking." Despite using votes as its inputs it works out be a measure of seat share.
- **Declination** captures partisan differences in the pattern of district vote shares when plotted on a graph. In a fair plan the points will have similar patterns because they were drawn without the win/loss threshold specifically in mind. It is not measured in either vote share or seat share, so magnitudes are more difficult to compare with the other metrics.

Partisan Symmetry

Broadly speaking, partisan symmetry is the idea that parties with equal vote shares should receive equal seat shares. It is most commonly assessed at the point when both parties have half the votes, in which case both parties should also have half the seats. If one party has more than half the seats for exactly half the votes, it has managed to claim an outright majority of the legislative power without receiving majority support from the public. This violates the fundamental principle of majority rule.

Both parties rarely have exactly half the seats, so partisan symmetry almost always imagines what the outcome *might* be *if* both parties had equal vote shares. The most typical way of calculating this counterfactual is to simulate a uniform partisan tide that shifts each district's actual vote share an equal amount, to the point where both parties end up with half the votes in the aggregate. After shifting the vote shares this way, the analyst identifies how many seats have changed hands and records the new seat share. The difference between this new seat share and 50 percent is partisan symmetry. The larger the difference between the two numbers, the more unfair the plan. *This means partisan symmetry is measured in terms of seat share*: the seat share above or below 50%.

As an example, suppose there are two parties: Party A and Party B. Party A has 53 percent of the vote and Party B has 47 percent, and with that outcome Party A wins 60 of 100 seats in the state legislature. Partisan symmetry requires understanding what might happen if both parties had half the votes, so Party A must lose three percent of the vote in every district—a "uniform swing"—to bring its overall vote share down from 53 percent to 50 percent (and bring Party B's vote share up three points from 47 percent to 50 percent). If Party A loses, say, four seats as a result of this counterfactual, then it will end up with 56 percent of the seats (56/100) for 50 percent of the vote, a fairly clear violation of the principle of majority rule.

This is not the only way that partisan symmetry has been implemented. Since partisan symmetry says that

parties with equal vote shares should receive equal seat shares, it can actually be calculated for any pair of vote shares. In the example above, one might explore what happens if the party's roles were reversed, and Party A received 47 percent of the vote and Party B 53 percent. In a symmetric plan, Party A should receive 40 percent of the seats just as Party B did when *it* had 47 percent of the vote. If it has more than 40 percent, the plan favors Party A; if it has less than 40 percent, the plan favors Party B. This is the approach to symmetry used by plaintiffs Harper et al. and their experts.

As a practical matter this version of symmetry tends to give similar answers to the one calculated at 50 percent, but the two can diverge in any given plan. I prefer the "50%" version for two reasons. First, this "vote-swapping" version of symmetry necessarily requires a counterfactual twice as large as the 50% version. In the example above, the uniform swing is 3 points for the version assessed at 50%, and 6 points (53% - 47%) for the vote-swapping version. Sometimes this makes sense, but often it can take the counterfactual into highly fanciful scenarios.

Second, the vote-swapping version has a more tenuous connection to the principle of majority rule. The 50% version tests a normatively critical threshold: the point at which more voters support a party than support its opposition. If a party manages to acquire more power—in the form of seats in the relevant legislative body—without clearly winning more support, there is a universal sense that something is amiss.

When the parties are competitive with each other statewide, each has a vote share close to 50 percent already and the counterfactual is not very large—in fact, something close to it may have happened recently or could be expected to happen soon. This is the case with North Carolina, making partisan symmetry a useful metric for this state.

Mean-Median Difference

The mean-median difference is just like it sounds: it is the difference between the average district vote share and the median district vote share (where half the districts have a vote share higher and half lower). The mean-median difference favors a party when its median vote share is *higher* than its mean vote share.

Returning to the example above, Party A has an average (mean) vote share of 53 percent. Imagine that the median vote share is 60 percent, so that Party A has a vote share higher than 60 percent in half the seats, and lower than 60 percent in the other half. The mean-median difference in this case would be -7 percent: 53 percent minus 60 percent. This would be a substantial advantage for Party A.

The mean-median difference may appear to measure something purely mathematical, but in fact it captures something very straightforward: the vote share above (below) 50% that a party needs in order to capture

exactly half the seats. If a party needs more than 50 percent of the vote to claim half the seats, it is at a disadvantage. *This means the mean-median difference is measured in terms of vote share*: the vote share above or below 50% required to win half the seats.

The mean-median difference and the 50% version of partisan symmetry are close cousins. Partisan symmetry is the seat share above (below) 50% that a party receives when it has half the votes; the mean-median difference is the vote share above (below) 50% that a party needs in order to win half the seats. So the mean-median difference is a counterfactual every bit as much as partisan symmetry. Again, for a competitive state like North Carolina, the counterfactual is entirely plausible and so not an issue.

The Efficiency Gap

Every single-member district plan has "inefficient" votes that do not contribute directly to a victory. These include votes beyond the number needed for a candidate to win, and those cast for a candidate that will definitely lose. Though these votes do not change the outcome in the race where they were actually cast, they might be enough to help a candidate of the same party win in a neighboring district instead.

The party with fewer inefficient votes than the opposition has an advantage, because its support translates more efficiently into victories. This is why a gerrymander tries to "crack" most of the opposition's supporters across many districts where victory is close but still elusive, and "pack" the remainder in a small number of districts which that party will win by large margins.

The efficiency gap captures this partisan difference in efficient votes. It is the difference between each party's total inefficient votes, divided by all the votes cast in the election. If both parties have equal inefficient votes the efficiency gap is zero and the plan is as balanced as possible.

Despite using votes as its inputs, the efficiency gap works out mathematically to be a measure of excess seat share. A balanced efficiency gap expects that for every additional one percent of the vote beyond 50%, a party should get an extra two percent of the seats. This is the "winner's bonus" so common to single-member district electoral systems, where simple proportionality between seats and votes is extremely rare. The efficiency gap says a party can receive a higher seat share than vote share and still have a fair result, but there is a limit to how high that gap can go. *This means the efficiency gap is measured in terms of seat share*: the seat share above or below the ideal implied by balanced inefficient votes.

Though it comes at the problem from a very different direction, the efficiency gap also satisfies the symmetry principle. In fact, when both parties have half the votes, the efficiency gap says they should each have half the seats, and partisan symmetry and the efficiency gap are equal. For competitive states like North Carolina,

the two metrics tend to give very similar results, except that the efficiency gap is calculated at the actual election result instead of at a counterfactual.

The Declination

Another way of thinking about a gerrymandering party is that it pays very close attention to the win/loss threshold when designing the districts. It is acutely conscious of which party will win each seat because it is trying to extract as many winning seats as possible.

The declination captures this idea through some clever geometry. Imagine plotting a plan's districts in order from lowest to highest vote share for the minority party (left to right). The vertical axis is the vote share in each district. The districts are ordered by design, so each point is always higher than all the ones to its left. In a plan drawn without regard to which party wins each seat, these points will likely just climb higher until they cross the win/loss threshold. The constellation of points above and below the win/loss threshold will look similar. Examples of such patterns are in Figure 1.

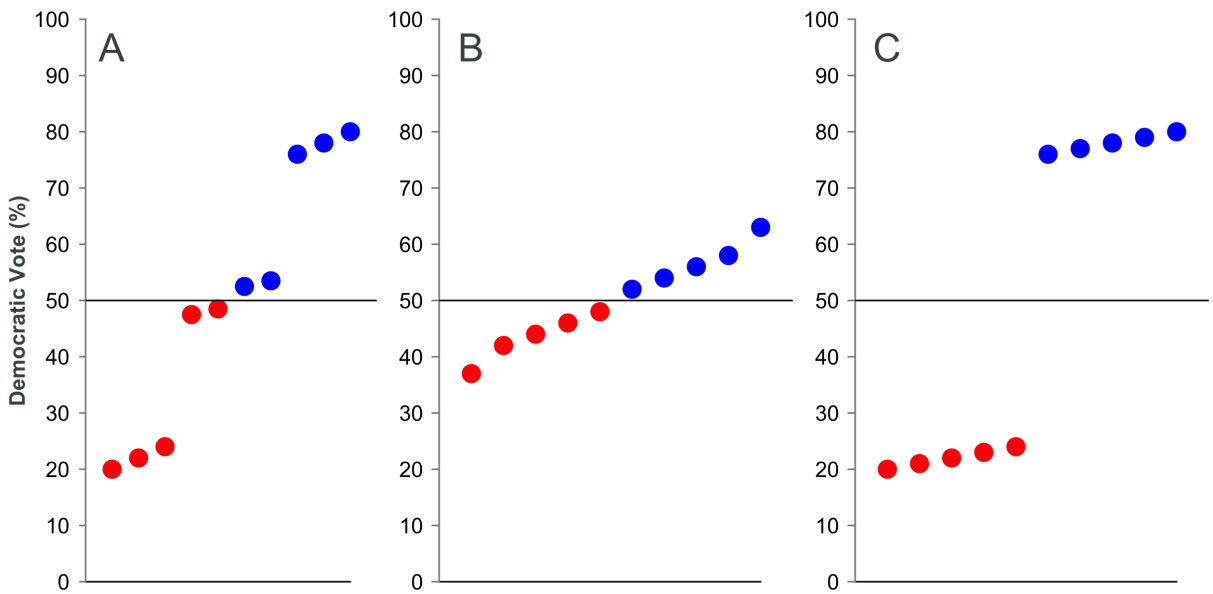

Figure 1: Fair plans are balanced around the win/loss threshold

In an unfair plan, the gerrymandering party will force as many seats on its side of the win/loss threshold as possible, leaving a few districts for the opposition to win by large margins. This is the basic logic of the efficiency gap. As an example, in Figure 2 the pattern of points on one side of the win/loss threshold is very different than on the other.[1]

[1] This example is an actual election from North Carolina's 2011 redistricting plan.

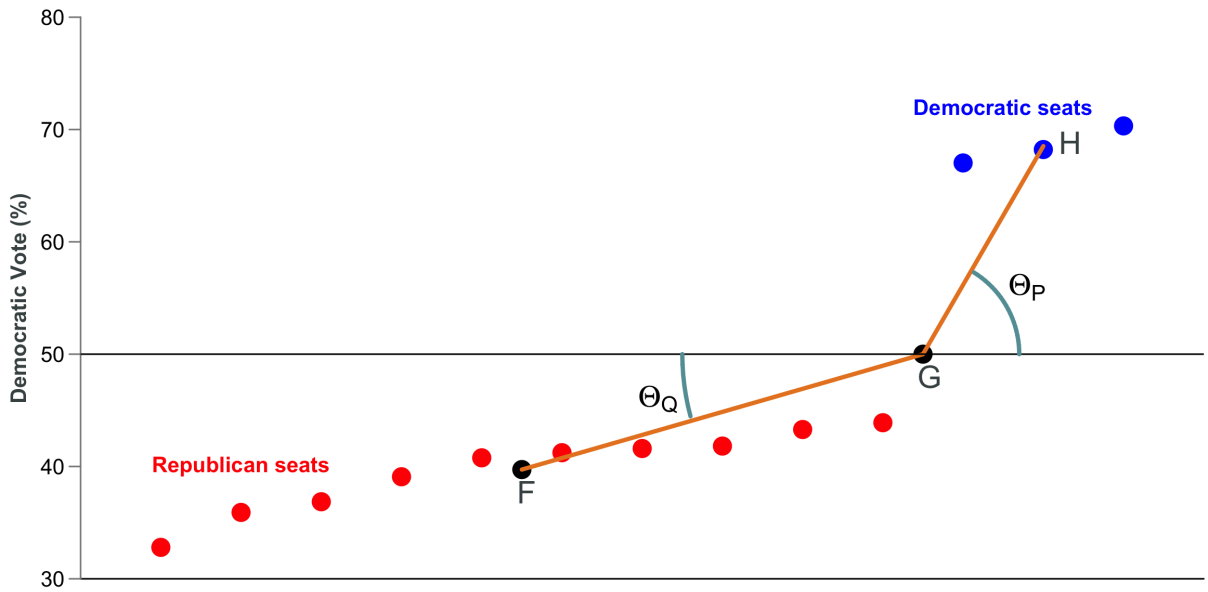

Figure 2: Unfair plans skirt the threshold in unequal ways

The declination connects the center of each party's points to a point on the win/loss threshold line. It turns out that differences between angles formed by these lines offer a way to summarize the difference between the two groups of points, and so to summarize the bias in the way the districts were drawn. However, *this means the declination is measured in units other than votes or seats*, making its magnitude difficult to compare to the other measures. It is best to first compare its direction to those of other measures of the same plan, and then compare its magnitude to the declination values of other plans.

Like the efficiency gap, the declination is calculated without a counterfactual, using the actual election results. The two metrics are in fact highly correlated, but the declination does not expect a particular relationship between vote share and seat share the way the efficiency gap does. The declination may best be considered a measure of intent: how much evidence is there that the pattern of election results is distorted in a way that implies an effort to win more seats for one side?