ELECTION LAW JOURNAL
Volume 14, Number 4, 2015

DOI: 10.1089/elj.2015.0311

Original Articles

Measures of Partisan Bias for Legislating Fair Elections

John F. Nagle

ABSTRACT

Several measures of partisan bias are reviewed for single member districts with two dominant parties. These include variants of the simple bias that considers only deviation of seats from 50% at statewide 50% vote. Also included are equalization of losing votes and equalization of wasted votes, both of which apply directly when the statewide vote is not 50% and which require, not just partisan symmetry, but specific forms of the seats-votes curve. A new measure of bias is introduced, based on the geometric area between the seats-vote curve and the symmetrically inverted seats-votes curve. These measures are applied to recent Pennsylvania congressional elections and to abstract models of the seats-votes curves. The numerical values obtained from the various measures of bias are compared and contrasted. Each bias measure has merits for different seats-votes curves and for different elections, but all essentially agree for most cases when applied to measure only partisan bias, not conflated with competitiveness. This supports the inclusion of partisan fairness as a fundamental element for election law reform, and some options are discussed.

1. INTRODUCTION

It is well recognized, in the popular press, in public opinion, and in the academic literature, that gerrymandering has been heavily practiced in redistricting, and that this is one of the likely factors in the widely perceived dysfunction of the federal government and some state governments (Mann and Ornstein 2013). While reform has occurred, notably in California (Mac Donald 2012), reform should be considered in many other states, and it is therefore pertinent to consider what such reform should accomplish. Does it suffice that the redistricting process be apolitical, not involving legislators whose interests are clearly at stake? Does it suffice that single member districts be compact and that political subdivisions be minimally subdivided? Is it important that minorities and communities of interest obtain districts likely to elect their representatives?

In this article, it is assumed that partisan fairness should be a primary outcome of redistricting.[1] Conventional redistricting criteria, such as compactness and respect for political subdivisions, even when used by non-partisan map makers, may nevertheless result in political bias via unintentional gerrymandering (Grofman and King 2007, Chen and Rodden 2013). Since these criteria do not appear to suffice to bring about partisan fairness, then reform legislation should aim to create a process that does produce it.[2] As modern computation allows the generation of many districting maps, reform could constrain the permissible pool to those

John F. Nagle is a professor emeritus in the Department of Physics at Carnegie Mellon University in Pittsburgh, Pennsylvania.

[1]It has been suggested that this is a more important consideration for congressional than for state and local districting (Gardner 2012).

[2]This article accepts the confines of single member districts mandated by Congress for congressional districts, noting in passing that much partisan bias could be relieved by alternative voting systems (Amy 2000).

that have acceptably limited political bias.[3] A prerequisite to this kind of reform is being able to measure partisan bias. That is the primary subject of this article.

Of course, measuring partisan bias is important for litigators to prepare suits claiming gerrymandered districting and for courts to decide such cases. Although gerrymandering is not justiciable in the opinions of many judges, other judges have written otherwise, but often with the reservation that a definitive measure of bias is lacking. The history has been thoroughly discussed in this journal (Grofman and King 2007). While litigation and the courts are important last resorts, partisan fairness legislation would be more expeditious, avoiding having to wait for an election and the subsequent time for a case to be decided and then wend its way through appeals. However, in the event that districting is challenged, either on partisan bias grounds or on other grounds, legislative mandating of partisan fairness would require courts to accept it as justiciable. Furthermore, legislative guidelines for an unacceptable degree of partisan bias would also facilitate court decisions by reducing the burden of litigators and courts having to construct their own criteria *de novo*.

Partisan fairness is sometimes conflated with the concept of competitiveness, aka responsiveness or representation. However, it has been appropriately stressed that these are separate concepts with separate measures (King and Browning 1987). One major difference between these concepts regards how to frame their respective criteria quantitatively. Obviously, the quantifiable criterion for partisan fairness is simply to minimize partisan bias, whereas maximizing competitiveness is not so obviously a good criterion (Hirsh and Ortiz 2005; Buchler 2011). Although such considerations motivate focusing primarily on partisan fairness, quantifiable competitiveness ideals do emerge from some measures of partisan bias (McGhee 2014; Stephanopoulos and McGhee 2015), so competitiveness/responsiveness will not be ignored.

The simplest qualitative diagnostic, frequently mentioned as evidence of partisan bias, is to compare the fraction of votes won by a party with the fraction of seats won. This is a very strong diagnostic of bias when the fraction of votes won exceeds half but the fraction of seats won is less than half. However, this diagnostic is insufficient to indicate partisan bias when fractions of votes and seats are both less than half. For example, there is not necessarily partisan bias if a party wins 40% of the votes and only 20% of the seats. This could indeed be a fair outcome when many seats are highly competitive, and it conforms to the so-called "cube law" (Kendall and Stuart 1950). In the limit of all seats being totally competitive, the expected outcome is zero seats for a party that wins 1% fewer votes than its opponent; this is the winner takes all extreme of competitiveness. The assumption that the fraction of seats should equal the fraction of votes assumes that proportional representation is the ideal, and that is highly arguable (King 1989), especially as it does not allow higher levels of competitiveness. As this simple diagnostic of partisan bias is incomplete at best and also difficult to quantify generally, it is appropriate to consider more specific measures. Let it be emphasized, however, that even this qualitative diagnostic remains powerfully persuasive.[4]

This article will consider several measures of bias in the literature and it will present some new ones.[5] While this may seem to be more than enough, having different measures of partisan bias is somewhat similar to having different measures of compactness. Although it would be convenient if there were a single, obviously superior and universally accepted measure for any redistricting criterion, complex social problems often do not yield to simple solutions. In the case of compactness, a reasonable course has been to employ different measures and compare the results. Such a practice can also be envisioned for measuring partisan bias and the comparisons in this article may help to promote that. Although this article presents a measure that the author suggests is superior

---

[3]McDonald 2007. As an example, Chen and Rodden 2013 generated many maps for Florida. Although the ostensible main theme of their paper was that political geography creates political bias, there were many maps that predicted less political bias than the average randomly generated map, so reform could constrain the allowed pool to those. Mattingly and Vaughn 2014 have recently provided another notable example of generating many compactness constrained maps for North Carolina; noting how far the map used in 2012 deviated from the mean bias in their generated pool, they call for a constrained pool.

[4]As emphasized by Baker 1990.

[5]Note that there have been many other proposed measures of partisan bias. Eight different measures were discussed by Grofman 1983.

mathematically to other measures, a conclusion of this article is that there are several measures that essentially agree, so having different measures should not block the inclusion of partisan bias in election reform.

In the next section, the well-known seats/votes (S/V) graph is reviewed as this has been a much-used construct for quantitative discussion of partisan bias. Examples are given that will be used in subsequent sections to evaluate and compare measures of partisan bias, starting with the bilogit model. Also included is an apparently new way to construct the S/V graph that is applied to the 2011 Pennsylvania congressional redistricting. This section also proposes two criteria that measures of bias should adhere to, and two examples are presented that are subsequently used to test measures. In Section 3, one of these examples reveals a defect in the simple diagnostic of bias and its corresponding measure that will be designated B_S, where the subscript S refers to simple. Also, a much-used variant is reviewed, designated B_{GK}, which ameliorates the defect in B_S. Another increasingly popular simple measure, designated $B_{M\text{-}m}$ in this section, is shown to violate the second general criterion for the last example in Section 2. A new quantitative measure of bias, B_G, is proposed in Section 4; it follows naturally from asymmetry in S/V graphs, being zero when the graph is symmetric. The measures B_S, B_G and B_{GK} are based on the principle that the expected outcomes should be symmetrical with respect to the parties, whereas $B_{M\text{-}m}$ shifts the focus to fairness for individual voters. Section 5 discusses several additional measures of bias that are based, in several different ways, on voter happiness/dissatisfaction. Like the $B_{M\text{-}m}$ measure, these are rooted in the concept that voters in each party should be treated equally. These measures differ from the symmetry measures, which are rooted in the principle that the outcomes should treat each party equally. These happiness measures each lead to an ideal S/V graph which is also symmetrical and which also prescribes an ideal competitiveness. The B_L measure, based on lost votes, leads back to proportionality and the B_W measure, based on wasted votes, leads to a more competitive S/V graph (McGhee 2014, Stephanopoulos and McGhee 2015). Section 6 focuses on comparing the different B measures and a more general discussion ensues in Section 7. Appendix A presents the derivation of the way that S/V curves can be obtained from election results.

2. SEATS/VOTES AND RANK/VOTES GRAPHS

As is well known, S/V graphs plot the predicted number of seats S versus the statewide percentage of votes V won for all districts combined. Let us begin by reviewing the bilogit family of model S/V curves. This family has been historically important because it emphasizes the distinction between partisan bias and competitiveness by embedding two independent parameters, λ for bias and ρ for competitiveness (King 1989; King and Browning 1987). Figure 1 shows some members of this family. The straight line for $\lambda = 0$ and $\rho = 1$ is the proportional S/V curve where the percentage of seats equals the percentage of the statewide vote. The $(\lambda,\rho) = (0,3)$ curve is like the so-called cubic law; it is more responsive than proportionality, having a larger slope of 3 at $V = 50\%$, whereas the less responsive (0,0.5) curve has a smaller slope of 0.5. All curves with zero bias, $\lambda = 0$, have $S = 50\%$ at $V = 50\%$, but this changes with non-zero bias. Positive values of λ make S less than 50% when $V = 50\%$, and by the same amount independently of the value of ρ.

While the literature clearly recognizes the importance of S/V graphs, there has been much appropriate discussion about how to construct them.

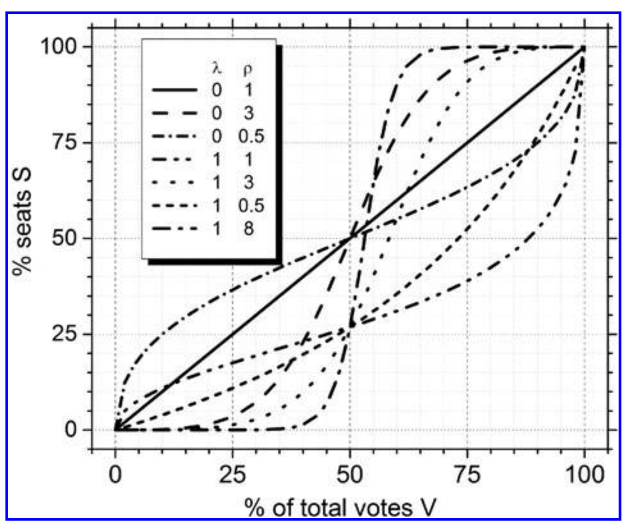

FIG. 1. Bilogit seats/votes (S/V) curves for several combinations of bias λ and responsiveness/competitiveness ρ.

Of course, to evaluate new plans before the first election, prior election returns and/or registration data by voter tabulation district (VTD) would be used.[6]

Our next example is an S/V graph for the 2011 Pennsylvania congressional districting. It is constructed using 2012 election returns.[7] The first step obtained a rank/vote (r/v) curve. To see how this was constructed, look first at the left hand axis of Figure 2. The districts are rank ordered from lowest percent Republican (R) vote to highest. The R vote for each district is shown on the upper horizontal axis. Only the 12th district had a close vote; it is sixth in the rank order and was won by the Republican with 51.8% of the vote. The X at 50.75% indicates the percentage of the statewide Democratic (D) vote.

If the statewide vote had been 50% for both parties, and if we assume a uniform shift of 0.75% in each district, the squares in Figure 2 would be shifted right to the corners of the solid line. That line is often taken to be the S/V curve where the number of Democratic seats S would be given on the right hand axis for the statewide Democratic vote V on the lower horizontal axis. However, this is not a valid S/V curve because it assumes that a uniform statewide swing of 40% to the Republicans would be required before they win all the seats, the last one being the one that had voted 90% Democratic in 2012. Such a uniform shift in all districts leads to the impossibility of requiring more than 100% Republicans in five of the districts.[8] The uniform shift assumption is quite unrealistic on its face because it assumes that the same number of D's would shift in districts with few D's as in districts with many D's.

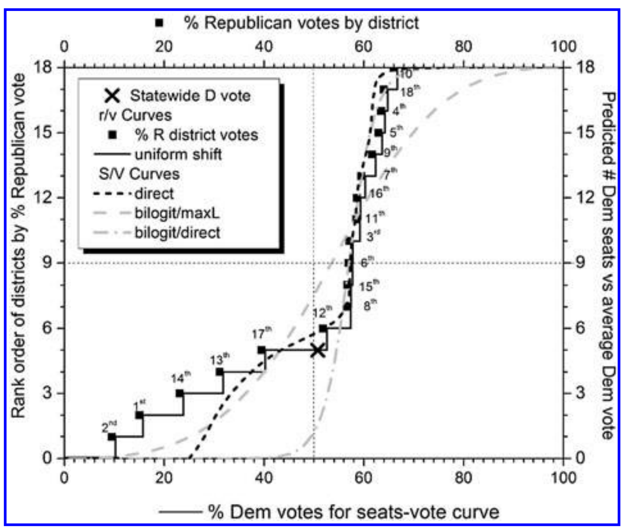

FIG. 2. The *squares* show the percentage Republican vote (*upper horizontal axis*) ranked by district (*left hand axis*) for the 2012 PA congressional election. A uniform shift of 0.75% corresponding to the statewide D vote of 50.75% gives the rank/vote (r/v) curve shown as the *stepped solid line*. The *smooth dashed* curve is a seats/votes curve with predicted number of Democratic seats on the *right hand axis* versus the statewide percentage of Democratic votes on the *horizontal axis*. The calculation is explained in the text and Appendix A. The *dash-dot grey* curve is a direct bilogit fit to the S/V curve and the *dashed grey line* is the maximum likelihood bilogit approximation to the election data. The *small numbers* next to the *squares* identify the PA districts.

A much more realistic model for obtaining an S/V curve from r/v data[9] is that a statewide percentage shift is equally likely to apply to any voter in any district; this is plausible and it has the merit that it does not allow the percentage of voters of either party in any district to fall below 0% or above 100% for any statewide vote shift. Simple math given in Appendix A[10] shows that any district, identified by subscript n, currently won by a party with v_n percent of the district vote and V percent statewide vote will be lost when the statewide vote falls to $V/2v_n$. This leads to the S/V curve for PA shown in Figure 2. Even if a district has 100% D vote when the statewide D vote is 50%, that seat would be lost when the statewide vote falls to 25%.

There is little difference between the S/V and the r/v curves for the central region where $40\% < V < 60\%$. In either case Figure 2 predicts that Democrats would have to obtain 58% of the statewide vote to obtain half the seats. (Assuming that the rank order remained the same, that would shift districts 12, 8, 15, and 6 to the Democrats, but which particular district seats would shift is not part of an S/V curve prediction.) A swing to 58% Democrats in PA would be a landslide swing. Supposing

[6]Backstrom, et al. 1990 advocated using election returns for a low-profile statewide base race, Gronke and Wilson 1999 averaged three races, and McDonald 2014 uses presidential votes.
[7]There were only a few third party candidates and those votes were pro-rated into the percentages shown.
[8]The uniform shift construction of S/V curves has been appropriately criticized by many, particularly King 1989.
[9]Note that the S and the V have quite different meanings in the S/V curve, for which S is the number of expected Democratic seats S and V is the statewide Democratic vote V, than the r and v in the r/v curve, for which the r is the rank order of Republican vote and v is the district level R vote percentage.
[10]Appendix A includes examples that illustrate the method and that also supplement the examples in this section.

that the mean probability of normally distributed swings is as large as 5%, there would be less than 8% probability of Democrats getting 58% or more of the total vote, so Figure 2 predicts that Democrats would have six or fewer seats 92% of the time with the current districts. Figure 2 also emphasizes that there is little responsiveness when the vote is in the 45%–55% range. In the 2014 midterm election, the overall congressional vote swung to 55% R, but there was no change in the number of party seats, nor in any of their districts. This S/V curve is even less responsive (less competitive) than proportionality; this is consistent with sweetheart bipartisan gerrymandering–incumbents of both parties like safe seats.

Another way used here to obtain an S/V curve from the PA election data employed the maximum likelihood method of King 1989. This method assumes that the underlying voting system is a member of the bilogit family. Application of this method to the r/v district data in Figure 2 gives the values $(\lambda,\rho)=(0.32,2.3)$. Using these values gives the dashed grey bilogit curve shown in Figure 2.[11]

It will be informative to consider simpler S/V graphs for testing measures of bias against objective criteria. Our first criterion says that a measure of bias should vary gradually when the r/v data vary gradually. It is based on the premise that a districting plan has an underlying value for its bias. If bias has any meaning, then its value should not change drastically, especially not discontinuously, with small variations in v_n. Small variations would naturally be brought about by subsequent elections, between which the change in fundamental bias would be expected to be small. Furthermore, when evaluating competing districting plans, a good bias measure should give a sensibly small range of values from the small variations in v_n that would ensue from using different baseline prior election results. This will be referred to as the smoothness criterion. A simple example to test the smoothness criterion has only three districts. District 1 always has $v_1=35\%$. District 2 has $v_2=50+x$ and $v_3=65-x$, where x is a parameter to bring about small variations. For simplicity this example has overall $V=50\%$ for all x. The r/v graph for this example is shown in Figure 3.

As is well known, the best way for a party to gain advantage is to pack opponent voters into a few districts while providing a comfortable winning margin of supporters in the majority of the remaining districts. Our second criterion, to be called the packing criterion, says that a measure of bias should be responsive to different levels of packing. The following example (number 4) to test this criterion has many districts. One set of districts is gerrymandered to give $v_n=60\%$ for the R party and the remaining set has $v_m=40\%-x$ where x is a variable percentage. Compatible with the assumption that the overall vote shouldn't vary with which district the voters are put into, this kind of example requires the overall vote to remain constant for all x and for simplicity the constant overall vote is taken to be $V=50\%$. Then, simple algebra shows that the fraction f of R seats is given by $f=(x+10)/(x+20)$. For $x=0$ each party obtains half the seats

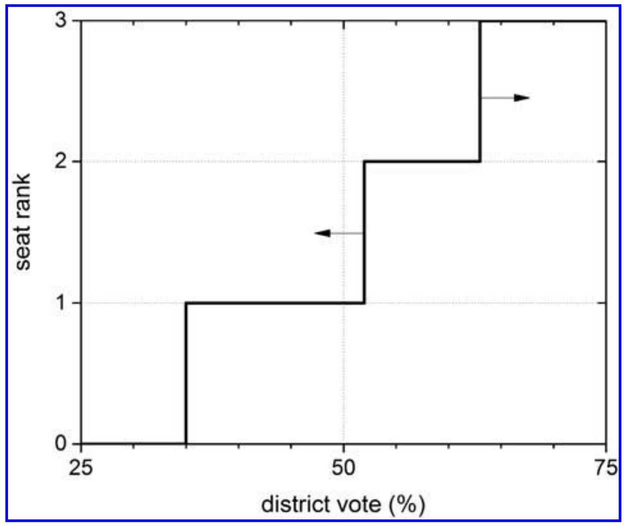

FIG. 3. The rank versus district vote plot for example 3 in the text with $x=2$. The *arrows* show how the votes in the 2nd and 3rd districts may be varied in a coordinated fashion to maintain statewide $V=50\%$.

[11]Note that this curve is not the straightforward least squares fit to the S/V data points. That fit gives $(\lambda,\rho)=(2.6,9.3)$ which is shown as the dash-dot grey line in Figure 2. It is not surprising that the bilogits fail substantially. All potential S/V curves consist of a non-zero fraction of all functions that increase monotonically with appropriate constraints on V at $S=0$ and the maximum S. Even if the S/V curve is only required to fit N discrete districts, the corresponding S/V family requires at least N-2 parameters as seen by considering just the finite Taylor series family consisting of N terms. Therefore, a possible S/V curve will not necessarily be well fit by the two parameter bilogit family when $N>4$. Nevertheless, the bilogit remains a useful model for comparing measures of bias. Rather more sophisticated methods to obtain S/V curves, at least for central ranges of V, have been proposed by Gelman and King 1994 and employed in the JudgeIt software (Gelman et al. 2012).

(f = 1/2). The fraction f of R seats increases as x increases and more of the B voters are packed into fewer districts. The r/v graph for this example 4 is shown in Figure 4.

3. SIMPLE BIAS MEASURES B_S, B_{GK}, AND $B_{M\text{-}M}$

S/V curves suggest a simple measure of partisan bias, designated B_S, based on the difference in the percentage of seats from 50% on the S/V curve when the statewide partisan vote V is 50%. For example, in Figure 1, for all the biased cases with a bias $\lambda = 1$ against the party in question, the expected percentage of seats is 26.9% when the statewide vote is 50%, so the value of B_S would be set to 50%–26.9% = 23.1%. For Figure 2 the expected number of seats is 5.8 and the total number of seats is 18, so $B_S = (9–5.8)/18$ which is a 17.8% bias in favor of the Republicans. However, the example in Figure 3 shows something that can go wrong with the B_S measure. As x passes through zero B_S passes from (2–1.5)/3 = 17% bias to (1–1.5)/3 = −17% bias which dramatically fails the smoothness criterion in the previous section.

To overcome this deficiency in the B_S measure, one may follow Gelman and King 1994 and Gelman and King 1990 by using a modified bias, to be designated here as B_{GK}. This first calculates the average number of seats from S(V) in the 45%–55% V interval before performing the same B_S calculation. However, a concern about this measure is the arbitrariness of choosing a 45%–55% interval rather than some other interval, like 47%–53%.The B_S measure of bias measures the deviation in seats from 50% in a vertical slice of the S/V graph at the V = 50% level. Another measure essentially slices the S/V graph horizontally at the S = 50% level and obtains the deviation of the vote from 50%. This measure can be, and usually is, applied directly to the r/v data by subtracting the mean vote V from the vote for the S = 50% seat (McDonald 2009, McDonald and Best 2015, Wang 2015).[12] Let us designate this median minus mean method as $B_{M\text{-}m}$.[13] The example in Figure 4 shows something that can go wrong with the $B_{M\text{-}m}$ measure. As x increases from 0 the vote for the median seat remains at 60% so $B_{M\text{-}m}$ remains constant at 60% – 50% = 10% even though packing increases continuously as x increases.[14]

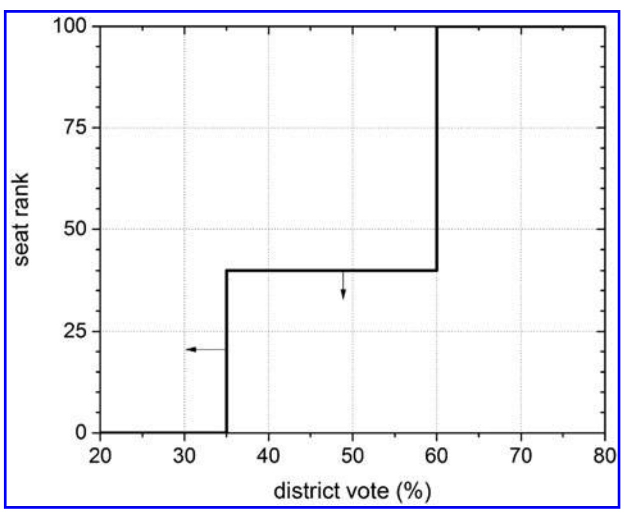

FIG. 4. The rank versus district vote plot for example 4 in the text with x = 5. The *arrows* show how the votes in the 2nd and 3rd districts may be varied in a coordinated fashion to maintain statewide V = 50%.

Values for these measures applied to the examples are given in Table 1.

4. GEOMETRICAL MEASURE OF BIAS; B_G

Symmetry has been well recognized as indicating an absence of bias in that if one party gets 50 + y% of the seats with 50 + z% of the votes, then so should the other party.[15] Correspondingly, symmetry in an S/V curve means that if there is a point on the curve at 50 + y% of the seats and 50 + z% of the votes, then there is another point on the curve at 50 – y% of the

[12]Subtracting the mean essentially performs the uniform shift near the center of the r/v plot where the uniform shift is valid.

[13]McDonald 2009 emphasized that total bias, including gerrymandering bias and also turnout bias, uses the statewide mean and that one can subtract the turnout bias by using the mean district %. We will use the statewide mean, thereby ignoring the small 0.3% turnout bias against the Democrats in the PA 2012 congressional election, and our model examples will assume no turnout bias.

[14]$B_{M\text{-}m}$ also fails the smoothness criterion for the example in Figure 4 because it jumps from +10% for x > 0 to −10% for x < 0. Just as the B_{GK} measure ameliorates the violation of the smoothness criterion for the B_S measure, an average that includes not just the median seat but additional seats on either side of the meridian seat would ameliorate this violation of the smoothness criterion.

[15]The ratio of y to z is a measure of responsiveness. The limit for small z gives the derivative in the S/V curve at V = 50%. The bilogit ρ gives exactly this value when λ is zero and ρ becomes only 30% larger as λ increases from 0 to 1.

Table 1. Percentage Values Obtained by Bias Methods Applied to Examples in Second Section

Examples / Methods		*(λ,ρ) Bilogits*				*PA 2012*		*(x) Fig. 3*			*(x) Fig. 4*		
		(½,3)	(1,3)	(½,8)	(1,8)	S/V	r/v	(±0)	(2)	(5)	(±0)	(5)	(10)
B_S	Simple 50%	12.3	23.1	12.3	23.1	17.8	22.2	±16.7	16.7	16.7	0	10.0	16.7
B_{GK}	Smeared 45–55%	11.9	22.6	10.3	19.9	17.6	20.8	0	6.7	16.7	0	10.0	16.7
$B_{M\text{-}m}$	Median-mean	4.1	8.3	1.6	2.9	7.8	7.8	0	2.0	5.0	±10	10.0	10.0
B_G	Geometric	7.9	15.6	3.1	6.3	9.4	13.7	0	2.7	6.7	0	8.0	13.3

seats and 50 – z% of the votes. This is called inversion symmetry about the (50,50) midpoint of the S/V graphical space. Clearly all the bilogits in Figure 1 with $\lambda = 0$ are symmetric, whereas none of the others are nor are any of those in Figs. 2, 3, and 4, although examples 3 and 4 become symmetric when $x = 0$.

While the symmetry test diagnoses whether there is bias, one needs to actually measure it. We show in this section how to quantitatively measure the asymmetry in S/V curves geometrically. Our procedure is to first invert the original S/V curve. This maps each (S,V) point on the original S/V curve to a point at (100 − S, 100 − V). For example, Figure 5 shows two bilogit S/V curves and their inversions.[16] Then, we propose that bias be measured as the geometric area between an S/V curve and its inversion. We further divide by the total number of seats so bias is given as a percentage, thereby giving a value appropriate for comparing different states with different numbers of seats. Figure 6 illustrates the method for obtaining B_G from the S/V curve for PA shown in Figure 2.[17] Figure 7 shows the area of bias for one of the curves in Figure 3. Values of B_G for selected values of the examples are given in the last column of Table 1.

A criticism of using the tails of the r/v distribution is that often those elections are uncontested. This concern is alleviated somewhat by the way we construct the S/V curve, so that even uncontested districts appear at V = 25% and 75%. One could also weight each district by its actual vote, thereby applying a turnout correction.

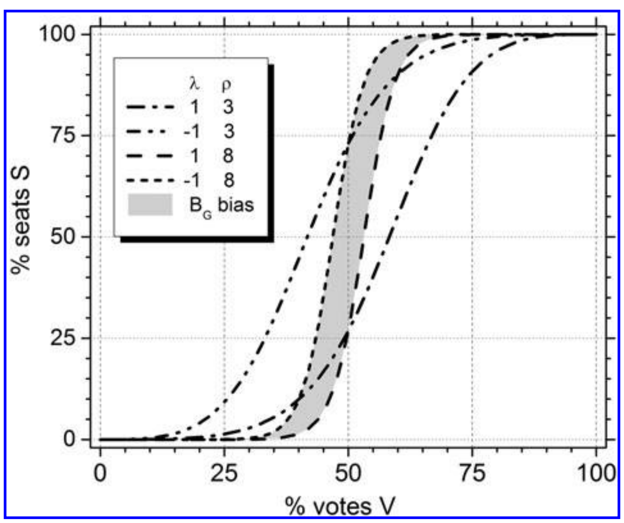

FIG. 5. Two bilogit S/V curves with $\lambda = 1$ and two values of ρ, each with its symmetric inversion that has the same value of ρ with $\lambda = -1$. The *shaded* area illustrates the proposed B_G measure of bias for the $\rho = 8$, $\lambda = 1$ case. Its value is 6.3% of the total area of the plot and changes sign for the $\lambda = -1$ case. For the $\rho = 3$ case, the B_G bias (not shown shaded) is 15.6%.

5. MEASURES BASED ON PARTISAN SATISFACTION (B_L, B_W, B_F)

The previous two sections focused on measures based on symmetry. This section considers measures based on a different kind of principle.[18]

As voters are happier when their candidate wins, this suggests that maximizing the total number of winning voters or minimizing the total number of losing voters would be a desirable feature of a voting system. However, this would lead to deliberate polarization by placing all like-minded voters in the same districts. In the limit where contiguity

[16]If the original bilogit curve has parameters (λ,ρ), then the inverted curve is a bilogit curve with parameters (-λ,ρ) so the original bias has been reversed from one party to the other. This is a general feature of all inverted S/V curves.

[17]The computation is simple. The % vote in the S/V Democratic data for seat number s is reassigned to seat number 18-s to obtain the S/V Republican curve. To obtain B_G in a spreadsheet, one just reverses the order of the votes column, subtracts from 100%, calculates the difference in the two columns, sums the absolute values, and averages. When this procedure is applied to an r/v curve, it should first be uniformly shifted to statewide v = 50%.

[18]However, it has also been argued that the median-mean $B_{M\text{-}m}$ method is more focused on fairness to the voter (McDonald and Best 2015) which draws it closer to the principles that motivate the measures in this section.

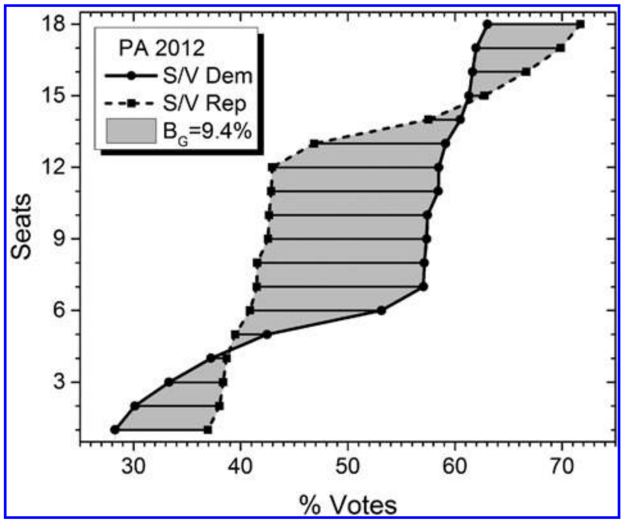

FIG. 6. The S/V curve for Democrats from Figure 2 is shown by a *solid line* and the *dashed line* shows its symmetrical inversion which is the S/V curve for Republicans. The B_G measure of bias is the *shaded* area between the two curves. It is computed to be 9.4% from the average of the length of the *horizontal black lines*.

and compactness could be ignored, idealizing voter happiness would create districts that would be 100% likely to vote one way, so this would be a very unresponsive system, requiring an enormous swing in votes to alter the percentage of seats away from the one used in the last districting. Most reformers consider the lack of competitiveness undesirable, although the opposite view has been persuasively argued by Buchler 2011, who notes that maximizing happiness is one of the many arguments that support his position. If 100% packing could be assured, this system would be proportionately representative, but it fails in other cases. For example, if V = 60% and the maximum attainable packing is 60%, then each district would have to be 60/40 and the majority party would likely win all the seats. While this is fair in that the opposite occurs when V = 40%, it certainly decreases the average happiness of minority voters. As this is not desirable, and as packing like-minded voters is geographically conflicted, and as such a system is so unresponsive to changes in statewide voter preference until a subsequent redistricting, let us not further pursue total voter happiness as a goal for districting.

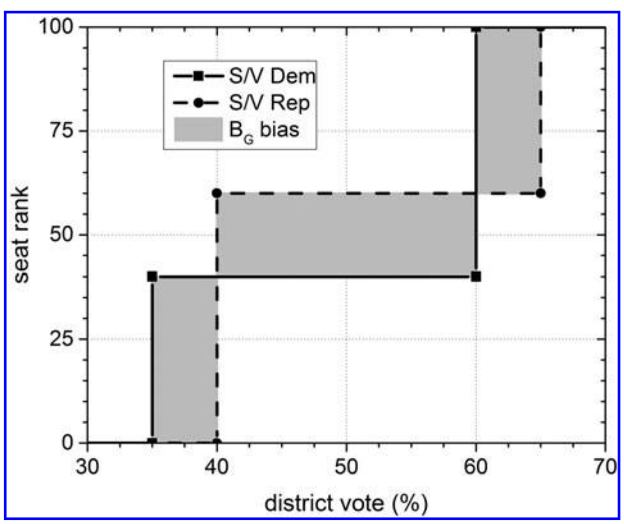

FIG. 7. The *solid line* shows the r/v plot for the simple model in Figure 4 with x = 5. The *dashed line* is the inverted plot and the *shaded* area is the B_G bias which equals 20%.

Nevertheless, the basic idea that a system should be most concerned with voters (Stephanopoulos and McGhee 2015; McDonald and Best 2015), essentially with their individual happiness and fair treatment, is appealing and leads to interesting measures of bias. The difference, compared to maximizing total voter happiness, is to use the concept in ways that also include partisan symmetry. First, let us consider equalizing the statewide percentage of unhappy, losing voters, L_D and L_R, in the D and R parties, respectively. The corresponding bias is then $B_L = L_D - L_R$. Interestingly, making B_L zero for all values of the statewide vote V_D leads to the proportional representation S/V curve. This is sufficiently important that a mathematical proof follows:

> Let v_n be the percentage D vote in district n. Then L_D is the sum of v_n over all districts n such that v_n is less than 50%. Similarly, L_R is the sum of $(100–v_n)$ over all districts n such that v_n is greater than 50%. The difference $B_L = L_D - L_R$ is the sum of v_n over all districts minus 100 times the number of districts won by D. The first number is just N times the percentage vote V_D, where N is the number of districts. The second number is just N times the percentage seats S_D, so $B_L = V_D–S_D$. Therefore, setting B_L to zero gives the proportional S/V curve $S(V) = V$.[19] Q.E.D.

The B_L measure finds bias in any deviation from proportional voting. Given an S/V curve, or even

[19]Of course, to compare B_L for different numbers of districts, one divides by N so that both S and V are represented by percentages.

just the r/v data from a single election, B_L is just $V - S$, which is the same as the simple B_S measure when $V = 50\%$.[20]

There is an important variant of the lost votes bias B_L that instead minimizes the difference in so-called "wasted" votes (McGhee 2014). Wasted Democratic votes W_D consist of the sum over all districts of the lost votes L_D and the surplus or excess votes E_D, where E_D is defined as the winning vote percentage in excess of 50%. A rationale for including excess votes is that they specifically focus on packing. As shown by McGhee 2014, equalizing W_D and W_R also leads to a particular S/V curve, namely, $S = 2V–50\%$ when $25\% < V < 75\%$, $S = 0$ when $V < 25\%$ and $S = 1$ when $V > 75\%$. This curve is shown by the dashed line in Figure 8. As the slope of this curve is 2 in the midrange of V, it is more responsive than proportionality and it is more realistic for the extreme values of V.[21] Let us designate the corresponding wasted votes measure of bias B_W as $W_D - W_R$.[22] This is easily calculated as $B_W = 2V - S - 50\%$. When $V = 50\%$, similarly to B_L, B_W is the same as the simple measure B_S. As with that measure, the smoothness criterion is also violated for example 3 by these measures. One important feature of both the B_L and the B_W measures is that, given an S/V plot, they generally obtain different values of bias for different V. Since the S/V plot purports to represent a districting plan, and since each plan should have a unique bias value, this raises the issue that these measures of bias and S/V plots are incompatible. Therefore, it has been asserted that one should only use r/v data to evaluate bias (Stephanopoulos and McGhee 2015).

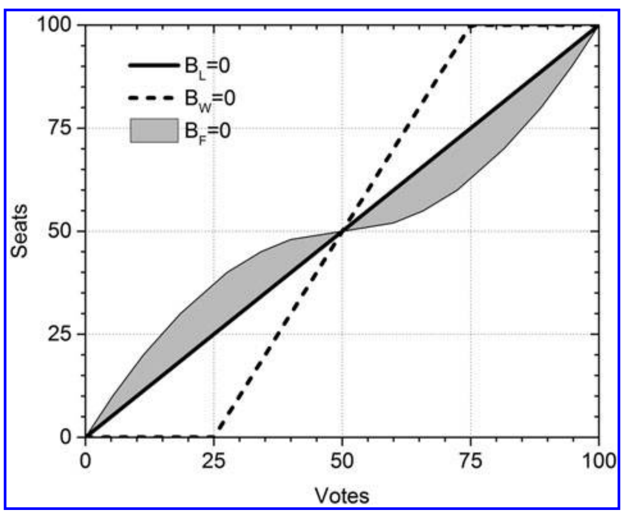

FIG. 8. The *bold solid straight line* is the proportional S/V plot given by $B_L = 0$. The *dashed line* shows the more responsive S/V plot given by $B_W = 0$. The *shaded* region shows a range of S/V combinations that can give $B_F = 0$ with different points at the same V corresponding to differences in the way the district votes v_n are distributed.

Our last measure of bias returns to equalizing happy voters in both parties. Although equalizing the number of winning ("happy") votes, $H_D = H_R$, did not work, as emphasized in footnote 20, an appealing variation, at least *a priori*, is to equalize the fractions of happy voters; let us designate these fractions as F_D and F_R. These are defined as the statewide number of winning votes for a party divided by the total number of votes for that party, so $F_D = H_D/V$ and $F_R = H_R/(1 - V)$ and the corresponding measure of bias will be designated by $B_F = F_R - F_D$. The appealing feature is that an R voter and a D voter have the same probability of being happy when $B_F = 0$. Interestingly, unlike the B_L and B_W measures, this one does not lead to a single ideal curve when $B_F = 0$. Instead there are areas of the S/V plot where one can obtain $B_F = 0$, as shown in Figure 8. In principle, this allows some flexibility not allowed by the $B_L = 0$ and the $B_W = 0$ criteria, and it shares with those measures that it can be applied when V is not close to 50%. Unfortunately, it requires packing an unrealistically large fraction of the majority party voters when V is much different from 50%. For example, for $V = 60\%$, the average fraction of majority voters in all its won districts must exceed 87%. Also, as Figure 8 shows, the ideal $B_F = 0$ region indulges a deeply minority party by giving it at least as many seats as proportional representation; for example, constraining $B_F = 0$, a districting plan for a state with ten districts would

[20]As equalizing losing votes leads to the popular, proportional representation, S/V curve, one might also ask what one is led to by equalizing winning votes. The same kind of analysis shows that setting the corresponding bias of winning votes to zero leads to $S(V) + V = 100\%$ which would idealize a bizarre S/V curve that would require a party to lose seats as it gains votes.

[21]If there are 100 seats, to obtain proportionality with 1% of the vote requires having at least half of the party's voters in one district. Generally, the slope of S/V curves must not exceed 2 when $V = 0$. The bilogit S/V curves violate this constraint when ρ is less than 1.

[22]This has recently been called the efficiency gap (Stephanopoulos and McGhee 2015).

ideally design a highly competitive district that would be won by the minority party with only $V = 5.3\%$ of the projected vote while giving it 10% of the seats.

6. COMPARISON OF MEASURES OF BIAS

Let us first compare the values shown in Table 1 given by the different bias measures. Although only four measures are shown, the three happiness measures, B_L, B_W, and B_F have essentially the same values as B_S for these particular cases. For the four bilogit examples shown in the table, both B_S and B_{GK} track the parameter λ that was used to construct the bilogit curves and that was originally described as their bias (King and Browning 1987). Of course, a normalization factor is always required to compare any two measures of bias, and a factor of about 25% applied to λ makes it numerically similar to B_S and B_{GK}.

In contrast, the $B_{M\text{-}m}$ measure tracks both the responsiveness parameter ρ and the putative bias parameter λ; indeed, $B_{M\text{-}m}$ nearly equals B_S/ρ Supposing that λ and B_S give the true measure of bias would therefore suggest that the $B_{M\text{-}m}$ measure is contaminated with responsiveness. Instead, the true bias should decrease when ρ increases because more seats become competitive and therefore statistically more likely to shift with a statewide vote shift. Partisan districting is more effective when the margin is not too thin (Owen and Grofman 1988), so the bilogit curve with $\rho = 3$ and $\lambda = 1$ in Figure 1 should have a smaller bias than the one with $\rho = 8$ and $\lambda = 1$ because the latter has about a third of its seats in the dangerous 50%–55% range, nearly three times as many as the less responsive bilogit.[23]

As previously mentioned, the basic difference between the B_S and the $B_{M\text{-}m}$ measures is that B_S measures the distance between the S/V curve and the symmetrical midpoint ($V = 50\%$, $S = 50\%$) along a vertical line and $B_{M\text{-}m}$ measures it along a horizontal line. As the underlying bias involves both directions, the product of B_S and $B_{M\text{-}m}$ would afford a combined measure that would correspond to an area near the center of the S/V graph.[24] The new B_G measure calculates an area in the S/V graph, but that area is not limited just to the center of the graph but includes the tails where the uncompetitive and packed districts reside. For the bilogits in Table 1, B_G tracks very well with the $B_{M\text{-}m}$ measure, after applying a normalization factor of 0.5.

Turning now to comparing these measures of bias for the PA 2012 election results in Table 1, B_S and B_{GK} give nearly the same results, primarily because the S/V curve is so uncompetitive that the averages between 45%–55% employed by B_{GK} make little difference. The difference between using r/v data and the S/V construct is simply due to the lumpiness in the r/v data compared to the smooth S/V construct, as can be seen in Figure 2. In contrast, lumpiness does not lead to a difference in the $B_{M\text{-}m}$ values because the S/V curve necessarily passes near the median seat vote due to the pileup of safe, but not too safe, R districts near $V = 60\%$. Unlike the bilogit examples, $B_S/B_{M\text{-}m}$ can not be construed to be the competitiveness because the slope of the graph in Figure 2 varies so rapidly with V, ranging from a highly uncompetitive value for the pertinent range near $V = 50\%$ to a very large value near 60%. Like the B_S measure, the B_G measure is also smaller for the S/V construct than for the r/v uniformly shifted data, but for a different reason. The S/V construct shortens the tails of the curve as seen in Figure 2 and this simply decreases the area between the S/V curve and its symmetrically opposite curve. Interestingly, the ratio $B_{M\text{-}m}/B_G$ using the r/v data remains reasonably close to 0.5 as it is for the bilogit examples.

Geometrically, the difference between B_S, B_{GK}, and B_G is the width of the area between the S/V plot and its symmetrical inverse used in the respective calculations of the different measures of bias. B_G uses all that area, B_{GK} uses only the portion between 45% and 55% of the vote, and B_S uses only an infinitely narrow strip at $V = 50\%$. Like B_S, the median-mean $B_{M\text{-}m}$ also uses only a narrow strip, but in the horizontal direction at $S = 50\%$ instead of in the vertical direction at $S = 50\%$. Possible perils of using narrow strips are revealed by the examples in Figures 3 and 4. $B_{M\text{-}m}$ successfully treats our

[23]Even when the bias λ is not zero, in the limit of infinite ρ the bilogit curve becomes a vertical step at $V = 50\%$, which is the unbiased, totally competitive, winner-take-all case. This further implies that the putative bias parameter λ is not proportional to the true bias.

[24]We will bypass this hitherto unproposed measure of bias, simply noting that it can be calculated as proportional to λ^2/ρ for the bilogits.

example in Figure 3 but it fails the packing criterion for our example 4 in Figure 4.[25] The seats based, narrow slice, measures B_S, B_L, B_W, and B_F successfully treat example 4 in Figure 4 but fail the smoothness criterion for example 3 in Figure 3. Only B_G passes both tests. While the existence of counterexamples and mathematical logic would eliminate all the other measures of bias, we do not advocate such a harsh conclusion. Counterexamples are necessarily special cases that may not represent the typical plan.[26] Nevertheless, counterexamples can be cautionary regarding cases where one should not apply one's favorite measure of bias. Example 3 emphasizes that care must be taken for the seats based measures when there are only a few districts with even fewer competitive seats and example 4 emphasizes that care must be taken for the $B_{M\text{-}m}$ measure when a party manages to secure a majority of districts with the r/v pattern shown for PA 2012 in Figure 2.[27]

One aspect of comparing the different measures that is not adequately contained in Table 1 is the apparent advantage the B_L, B_W, and B_F measures have when the statewide vote deviates substantially from $V = 50\%$. For example, a party with $V = 70\%$ could obtain all the seats if it can obtain a bilogit S/V curve with $\lambda = 0$ and $\rho = 8$.[28] This is a perfectly symmetrical curve that has no bias, but the high responsiveness, approaching the winner-take-all curve, gives 99.9% of the seats to the majority party at $V = 70\%$. Minimizing the lost votes bias B_L, by achieving proportionality, would seem fairer, but this becomes difficult to implement when V increases further. As noted before, a similar problem arises for B_F. This problem is avoided by choosing instead to minimize the wasted votes bias B_W and this has the advantage to those concerned more with competitiveness of being more responsive to swings in the statewide vote. However, this measure would freeze out parties that have less than 25% of the vote. While these are very real considerations, it must be emphasized that these issues are concerned with competitiveness/responsiveness, not partisan bias. For partisan bias, the B_L and B_W measures should be applied only to the $V = 50\%$ point on the S/V curve, at which they agree with the simple B_S measure. From this point of view, applying the B_L or B_W measures when V is not equal to 50% attempts to do too much in addition to assessing partisan bias.

It is easy to calculate B_F from the PA data in Figure 2. For that election only 45% of the D voters voted for a winning candidate while 87% of the R voters did. In other words, R voters in PA were nearly twice as likely to be happy with their representative as D voters. Subtracting their percentages yields a bias $B_F = 42\%$ favoring the Republicans. In percentage terms, this is the most sensitive measure of bias for PA. For the examples 3 and 4 in Table 1, B_F is exactly twice B_S. More r/v details would have to be provided to calculate B_F for the S/V bilogits.

Comparing results using r/v versus S/V is of interest because r/v are direct data, whereas the S/V curve requires a model and a layer of intervening manipulation. It has been argued forcefully that such manipulations should be avoided and that one should use unmodified r/v data exclusively. (Stephanopoulos and McGhee 2015, McDonald 2009, McDonald and Best 2015, Wang 2015) This argument favors the happiness measures in Section 5 and the median-mean measure $B_{M\text{-}m}$ in Section 3. However, the seats based r/v values for PA 2012 in Table 1 are only modestly larger than the S/V values.[29]

7. DISCUSSION

There are several aspects regarding inclusion of partisan fairness in the law. The first aspect distinguishes between making it justiciable, in order to challenge blatantly biased districting plans on the one hand and, on the other hand, including it in legislation and/or state constitutional amendments in order to avoid having biased plans in the first place. In either case, conciseness would be helpful. As Grofman and King 2007 have focused on the judiciary, this discussion concentrates on legislation.

[25]For an additional example, suppose half the districts have 50% R, 1/3 have 65% R and 1/6 have 20% R. Then, the statewide vote is 50% and the median seat also has 50% vote, so $B_{M\text{-}m}$ would therefore be zero whereas the expectation is for R to obtain 58% of the seats. Importantly, the r/v plot and its symmetrical inverse are identical in the central (50,50) region and the asymmetry only appears in the tails, thereby supporting the B_G measure.

[26]There are also subsequent considerations that preclude discarding all the other measures.

[27]For example, if neighboring districts 14 and 12 were slightly rearranged in Figure 2, 14 would be less strongly packed and 12 would shift to Democratic, but $B_{M\text{-}m}$ would not change.

[28]This curve is halfway between the two $\rho = 8$ curves in Figure 5.

[29]The S/V and r/v results would have been exactly the same if a smooth interpolation had not been made for the S/V curve in Figure 2. Furthermore, the r/v results are the same whether a uniform shift is applied or not.

Of course, to reform existing election law, it is most important that partisan bias simply be recognized as a restrictive criterion in state constitutions. This could be done just with a phrase in a constitutional amendment. Unfortunately, that could allow it to be practically disregarded just as compactness has been largely disregarded in the past, so something more might be contemplated. Reform could also come through legislative statute; that has the advantage of being easier to change as conditions change, but also easier to weaken for partisan advantage. Reform could also come through an independent districting commission. Such a commission would likely be more inclined to include partisan fairness in its deliberations if there is at least a phrase in the constitution mandating it, and it would be largely prohibited from doing so if political data, like election returns or voter registration, were constitutionally forbidden, as some reformers advocate.[30] An independent commission might also find it helpful if a clear set of guidelines is worked out by scholars.

The second aspect regards how measures of bias should be applied. In their important paper Grofman and King 2007 discuss five possibilities: (1) “require plans with as little partisan bias as practicable,” (2) “disqualify plans with partisan bias that deviate from symmetry by at least one seat,” (3) “disqualify only those plans with egregious levels of partisan bias (defined in terms of a specified percentage threshold),” (4) “disqualify only those plans that (can be expected to) translate a minority of the votes into a majority of the seats,” and (5) “disqualify only those plans whose partisan bias is both severe and greater than that in the plan being replaced.” Although Grofman and King 2007 discussed these in the context of justiciability, these are also possibilities for the legislative approach. All except number (4) require a quantifiable measure of bias. Possibilities (2) and (3) also address a third aspect, namely, how much partisan bias should be allowed.

The fourth aspect of including partisan fairness regards measures of bias to be used.[31] That has been the focus of this article, and several different measures have been presented and compared to each other. Recapitulating, when applied appropriately to the S/V or the shifted r/v curves, all the viable measures essentially agree.[32] Apparent divergences that occur when the B_L (proportional representation) and the B_W (super-proportional representation) are applied to statewide votes that deviate considerably from 50% are due to these measures also attempting to establish a standard for competitiveness. As noted in the introduction, that is not an easy issue; it should not impede going forward with the issue of partisan fairness. In the remainder of this section, the different bias measures will be discussed in relation to the other three aspects of inclusion of partisan fairness in the law.

Conciseness is an advantage when one considers that inclusion of any measure of bias in legislation would already have to involve quite a lot of nontrivial language if it is to be spelled out in detail. Any of the happiness measures have the merit that they can be more concisely stated than the others because they can be applied directly to r/v data and their formulas are simple. The $B_{M\text{-}m}$ formula is a bit more complicated in that one has to define which mean vote to use.[33] While the formula for B_S is simple, it should not be applied directly to r/v data when V is not 50%, but simply stating that it should be applied at V rather than at 50%.[34] Language describing the formula for the new B_G measure would be unwieldy, although it would at least be allowed if it were simply stated that symmetry should be applied. Equally for all measures of bias, there is the consideration of what kind of r/v data to obtain (see footnote 6). These choices would be tempered depending upon whether they

[30]However, it may be noted that Belin et al. 2011 propose an interesting way to promote partisan fairness that does not require political data, only population density data.

[31]Grofman and King 2007 essentially assume that bias would be measured using JudgeIt software (Gelman, et al. 2012). This would be a good choice, although it would probably be considered inappropriate to name it in legislation. It may also be noted that JudgeIt is apparently not designed to reliably obtain the S/V curve over the full range of V from 0 to 100%, so it is incompatible with the B_G measure. As mentioned at the end of Section 6, it might be simpler just to use the shifted election data for measuring partisan bias, and JudgeIt does seem to involve a uniform shift. Of course, JudgeIt is an invaluable statistical tool for measuring many other election elements and for including other kinds of data, especially incumbency. In this latter regard, it may also be noted that the new incumbent in the 12th PA district in 2014 received 3.8% more than the average R vote in 2014 and 2.6% more in 2012, consistent with a 1.2% incumbency advantage based on a very small sample.

[32]The B_G and $B_{M\text{-}m}$ measures can be made to agree better with the others if they are normalized by factors of 1.5 and 3 respectively.

[33]See footnote 13.

[34]This is formally the same as a uniform shift, but for the application of B_S, the more sophisticated adjustment in Appendix A makes little difference.

would be embedded in a constitution, in legislation, or just in guidelines for redistricting commissions. Codifying some choice would, of course, facilitate litigation to challenge a proposed or adopted plan.

Regarding the second aspect of how to apply a measure of bias for possibilities (1,2,3,5, listed above) of Grofman and King 2007, none of the viable measures is any less preferred than the others. Also regarding the third aspect of how much bias to allow, the only distinction is that any percentage allowed by B_G should be set smaller than the percentages allowed by B_S, B_{GK}, B_L, and B_W, and the percentages for B_F should be set higher.[35] Possibility (2) of Grofman and King 2007 above requires that the allowable bias would be smaller for a body with more seats. For PA with 18 congressional seats, it would constrain bias to 5.5% using the B_S, B_{GK}, B_L, and B_W measures, to about 4% using the B_G measure, and to about 2% using the $B_{M\text{-}m}$ measure.[36]

A final issue then is whether a percentage bias number should be embedded in the state constitution or whether it is best done by legislative statute. While reform would be more robust if it could be embedded in the constitution, the language to cover all the contingencies could be unwieldy and, to ensure passage, the bias range may be set too high and then not be easy to change. This consideration favors possibility (1) above, but the clause "as practicable" could vitiate implementation if it were claimed that other factors make fairness not practicable. As usual, reform law would require political deliberation, but that and having a choice of possible measures of bias should not dissuade reformers from pressing the issue of partisan fairness in election law.

REFERENCES

Amy, Douglas J. 2000. *Behind the Ballot Box: A Citizen's Guide to Voting Systems*. Westport, CT: Praeger Publishers.

Backstrom, C., L. Robins, and S. Eller. 1990. "Establishing a Statewide Electoral Effects Baseline." In *Political Gerrymandering and the Courts*, B. Grofman, ed. New York, NY: Agathon Press. 145–70.

Baker, G. E. 1990. "The 'Totality of Circumstances' Approach." In *Political Gerrymandering and the Courts*, B. Grofman, ed. New York, NY: Agathon Press. 203–11.

Belin, T. R., H. J. Fischer, and C. M. Zigler. 2011. "Using a Density-Variation/Compactness Measure to Evaluate Redistricting Plans for Partisan Bias and Electoral Responsiveness." *Statistics, Politics, and Policy* 2.

Buchler, Justin. 2011. *Hiring and Firing Public Officials: Rethinking the Purpose of Elections*. New York, NY: Oxford University Press.

Chen, J. W. and J. Rodden. 2013. "Unintentional Gerrymandering: Political Geography and Electoral Bias in Legislatures." *Quarterly Journal of Political Science* 8: 239–69.

Gardner, J. A. 2012. "How to Do Things with Boundaries: Redistricting and the Construction of Politics." *Election Law Journal* 11: 399–419.

Gelman, A. and G. King. 1990. "Estimating the Electoral Consequences of Legislative Redistricting." *Journal of the American Statistical Association* 85: 274–82.

———. 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*. <http://gking.harvard.edu/judgeit>.

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.

Hirsch, S. and D. Ortiz. 2005. *Beyond Party Lines: Principles for Redistricting Reform*. Alexandria, VA: Reform Institute. <http://kms1.isn.ethz.ch/serviceengine/Files/ISN/90778/ipublicationdocument_singledocument/A20DD06E-DB8E-405B-BB8E-8E504CFE1B96/en/2008-06-26_Principles_for_redisctricting_reform.pdf>.

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.

Leach, Daylin. 2014. "Dismantling Democracy: How Gerrymandering Has Stolen Your Voice and Undermined Our Political System." *Roxbury News*, March 1. Available at <http://www.roxburynews.com/index.php?a=6830>.

[35]Different levels for different measures is analogous to obtaining different numbers when measuring the same length with a meter stick or a yard stick, so this is not a valid criticism against the quantification of partisan bias.

[36]One concern could be that political geography would not allow any plans to fall in the allowed bias range. This occurred with the districts generated by Chen and Rodden 2013 for Florida. For PA that concern might be moot as it has been asserted that a districting plan has been made that would reverse the 2012 result, giving the Democrats 13 seats if all voters would have voted for candidates of the same party (Leach 2014).

Mac Donald, K. 2012. "Adventures in Redistricting: A Look at the California Redistricting Commision." *Election Law Journal* 11: 472–89.

Mann, T. E. and N. J. Ornstein. 2013. "Finding the Common Good in an Era of Dysfunctional Governance." *Daedalus* 142: 15–24.

Mattingly, J. C. and C. Vaughn. 2014. "Redistricting and the Will of the People." *arXiv* 1410.8796v1 [physics.soc-ph] 29.

McDonald, M. P. 2014. "Presidential Vote within State Legislative Districts." *State Politics and Policy Quarterly* 14: 196–204.

———. 2007. "Regulating Redistricting." *Ps-Political Science and Politics* 40: 675–79.

McDonald, Michael D. 2009. "The Arithmetic of Electoral Bias, with Applications to U.S. House Elections." *APSA 2009 Toronto Meeting Paper*. Available at SSRN: <http://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: preprint.

McGhee, E. 2014. "Measuring Partisan Bias in Single-Member District Electoral Systems." *Legislative Studies Quarterly* 39: 55–85.

Owen, G. and B. Grofman. 1988. "Optimal Partisan Gerrymandering." *Political Geography Quarterly* 7: 5–22.

Stephanopoulos, N. O. and E. M. McGhee. 2015. "Partisan Gerrymandering and the Efficiency Gap." *University of Chicago Law Review* 82: 831–900.

Wang, S. S. H. 2015. "A Three-Pronged Standard for Practical Evaluation of Partisan Gerrymandering." Preprint. Available at SSRN: <http://ssrn.com/abstract=2671607>.

Address correspondence to:
John F. Nagle
Department of Physics
Carnegie Mellon University
Wean Hall
Pittsburgh, PA 15213
E-mail: nagle@cmu.edu

APPENDIX A: CONSTRUCTION OF S/V CURVES FROM VOTING DATA

Given N districts, labeled $n = 1, \ldots, N$, and given Democratic (D) past (or projected) vote v_{nD} in current (or proposed) district n, the statewide D vote is $V_D = (1/N)\sum v_{nD}$ where $\sum$ connotes the sum over all districts. First consider a statewide shift from D to R (Republicans) and denote the new statewide vote as $V_D' = V_D - \delta_D$, where δ_D is the average percentage of shifted Democrats statewide. We assume that there is an equal probability x_D for any Democrat in any district to shift,[37] so the new district vote is $v_{nD}' = v_{nD}(1 - x_D)$. Summing v_{nD}' over all districts gives $V_D' = V_D - x_D V_D$; this identifies $x_D = \delta_D/V_D$. Then, the shift δ_{nD} in statewide vote that is necessary to shift a seat n that was won by the Democrats must satisfy $v_{nD}' = 50\%$ which gives $\delta_{nD} = + V_D - V_D/2v_{nD}$. Therefore, the statewide vote required to shift the nth seat from D to R is $V_{nD}' = V_D - \delta_{nD} = V_D/2v_{nD}$, as stated in the main text. This is the value of V that is plotted on the seats/votes (S/V) curve for the nth seat originally won by the Democrats. For seats originally won by the Republicans, the same derivation gives the symmetrically equivalent formula $V_{nR}' = V_R/2v_{nR}$ when the Republican vote v_{nR} exceeds 50%. For those seats, the value of V that is plotted on the S_D/V_D curve for district n is, of course, $100 - V_{nR}'$, corresponding to V_{nD}'.

It may be of some interest to view some simple r/v examples and their corresponding S/V curves. Figure A1 shows that uniform distributions of the votes by ranked seats result in S/V curves that all have $B_S = 0$. Uniform shifts of these r/v lines to $V = 50\%$ (not shown) are all symmetrical. Slight asymmetry is evident for the two S/V examples derived from r/v curves with statewide vote different from 50%, so the B_G measure is small but not zero, gradually reaching 5% as V increases to 75%.

In contrast to Figure A1, Figure A2 shows an example of an asymmetrical r/v curve where the first party wins only 25% of the seats with 37.5% of the vote. Nevertheless, this outcome is biased against the second party because the S/V curve shows that the second party would, on average, only win 38% of the seats when the statewide vote is 50% for a bias $B_S = 12\%$.

Finally, it may be of interest to view a comparison of the S/V curves for the 2014 and 2012 PA congressional elections. Figure A3 shows that three uncontested seats moved the largest and smallest

[37]The model assumes that only Democrats shift when there is a statewide shift towards Republicans. Supposing that Republicans also shift with a uniform probability leads to the unreasonable result that a massively Republican district would always become more Democratic even when the statewide vote was shifting in favor of the Republicans.

district vote percentages v_n to more extreme values. More interestingly, it also appears that the "wall" of Republican seats built up near $V_D = 60\%$ in 2012 may be moving into the more competitive range, even though those seats were comfortably held by the Republicans at the 60% district level in 2014 due to a 5.5% statewide swing. This would be consistent with studies that have reported gradual erosion in the effects of gerrymanders (McGhee 2014). This should not, of course, be used as an excuse to ignore gerrymandered partisan bias when redistricting, even if the effect only lasts for a few elections.

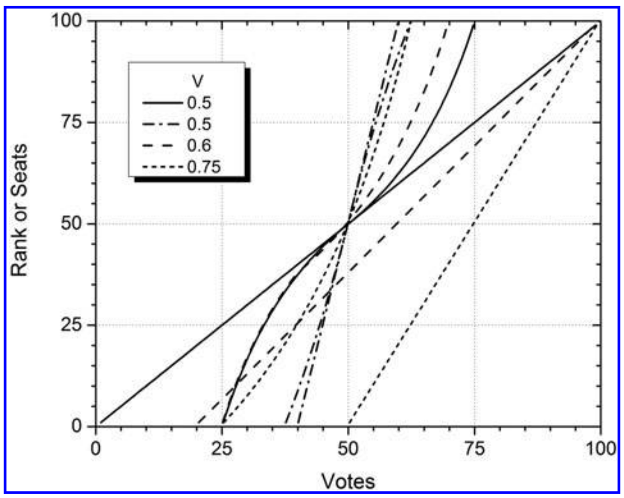

FIG. A1. Four simple r/v examples are shown by the *straight lines* with different line type as indicated in the legend. For the r/v lines the votes axis is the vote for the first party in each ranked seat. For each r/v line, the corresponding curved S/V plot has the same line type, the votes are the statewide votes required for the second party to obtain the number of seats on the S/V curve. For the last example, the first party received 75% of the statewide vote, which is the average of the vote in 100 districts starting with 50% for the seat won narrowly and moving up to the last seat which was won unanimously. All of these examples are unbiased using the B_S measure r/v, rank/votes; S/V, seats/votes.

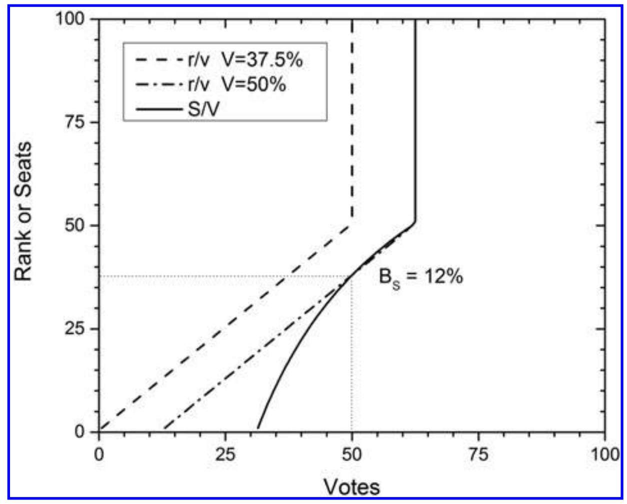

FIG. A2. A highly asymmetric r/v curve that gives 37.5% of the statewide vote to the first party is shown by the *dashed line*. That party loses ranked seats 1–50 and splits seats 51–100. The *dash-dot line* shows the shifted r/v curve to 50% statewide vote and the *solid* curve shows the S/V plot computed from the unshifted r/v *lines*. The S/V and the shifted r/v curves agree in the vote region centered at 50%.

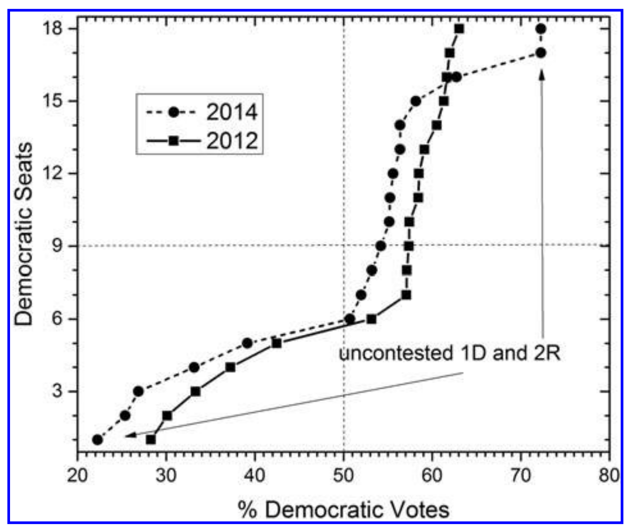

FIG. A3. Comparison of S/V curves for PA 2012 and 2014 congressional elections. Three uncontested seats, one Democratic and two Republican, are at the ends of the 2014 curve.