IMPROVING THE EFFICIENCY GAP

Cover image by: Cole Youngblood.

Jeffrey T. Barton

Partisan gerrymandering and how to measure it have been the subject of several recent court cases and a lot of media coverage. The *efficiency gap* in particular has garnered a lot of attention. The creators of the efficiency gap, Nicholas Stephanopoulos and Eric McGhee, claim that the measure is a "tally of all the cracking [spreading voters of a party across many districts to diminish their influence] and packing [stuffing voters of a party into few districts to concentrate their influence] decisions in a district plan" ("Partisan gerrymandering and the efficiency gap." *U. Chi. L. Rev.* 82, no. 2 [Spring 2015]: 831–900).

The efficiency gap is compelling in its simplicity, and it has produced reasonable results in practice. It has made its way to the courts as part of a considered, multifaceted approach to detect gerrymandering.

However, experts such as Mira Bernstein and Moon Duchin have pointed out that the measure has several "grave limitations" ("A formula goes to court: Partisan gerrymandering and the efficiency gap." *Notices Amer. Math. Soc.* 64, no. 9 [October 2017]: 1,020–1,024).

In this note we suggest a fix for many of these limitations—a metric that requires no new data or tools to compute, is no more complex than the original, fixes the problems, and is more in line with our intuition and with court findings on proportional representation.

The Efficiency Gap

In a two-party election, the efficiency gap measures the disparity between the wasted votes of parties A and B as a proportion of total votes. In any district, any vote in excess of the majority or any vote for a losing candidate is wasted. We denote the total number of wasted votes across all districts for parties A and B by W^A and W^B, respectively. Assuming N equal-sized districts of size D, the total number of votes in the election is $V = ND$. The efficiency gap is defined to be:

$$EG = \frac{W^A - W^B}{V}.$$

Note that $W^A + W^B = V/2$ and $-1/2 \le EG \le 1/2$.

If party A's votes have been unfairly packed or cracked by a given map, then A will waste more votes than B and EG will be positive. If party B is treated unfairly, B will waste more votes than A, and EG will be negative. If $EG = 0$ we have "perfect fairness."

For simplicity, we allow "fractional people" (we don't round to the nearest person), and we ignore ties. For example, we take $D/2$ to be the number of votes necessary to win a district despite the fact that if $D = 10{,}000$ we would need 5,001 votes, not 5,000. Neither assumption causes any real difficulty, and even though we ignore ties in our definitions, our measure handles them sensibly.

Table 1 presents the vote splits for 10 districts of 100 voters each. It gives the numbers of votes and wasted votes for A and B. For district 1, A wastes 10 votes because it needs only 50 to win, and B wastes all 40 of its votes for a losing candidate. The efficiency gap is $(146-354)/1000=-0.208$.

Stephanapolous and McGhee suggest $|EG|>0.08$ as the threshold for presumptive gerrymandering. Thus, the map behind this example would be considered an illegal gerrymander at the expense of B, which wins 45.4 percent of the vote but only 20 percent of the seats. So far, so good. This map seems to be the result of gerrymandering, and the efficiency gap has detected it.

However, it is easy to break the formula. Suppose there are eight districts, and A receives 90 votes to B's 10 in each district. Party A wastes $40\cdot 8=320$ votes and B wastes $10\cdot 8=80$, so the efficiency gap is $(320-80)/800=0.30$. Thus, it considers this map biased against party A—which won all eight seats.

The efficiency gap also has trouble with *granularity*; when there are few districts, each seat swing moves the fairness needle a lot. This can lead to problematic results even in competitive elections. Consider the following scenario: A wins three districts 51 to 49 and B wins two districts 51 to 49. So, A wastes $1+1+1+49+49=101$ votes and B wastes $49+49+49+1+1=149$. The efficiency gap, $(101-149)/500=-0.096$, reports a bias against B. But B is the minority party, so winning two seats is the best they should expect. To award B another seat would give the minority party a majority of seats.

Even worse is the case of a single district election in which the majority party A wins the only seat: The efficiency gap can flag the election as very unfair to A, very unfair to B, or anything in between, depending on the size of A's majority. Clearly this result is problematic.

Stephanopoulos and McGhee showed that the efficiency gap is related to the proportions of votes and seats won by a party. With V^A as the total number of votes received by A, and S^A the total number of seats won by A, we write $v^A=V^A/V$ and $s^A=S^A/N$ for the proportions of votes and seats won by A. Then the efficiency gap is $EG=2v^A-s^A-1/2$. (Prove it!)

In our first example, party A won 54.6 percent of the votes (v^A) and 80 percent of the seats (s^A), and indeed, $EG=2(0.546)-0.8-0.5=-0.208$.

Writing EG in this way reveals the relationship between votes and seats that the efficiency gap considers to be perfectly fair. Setting $EG=0$ yields the *seats-votes equation*: $s^A=2v^A-1/2$. Figure 1 shows the seats-votes line in red and the line of proportional representation ($s^A=v^A$) in green.

Table 1. Ten-district map with majority party A.

District	A votes	B votes	Winner	Wasted A votes	Wasted B votes
1	60	40	A	10	40
2	65	35	A	15	35
3	75	25	A	25	25
4	65	35	A	15	35
5	55	45	A	5	45
6	51	49	A	1	49
7	55	45	A	5	45
8	60	40	A	10	40
9	30	70	B	30	20
10	30	70	B	30	20
Total	546	454		146	354

Figure 1. The seats-votes line is red, and the red band shows $|EG|\le 0.08$. The green line indicates proportional representation.

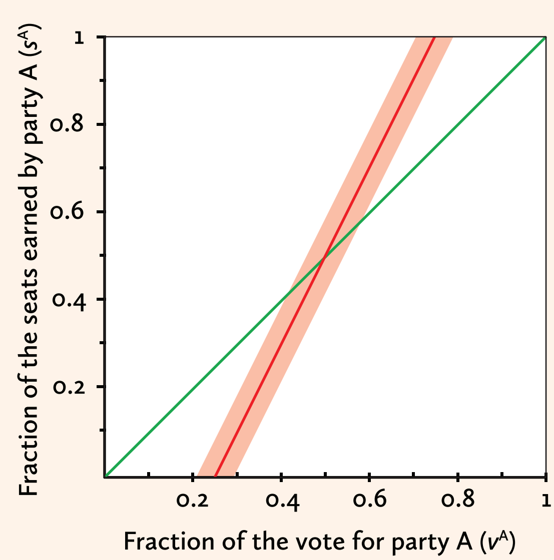

Notice that no map can achieve perfect fairness when $v^A < 0.25$ or $v^A > 0.75$. Thus, the fairness criterion is *impossible* to meet for half of the possible values for v^A.

We also see that the seats-votes line is quite different from the line for proportionality. Taking $|EG| \le 0.08$ (the red region in figure 1) as the requirement for fairness, the efficiency gap flags proportionality as unfair when a party receives less than 42 or more than 58 percent of the vote. Though the courts have held that proportionality is not a requirement for fairness, the efficiency gap rejects proportionality in a way that the courts have not requested (see Bernstein and Duchin).

Modifying the Efficiency Gap

Many of the issues with the efficiency gap stem from the fact that wasted votes are not the appropriate quantity to count if one is interested in unfairness due to gerrymandering. Such a count fails to take into account that each party must waste a certain number of votes *regardless of how districts are drawn*. If we know the number of districts our map must contain, then we can find the minimum number of votes that must be wasted by A and B before drawing any line. These are wasted votes that should be blamed on mathematics rather than any map, yet the efficiency gap takes no notice of them.

Suppose A is the majority party. It requires $D/2$ votes to win a district of size D. So A needs at least $DN/2 = V/2$ votes to win all N seats. Because A cannot win more than all of the seats, all votes for A above $V/2$ *must* be wasted. If A receives V^A votes, it wastes $M^A = V^A - V/2$ of them purely because of the mathematics.

The situation is different for B. The minority party does not have enough votes to win all districts. But, B can win its maximum number of districts by allocating $D/2$ votes to districts until it runs out of votes. Thus, B can win at most $\lfloor 2V^B/D \rfloor$ seats, where $\lfloor x \rfloor$ is the greatest integer less than or equal to x. Once B allocates its votes to win its maximum number of seats, all leftover votes are wasted because they must be placed in a district won by B or in one that B must necessarily lose. This means that

$$M^B = V^B - \frac{D}{2}\left\lfloor \frac{2V^B}{D} \right\rfloor.$$

Having defined M^A and M^B, we define our new measure of gerrymandering as the difference of the *unnecessarily wasted* votes as a proportion of total votes. We call the new measure the *modified efficiency gap*:

$$MEG = \frac{(W^A - M^A) - (W^B - M^B)}{V}.$$

Note that as with EG, $-1/2 \le MEG \le 1/2$.

Table 2 compares EG and MEG for the previous examples. In the first example, the two measures report similar values. In the second example, B received 10 percent of the vote but won no seats. The efficiency gap marks this as unfair to A, while the modified efficiency gap reports it as slightly unfair to B. This example highlights the fundamental difference between the two metrics. The efficiency gap ascribes all of the 320 wasted A votes and the 80 wasted B votes to this specific map, while the modified efficiency gap recognizes that due to its large majority, A *had* to waste those 320 votes while B had to waste 30 votes. Thus, MEG attributes only the remaining 50 wasted B votes to the map, and as a result it detects that the map was actually unfair to B.

Finally, in the third example we have a map that the efficiency gap considers to be very unfair to B, but the modified efficiency gap gives the more sensible response that the map is completely fair because it takes the granularity of the situation into account.

Like the formula for EG, the formula for MEG boils down to a seats-votes relationship:

$$MEG = 1 - s^A - \frac{1}{2N}\left\lfloor 2N(1 - v^A) \right\rfloor.$$

Table 2. The differences between *EG* and *MEG* for our three examples

Example	W^A	W^B	M^A	M^B	*EG*	*MEG*
1	146	354	46	4	−0.208	−0.25
2	320	80	320	30	0.30	−0.0625
3	101	149	1	49	−0.096	0

With $MEG = 0$ for "perfect fairness," the seats-votes equation for the modified efficiency gap is

$$s^{A} = 1 - \frac{1}{2N}\lfloor 2N(1 - v^{A}) \rfloor.$$

Figure 2 shows the seats-votes graph for eight districts and the $|MEG| \leq 0.08$ region in red. The green line is the line of proportionality. In contrast to the efficiency gap, the modified efficiency gap requires no impossible values for fairness. Also, the maximum departure of the seats-votes graph from proportionality is $\frac{1}{2N}$. So, it is more in keeping with what the courts have held regarding proportionality: It neither rejects nor requires it.

By acknowledging the granularity of districting in its definition, the modified efficiency gap is better able to handle cases of few districts. We saw this in the five-district example, but the difference is even more stark in the extreme case of a single district. As $2(1 - v^{A}) < 1$ when $v^{A} > 1/2$, $MEG = 0$. In contrast to the efficiency gap, the modified efficiency gap reports an election with one district in which the majority wins as perfectly fair.

The One-Liner

In every election it is possible for the minority party, B, to win no seats if the districts are drawn perfectly for A. On the other hand, B can win at most $\lfloor 2V^{B} / D \rfloor$ districts if they are drawn perfectly in its favor. We can think of the maximum number of districts that B can win as the number that are "in play" when drawing a map. Because the best map for B is the worst map for A and vice versa, the modified efficiency gap says that the fairest approach is to split the in-play districts equally between the two parties. Thus, we get the rather appealing one-line description: *A perfectly fair map is one that awards the minority party half of its maximum number of seats*. The formula for *MEG* measures the degree to which an actual map departs from this standard.

To verify this interpretation, multiply both sides of the seats-votes equation by N and simplify to obtain $S^{A} = N - \frac{1}{2}\lfloor 2V^{B} / D \rfloor$. Hence, $S^{B} = \frac{1}{2}\lfloor 2V^{B} / D \rfloor$.

In our five-district example, B could win four districts with a perfect B map, so we award B two districts for perfect fairness. In our one-district example, B can win no districts, so B is awarded none. In our eight-district example, B can win a maximum of one seat, so it is awarded half a seat. In such an impossible situation, the modified efficiency gap is indifferent to awarding zero or one seat to B. This is the best we can expect within the limitations imposed by the granularity of districting.

Figure 2. The eight-district seats-votes graph is red, and the red band shows $|MEG| \leq 0.08$. The green line indicates proportional representation.

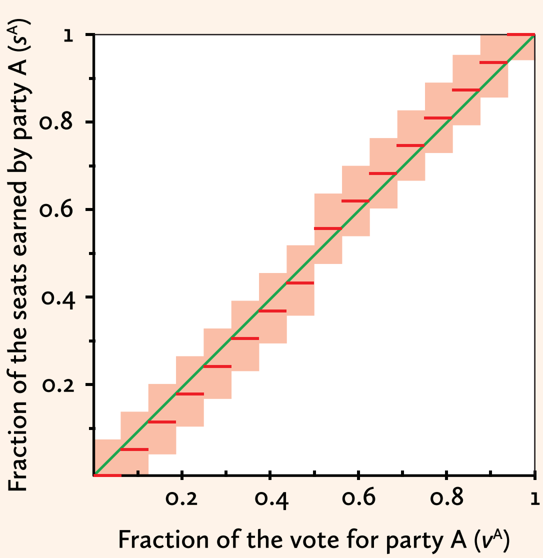

Which Party Benefits?

As a final comparison between the metrics, look at the red regions in figures 1 and 2 that correspond to $|EG| \leq 0.08$ and $|MEG| \leq 0.08$. We see that, in most cases, the efficiency gap is biased in favor of the majority party and that switching to the new metric could very well turn into seats for the minority party.

For instance, if there are 18 districts and A receives 55 percent of the votes, the efficiency gap would allow 10 to 12 seats for A, while the modified efficiency gap would only allow 9 to 11. The differences grow quickly as A's majority increases until around 70 percent when the acceptable ranges become disjoint.

When a party has a small majority, both metrics accept similar ranges for seats. This indicates one reason why the efficiency gap's conflation of votes wasted by math with votes wasted by gerrymandering has not yet surfaced as a problem in applications.

Jeff Barton *received his Ph.D. in mathematics from the University of Texas at Austin. He is a Project NExT fellow (Gold Dot '00) and teaches mathematics at Birmingham-Southern College.*

10.1080/10724117.2018.1468140