The Mathematics of Gerrymandering

Kevin Knudson

November 8, 2018

Department of Mathematics
University of Florida
kknudson@ufl.edu
http://kpknudson.com
@niveknosdunk

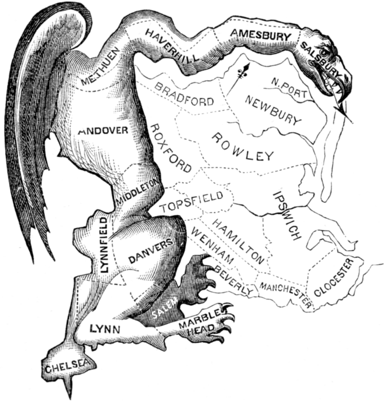

Outline

History

Outline

History

Constraints

Outline

History

Constraints

Algebra

Outline

Outline

History

Constraints

Algebra

Probability

Geometry

Outline

History

A Practice (Nearly) as Old as the Republic

Etymology: The first appearance of the word *Gerry-mander* was in the *Boston Gazette* on March 26, 1812, as part of an editorial excoriating Governor Elbridge Gerry's signing of a redistricting bill.

The district map for the Massachusetts state senate favored Gerry's Democratic-Republican Party, and one district near Boston looked a bit like a salamander when drawn on a map. Hence, the portmanteau Gerry-mander.

By the way, it worked. Gerry's party controlled the senate after the 1812 election, but they lost the house and governorship.

In the intervening 200+ years, politicians of all parties have employed these tactics.

The Constitution does not prohibit it, and in fact our system of government is susceptible to attempts to game the system. Countries that employ proportional representation systems generally do not have gerrymandering problems.

That's not to say it doesn't happen elsewhere. You can read about a lot of international gerrymandering at the Wikipedia page: `http://en.wikipedia.org/wiki/Gerrymandering`.

Modern Problems

Better computers and more accurate data gathering have turned what was once an art into a precision science, allowing legislators to surgically design districts to maximize their advantage. For example:

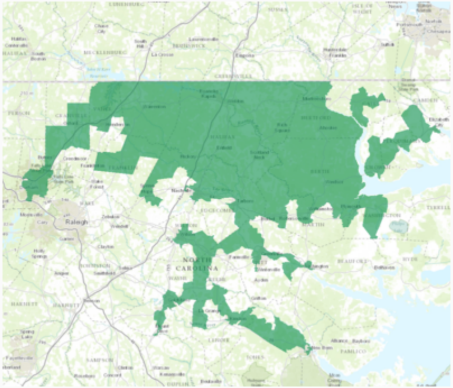

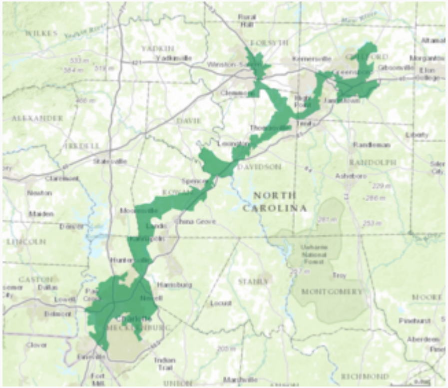

Constraints

What are the constraints on a districting map?

There are some statutory limitations:

- Apportionment Clause, Article I, Section 2 of the Constitution: Congressional districts must be as nearly equal in population *as practicable*.
- Equal Protection Clause, 14th Amendment: state legislative districts must be substantially equal.
- Voting Rights Act of 1965, Section 2: no plan may intentionally or inadvertently discriminate on the basis of race.

Traditional districting principles

Most states adhere to the following criteria when drawing district maps:

- Compactness
- Contiguity
- Preservation of counties or other political subdivisions
- Preservation of communities of interest
- Preservation of cores of prior districts
- Avoiding pairing incumbents

What about this district?

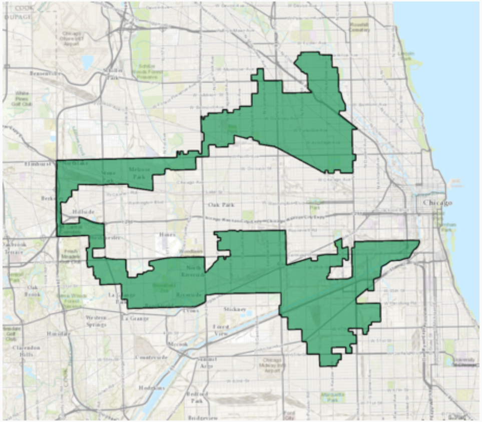

This is Illinois District 4. It looks bad. The western border runs along a stretch of I-294 to maintain contiguity, but on its face it seems suspicious.

What about this district?

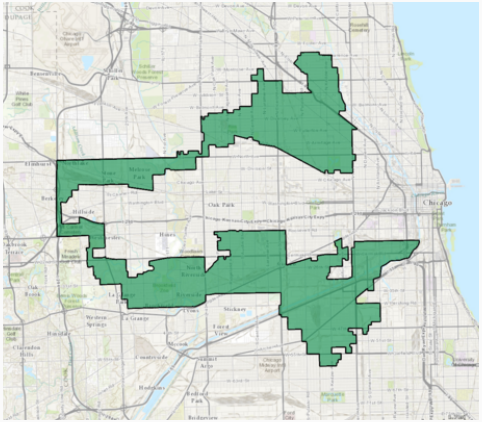

This is Illinois District 4. It looks bad. The western border runs along a stretch of I-294 to maintain contiguity, but on its face it seems suspicious.

Except it isn't. The two flares of the earmuffs consist of Hispanic neighborhoods in Chicago. It was created by federal court order in the early 1990s to create a majority Hispanic district.

Algebra

Packing and Cracking

Gerrymandering usually comes down to some combination of the following strategies.

Packing and Cracking

Gerrymandering usually comes down to some combination of the following strategies.

Packing. In this scheme, members of one group are concentrated into as few districts as possible. The majority then sacrifices a few seats but wins all the rest.

Packing and Cracking

Gerrymandering usually comes down to some combination of the following strategies.

Packing. In this scheme, members of one group are concentrated into as few districts as possible. The majority then sacrifices a few seats but wins all the rest.

Cracking. This is the opposite approach: divide a group among as many districts as possible to dilute their votes.

Three districting plans

Here is a simple state with 42 voters to be divided into 6 districts of 7 voters each. There are 21 voters in each "party" (red and blank).

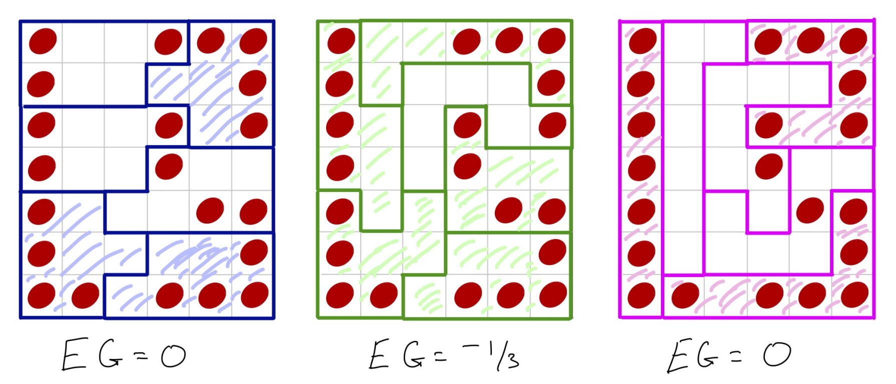

adapted from [Bernstein and Duchin, 2017]

The first plan is pretty fair. Each party wins three seats and the votes are fairly evenly distributed. The third plan acomplishes the same end, but notice that the districts are all packed so that the winning party does so overwhelmingly. The second, however, yields 5 wins for the red party.

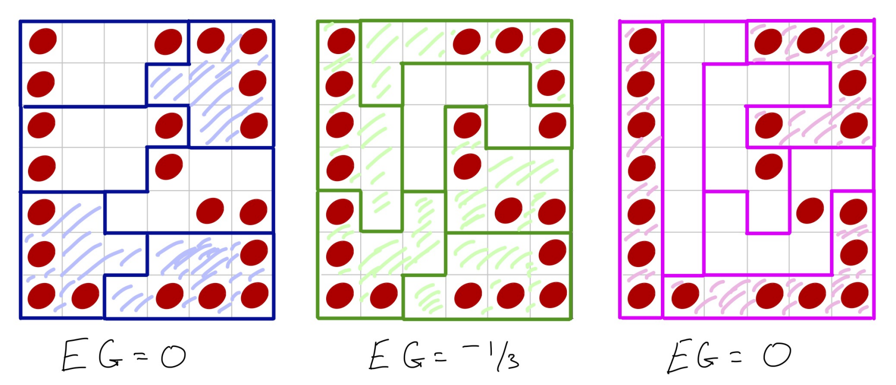

The Efficiency Gap

We'd like a way to measure packing and cracking numerically. The idea is to count *wasted votes*. In a majority-rules two-party system, the winner needs to get $50\% + 1$ vote; anything beyond that is wasted in some sense. Getting 80% of the vote might be good for a candidate's ego, but it does nothing to affect the outcome.

The Efficiency Gap

We'd like a way to measure packing and cracking numerically. The idea is to count *wasted votes*. In a majority-rules two-party system, the winner needs to get $50\% + 1$ vote; anything beyond that is wasted in some sense. Getting 80% of the vote might be good for a candidate's ego, but it does nothing to affect the outcome.

The *efficiency gap* is a mechanism to measure this.

Some definitions

Let's say we have two parties, A and B, and there are S legislative districts. Denote the set of districts by $\mathcal{D} = \{d_1, \ldots, d_S\}$. For a particular election, let $\mathcal{D}^P$ be the subset of districts won by party P. Let S_i^P be 1 if party P won district i and 0 otherwise; it follows that $S^P = |\mathcal{D}^P|$ is the number of seats won by P. Let T_i^P be the number of votes cast in district i for party P so that T^P denotes the total number of statewide votes for P.

Some definitions

Let's say we have two parties, A and B, and there are S legislative districts. Denote the set of districts by $\mathcal{D} = \{d_1, \ldots, d_S\}$. For a particular election, let $\mathcal{D}^P$ be the subset of districts won by party P. Let S_i^P be 1 if party P won district i and 0 otherwise; it follows that $S^P = |\mathcal{D}^P|$ is the number of seats won by P. Let T_i^P be the number of votes cast in district i for party P so that T^P denotes the total number of statewide votes for P.

Now define

$$\tau = \frac{T^A - T^B}{T} \qquad \sigma = \frac{S^A - S^B}{S}$$

to be the *vote lean* and *seat lean* for party A, respectively.

Proportional representation

Most people have an intuition about what is fair, and it would seem logical that

$$\sigma \approx \tau$$

is a desirable outcome. That is, parties should be represented *proportionally* to vote share.

Proportional representation

Most people have an intuition about what is fair, and it would seem logical that

$$\sigma \approx \tau$$

is a desirable outcome. That is, parties should be represented *proportionally* to vote share.

Our electoral systems *do not* guarantee this. In fact in a 1986 ruling (Davis v. Bandemer), the U.S. Supreme Court wrote that "the mere lack of proportional representation will not be sufficient to prove unconstitutional discrimination."

One might say that the ultimate goal of gerrymandering is to undermine proportionality by delivering the winning party more seats in the legislature than their proportion of the vote.

Wasted Votes

A *wasted vote* is one of the following:

- a vote cast for the losing side; or
- a vote cast for the winning side in excess of the 50% needed to win.

Wasted Votes

A *wasted vote* is one of the following:

- a vote cast for the losing side; or
- a vote cast for the winning side in excess of the 50% needed to win.

Thus, the number of votes wasted by A-voters in district d_i is

$$W_i^A = \begin{cases} T_i^A - T_i/2, & d_i \in \mathcal{D}^A \\ T_i^A & d_i \in \mathcal{D}^B \end{cases} = T_i^A - S_i^A \cdot \frac{T_i}{2}$$

Note that the total number of wasted votes in a district $W_i = W_i^A + W_i^B$ is always half of the turnout T_i. It's a question of distribution. If most of the wasted votes belong to the winning side then it's a packed district. If most of the wasted votes belong to the losing side then it's a competitive district. If there are several adjacent districts where most of the wasted votes are on the losing side then it may be a cracked plan.

Finally, the definition

The *efficiency gap* associated with the districting plan $\mathcal{D}$ is

$$EG = \sum_{i=1}^{S} \frac{W_i^A - W_i^B}{T} = \frac{W^A - W^B}{T}.$$

Finally, the definition

The *efficiency gap* associated with the districting plan $\mathcal{D}$ is

$$EG = \sum_{i=1}^{S} \frac{W_i^A - W_i^B}{T} = \frac{W^A - W^B}{T}.$$

How to interpret (in theory):

- If EG is large and positive, the districting plan is unfair to A.
- If $EG \approx 0$, then the plan is fair in the sense that both parties waste about an equal number of votes.

Back to the example

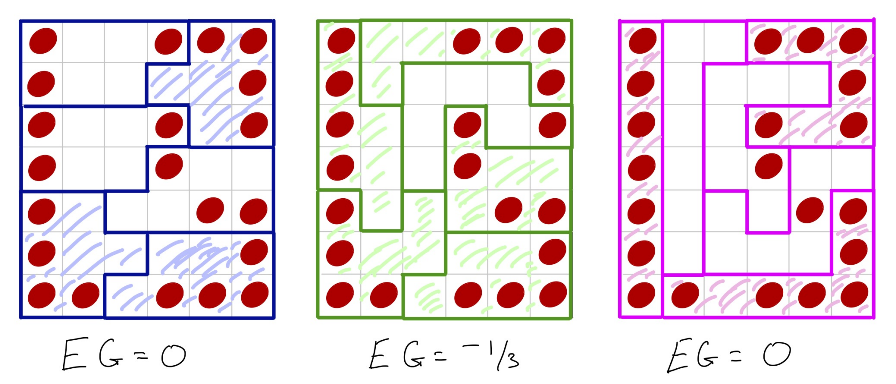

Fair

Unfair to *B*

Fair

How much is too much?

What's the "acceptable" threshold for EG? The designers of the measure, Stephanopoulos and McGhee argue that the right number for gerrymandering detection is $EG > 0.08$. This was used most famously last year in Whitford v. Gill, a case about partisan gerrymandering in Wisconsin that reached the U.S. Supreme Court.

How much is too much?

What's the "acceptable" threshold for EG? The designers of the measure, Stephanopoulos and McGhee argue that the right number for gerrymandering detection is $EG > 0.08$. This was used most famously last year in Whitford v. Gill, a case about partisan gerrymandering in Wisconsin that reached the U.S. Supreme Court.

The last three State Assembly elections had $0.1 < EG < 0.13$ in favor of the Republicans, who had drawn the maps.

How much is too much?

What's the "acceptable" threshold for EG? The designers of the measure, Stephanopoulos and McGhee argue that the right number for gerrymandering detection is $EG > 0.08$. This was used most famously last year in Whitford v. Gill, a case about partisan gerrymandering in Wisconsin that reached the U.S. Supreme Court.

The last three State Assembly elections had $0.1 < EG < 0.13$ in favor of the Republicans, who had drawn the maps.

Chief Justice John Roberts famously referred to this measure as "sociological gobbledygook" during oral arguments (this did not go over well with members of the academic community).

But what does it really measure?

We're all mathematicians here. Let's make a simplifying assumption: equal voter turnout:

$$T_i = \frac{T}{S}.$$

But what does it really measure?

We're all mathematicians here. Let's make a simplifying assumption: equal voter turnout:

$$T_i = \frac{T}{S}.$$

In this case the efficiency gap reduces, via some simple algebra, to

$$EG = \tau - \frac{1}{2}\sigma.$$

That is, the efficiency gap is just the statewide vote lean minus half the statewide seat lean favoring A.

UF

But what does it really measure?

We're all mathematicians here. Let's make a simplifying assumption: equal voter turnout:

$$T_i = \frac{T}{S}.$$

In this case the efficiency gap reduces, via some simple algebra, to

$$EG = \tau - \frac{1}{2}\sigma.$$

That is, the efficiency gap is just the statewide vote lean minus half the statewide seat lean favoring A.

Aside: while this assumption is unrealistic, dropping it only really makes the problems with EG worse.

The problems

- **Penalizes proportionality.** If party A has 60% of the statewide vote and 60% of the seats, EG calls this an unacceptable gerrymander in favor of party B: $\tau - \sigma/2 = 0.2 - 0.1 > 0.08$.
- **Incentivizes** $3:1$ **landslide districts.** Each district has 50% vote wastage; the only way to share that evenly in a single district is to have a 75-25 vote split.
- **Edge case breakdown.** 75% is an artificial sweet spot, but 80% breaks it completely: if A controls more than 79% of the vote then $\tau > 0.79 - 0.21 = 0.58$ so that getting $EG < 0.08$ would require $\sigma > 1$!

The upshot

Efficiency gap is a convenient, easy-to-calculate quantity which *might* suggest that gerrymandering has taken place. But it can be gamed in very undemocratic ways.

Probability

A simple question

Suppose a particular state has 13 Congressional districts and in a certain election the two parties have the following vote totals (by percentage), and number of seats won:

D	50.60%	4
R	48.75%	9

A simple question

Suppose a particular state has 13 Congressional districts and in a certain election the two parties have the following vote totals (by percentage), and number of seats won:

D	50.60%	4
R	48.75%	9

Suppose I also told you that the Republicans had drawn the district boundaries after the 2010 Census.

A simple question

Suppose a particular state has 13 Congressional districts and in a certain election the two parties have the following vote totals (by percentage), and number of seats won:

D	50.60%	4
R	48.75%	9

Suppose I also told you that the Republicans had drawn the district boundaries after the 2010 Census.

Would you be suspicious? How would you prove your suspicions?

Not a made-up example

This is exactly what happened in North Carolina in 2012.

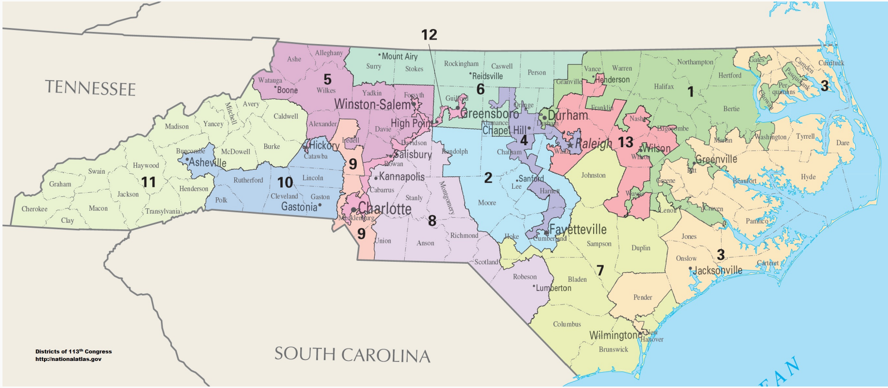

Before this election, the split had been 7-6 Democrat–Republican, which most people would agree was fair, given the vote proportions. Had the 2012 election yielded a 7-6 Republican advantage, despite the vote totals, few eyebrows would have raised.

Before this election, the split had been 7-6 Democrat–Republican, which most people would agree was fair, given the vote proportions. Had the 2012 election yielded a 7-6 Republican advantage, despite the vote totals, few eyebrows would have raised.

But 9 to 4? That just doesn't seem right. So here's a question:

What is the probability of this outcome?

Before this election, the split had been 7-6 Democrat–Republican, which most people would agree was fair, given the vote proportions. Had the 2012 election yielded a 7-6 Republican advantage, despite the vote totals, few eyebrows would have raised.

But 9 to 4? That just doesn't seem right. So here's a question:

What is the probability of this outcome?

Late edit: It happened again in 2018. Democrats received 200,000 more votes statewide and the Republicans won **10** of the 13 seats.

Building a model

States are divided into Voting Tabulation Districts (VTD). We can therefore represent a state as a graph G with vertices V representing the VTDs and edges E joining adjacent VTDs. For the 2012 North Carolina election, this graph has over 2500 vertices and over 8000 edges.

A *redistricting plan* is a function $\xi : V \to \{1, 2, \ldots, 13\}$ assigning each VTD to a Congressional district.

Pop Quiz. How many redistricting plans are there?

Answer

There are more than

$$13^{2500} \approx 7.2 \times 10^{2784}$$

possible plans!

Answer

There are more than

$$13^{2500} \approx 7.2 \times 10^{2784}$$

possible plans!

By contrast, current estimates of the number of atoms in the universe range from 10^{78} to 10^{82}.

But how many are there really?

Not all of the possible plans work. Recall that there are minimal requirements on a redistricting plan:

- Equal size
- Contiguity
- Compactness

So this cuts down the number, but it's still an enormous set.

To compute the exact probability of a 9 to 4 Republican victory, we would need to examine every possible legal redistricting plan and count how many yield this outcome.

To compute the exact probability of a 9 to 4 Republican victory, we would need to examine every possible legal redistricting plan and count how many yield this outcome.

This is unfeasible.

Let's go to Monte Carlo

Luckily, probabilists have figured out a way to estimate the answer when the sample space is unreasonably large. It's called the *Monte Carlo method* (fun fact: this process was invented by a topologist, Stanislaw Ulam, by playing solitaire while he recovered from surgery).

Let's go to Monte Carlo

Luckily, probabilists have figured out a way to estimate the answer when the sample space is unreasonably large. It's called the *Monte Carlo method* (fun fact: this process was invented by a topologist, Stanislaw Ulam, by playing solitaire while he recovered from surgery).

The basic idea is to take a reasonably large sample from the space of outcomes and count the number of times each result comes up. Of course there are lots of technicalities: How do you sample? What is "reasonably large"? How do you ensure uniform sampling?

So was the map gerrymandered?

All these technical questions can be addressed. Using the actual voting data from the 2012 election, simulations yield the following histogram

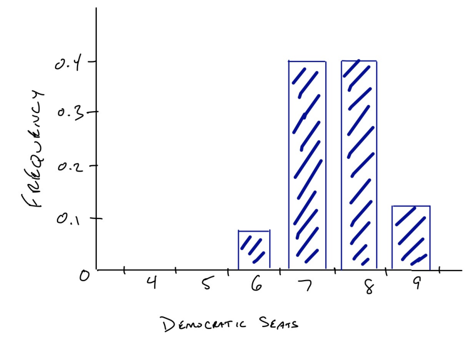

adapted from [Mattingly and Vaughn, 2014]

This is 100 samples drawn from a probability distribution on the space of allowable redistrictings. Not once did the Democrats win fewer than 6 seats, and more than 80% of the simulated elections yielded 7 or 8 Democratic seats.

The upshot

It is *extremely* unlikely that a random redistricting plan would have yielded the 9 to 4 outcome.

The upshot

It is *extremely* unlikely that a random redistricting plan would have yielded the 9 to 4 outcome.

This is a good method to detect gerrymandering *after the fact*, but it can also be used proactively to examine proposed plans before they are put into place using the most recent election (or two).

Geometry

If it looks like a duck...

All this algebra and probability is great, but can't we just *look* at a map and tell that it's a partisan gerrymander? I mean, if a district looks like this, isn't it obvious?

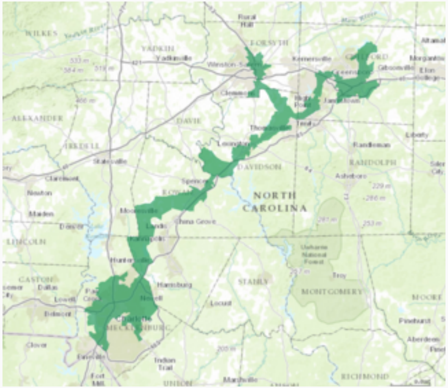

The Isoperimetric Inequality

One reason we are suspicious of districts like NC-12 is that they don't appear very "compact." That is, they tend to stretch out over unnecessary distances.

There's actually a classical way to measure this:

The Isoperimetric Inequality

One reason we are suspicious of districts like NC-12 is that they don't appear very "compact." That is, they tend to stretch out over unnecessary distances.

There's actually a classical way to measure this:

Theorem. Suppose X is a closed curve in the plane of length L bounding a region of area A. Then

$$\frac{4\pi A}{L^2} \leq 1,$$

with equality if and only if X is a circle.

The Isoperimetric Inequality

One reason we are suspicious of districts like NC-12 is that they don't appear very "compact." That is, they tend to stretch out over unnecessary distances.

There's actually a classical way to measure this:

Theorem. Suppose X is a closed curve in the plane of length L bounding a region of area A. Then

$$\frac{4\pi A}{L^2} \leq 1,$$

with equality if and only if X is a circle.

This gives us a way to quantify compactness. This ratio lies between 0 and 1; the closer to 0 it is the less compact the region is.

This forms the basis of the Polsby–Popper Test. Given a redistricting plan compute this ratio for all the districts and see if they are mostly compact. For example, for NC-12:

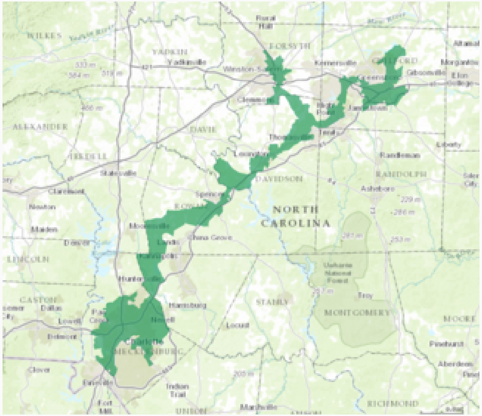

Here we compute the isoperimetric ratio to be

$$\frac{4\pi A}{L^2} = 0.0291$$

That's pretty small. In fact, it was the smallest such ratio of any Congressional district in 2012.

Other area-based measures

A couple of other things you might try:

Other area-based measures

A couple of other things you might try:

- Compare the area of the district to the area of the smallest circle circumscribing the district.

Other area-based measures

A couple of other things you might try:

- Compare the area of the district to the area of the smallest circle circumscribing the district.
- Compare the area of the district to the area of the convex hull of the perimeter of the region.

Other area-based measures

A couple of other things you might try:

- Compare the area of the district to the area of the smallest circle circumscribing the district.
- Compare the area of the district to the area of the convex hull of the perimeter of the region.
- Compute the "bizarreness" of the district: the probability that the straight line joining any two people in the district lies entirely within the district. (This is really about convexity, but it's related to compactness.)

Meanderingness

Here's something I worked on with my student Eion Blanchard [Blanchard and Knudson, 2018]. Can we measure how much a legislative district "meanders"? How about this district?

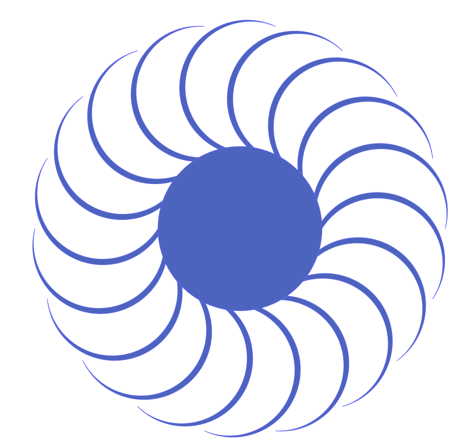

It's relatively compact. The area of the smallest circle containing it isn't much larger than the area of the region itself. And you could game that by making the spirals tighter. BUT it sure seems to wander around a lot. Can we measure that?

The Medial–Hull Ratio

Our idea was to measure the *medial axis* of a district. This is essentially the central spine of a figure. Computing it is not all that easy, but there are algorithms. One problem: some states are bigger than others and so their districts could have long medial axes just because they're large.

To fix that, we compute the medial axis of the convex hull of the district and take the ratio. This gives us a dimensionless quantity which we can use to search for evidence of gerrymandering.

Test case: Pennsylvania

Here's Pennsylvania Congressional district 11, before the maps had to be redrawn by order of the state supreme court.

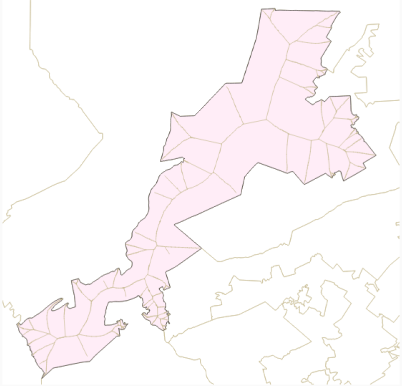

The medial–hull ratio here is 3.20. In our computational experiments, we determined that values of this ratio above 2.80 should flag a district for strong suspicion.

Here is the current North Carolina map with all the medial axes shown. The dark red districts are most suspicious under the medial–hull ratio.

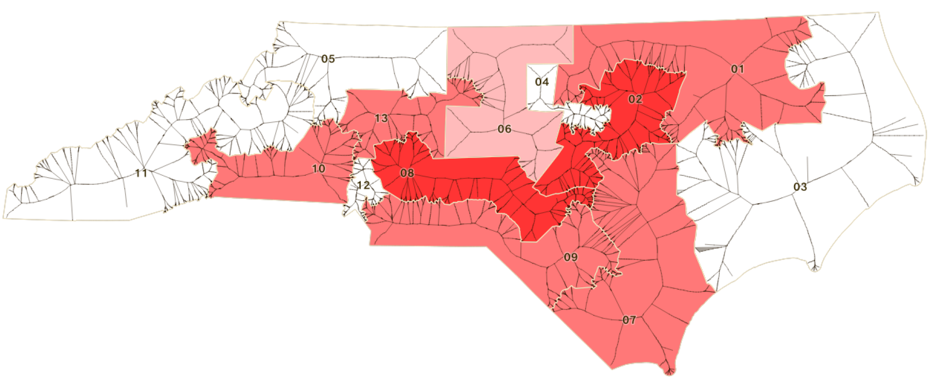

The Future

What's the solution?

First, the bad news. The solution to this problem is *political*. It requires the hard work of discussion, debate, knocking on doors, contacting legislators, recruiting new candidates, and, ultimately, voting.

The good news

Mathematics might not solve the problem completely, but it provides tools to help citizens and policymakers move toward a fairer system. In the end, that's what we all want.

References

- Bernstein, M. and Duchin, M. (2017).
  A formula goes to court: partisan gerrymandering and the efficiency gap.
  Notices AMS, 64:1020–1024.
- Blanchard, E. and Knudson, K. (2018).
  Measuring congressional district meandering.
  arXiv:1805.08208.
- Mattingly, J. and Vaughn, C. (2014).
  Redistricting and the will of the people.
  arXiv:1410.8796v1.