less than half a percent.* The number of districts located entirely within Milwaukee County dropped from eighteen in the old map to just thirteen. Democrats hold eleven of those assembly districts, and ten of those are so uncompetitive that Republicans didn't even field a candidate in 2018.

Politics, as the saying goes, ain't beanbag, and from one point of view, there's no unfairness to be found here. Legislating is a game where whoever's ahead gets to change the rules on the fly, and there's no right or wrong, just winning and losing. But most people see something to be leery of in the practice of gerrymandering, and some of those people are federal judges. Wisconsin's districts were the subject of court challenges almost from the moment Scott Walker signed them into law. Two of the districts were modified by judges in 2012 to make the map less hostile to Hispanic voters in Milwaukee, in a decision that starts out, "There was once a time when Wisconsin was famous for its courtesy and its tradition of good government," and which describes the map-drawers' claims that they worked without partisan bias as "almost laughable." Then, in 2016, a three-judge panel of the U.S. District Court for the Western District of Wisconsin threw out the whole map as a specimen of political gerrymandering in violation of the U.S. Constitution. That decision was appealed and made its way to the U.S. Supreme Court, which had long labored to find a reasonable legal standard for how much partisan gerrymandering was too much. What happened after that was a collision of math, politics, law, and motivated reasoning whose implications American politics is still absorbing.

"A MINORITY RULE IS ESTABLISHED"

If you know anything about gerrymandering, it is probably this pair of facts, which are glommed together right there in the name: first, that it was invented by Elbridge Gerry, who as governor participated in a dis-

* Added in proof: One more, the 13th, was flipped by Democrats in November 2020. The statewide presidential vote was almost exactly even, as the Walker-Evers race had been, and Repubicans maintained a 61–39 assembly majority.

tricting of Massachusetts designed to assist the Democratic-Republicans in fending off the Federalists in the 1812 election; and second, that it involves districts drawn with bizarrely sinuous boundaries, like the "salamander"-shaped district in Massachusetts that a cartoonist immortalized as the "Gerry-mander."

Both these facts are wrong. First of all, gerrymandering in America goes back well before the word and well before Gerry. According to Elmer Cummings Griffith's definitive study, his 1907 PhD dissertation in history at the University of Chicago, the practice dates back at least as far as the colonial assembly of Pennsylvania in 1709. And in early America, the most notorious example of politically motivated district making was carried out by Patrick Henry—"Give me liberty or give me death" Patrick Henry, whose pro-liberty attitude was tempered by his desire to maintain iron control over the Virginia legislature. Henry was a bitter opponent of the new U.S. Constitution, and was determined to keep one of its chief architects, James Madison, out of Congress in the 1788 election. At Henry's direction, Madison's home county was placed in a district with five counties that were seen as anticonstitutional, which Henry hoped would vote for Madison's opponent, James Monroe. Just how unfair this district was is disputed to the present day, but there's no question Madison and his allies felt Henry was playing dirty. Madison didn't get the easy path to Congress he'd hoped for, instead having to return home from New York to campaign for weeks throughout the district. He had a bad case of hemorrhoids that made travel difficult, and picked up frostbite on his face debating Monroe outdoors in January in front of a crowd of Lutherans. Gerrymander or no gerrymander, Madison prevailed, in part by winning his home base of Orange County by 216 votes to 9.

So by the time Gerry gerrymandered, it was no innovation, but an established political technology. (Stigler's Law strikes again!) By 1891, the practice, intertwined with other colors of electoral shenanigans, was so severe as to move President Benjamin Harrison* to warn in his State of the Union address:

* Who won a decisive majority of the Electoral College in 1888 while losing the popular vote, for what it's worth.

> If I were called upon to declare wherein our chief national danger lies, I should say without hesitation in the overthrow of majority control by the suppression or perversion of the popular suffrage. That there is a real danger here all must agree; but the energies of those who see it have been chiefly expended in trying to fix responsibility upon the opposite party rather than in efforts to make such practices impossible by either party. Is it not possible now to adjourn that interminable and inconclusive debate while we take by consent one step in the direction of reform by eliminating the gerrymander, which has been denounced by all parties as an influence in the selection of electors of President and members of Congress?

Harrison's description of democracy under the gerrymander has lost none of its aptness: "A minority rule is established that only a political convulsion can overthrow."

This raises a question. If legislators have been drawing district boundaries to suit their partisan interests for three hundred years, and democracy has more or less persisted, why, now, is urgent reform suddenly required?

That's partly a technology story. An old Wisconsin election hand once told me how redistricting used to be done. There was a person who, over decades of experience in Wisconsin politics, had memorized the idiosyncrasies of every voting ward in the state, from Kenosha to Superior. And the districting savant would have a big paper map spread out on a giant conference table, and would gaze at it, move a chunk here and a chunk there, denote changes with a marker, and the thing was done.

Gerrymandering used to be an art; advanced computation has made it a science. Joe "Aggressive" Handrick and his team of mapmakers tried out map after map, tweak after tweak, not on a wooden table but on a screen. They ran each potential districting through simulations that tested its performance in a wide range of political climates, until they converged on a map optimized to preserve Republican control in all but the most extreme circumstances. That process isn't just faster, it's better. A lawyer involved in the suits against the state told me that the

effectiveness of the Act 43 gerrymander went far beyond what any old-fashioned map master was able to accomplish.

What's more, a gerrymander that works well in the initial elections after the cycle creates a population of incumbents for the gerrymandering party, adding even more advantage to what the gerrymander provides. Opposition donors, assessing the map as too tilted to be overcome, allocate contributions elsewhere. And so the gerrymander feeds itself.

Justice Sandra Day O'Connor, writing in dissent in the 1986 case *Davis v. Bandemer*, argued that courts didn't need to intervene in redistricting cases. Remember, making a good gerrymandered map involves constructing a lot of districts where your party has a moderate advantage, set against a few where your opponent predominates. Doesn't that mean, O'Connor asks, that gerrymandering is an inherently risky strategy? In her account, parties will refrain from gerrymandering a state unreasonably strongly, because it puts their incumbents at too much risk of being knocked over by an unanticipated political gust. "[T]here is good reason to think," she wrote, "that political gerrymandering is a self-limiting enterprise."

Back then, she may have been right. But today's computational power has blown away the self-limiting nature of the enterprise, as it has with so many other limits. (Ask Marion Tinsley.) Just as the maps can be tuned to produce substantial partisan advantage, they can be jiggered to reduce risk to incumbents at the same time. And that's not just because a souped-up modern computer is faster than an Apple II. Voters have changed, too! We Americans like to think of ourselves as dispassionately assessing the ballot without prejudice, making a study of each candidate's policy platform and temperamental fitness for the office, and choosing the one who makes the best case for our vote. But actually we're pretty predictable, and getting more so. Most of us just vote for the person with the right letter after their name. The proportion of "floating voters" who switch parties from one presidential election to the next, which hovered around 10% from the mid-twentieth century through the 1980s, has dropped to half that. The more stable and predictable voters are in their choices, the greater the ability a party has to draw a map that preserves its majority *and* protects incumbents *and* persists long enough in its

effect to get to the next census, and thus to a brand-new map drawn by the same old legislative majority in the same old locked room.*

STOP KICKING DONALD DUCK

A traditional point of view on gerrymandering is that, as Justice Potter Stewart said in a very different judicial context, you know it when you see it. And yes, there are some weird-shaped legislative districts out there, like the Fourth Congressional District of Illinois, the "earmuffs," which consists of two distinct regions connected by a mile or two of freeway, or this beauty in Pennsylvania, colloquially known as "Goofy Kicking Donald Duck":

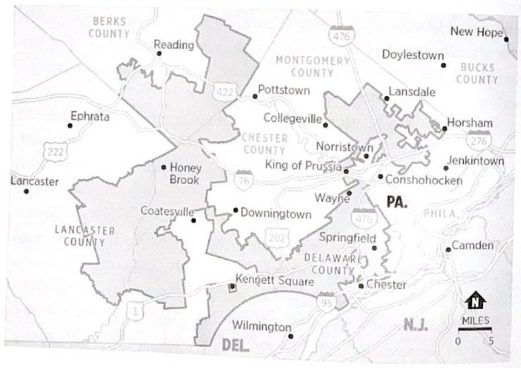

Pennsylvania's Seventh District was drawn this way in order to capture enough scattered Republicans to form a GOP-favoring district. The

* Compare what Elmer Cummings Griffith has to say about a similarly polarized era: "By 1840 two great parties had settled down for a steady and continuous struggle for political supremacy. With general political stability, elections could be predicted with greater promise of success. When parties become established, the change of voters from one to the other is a very small per cent. And the results of an election can be safely predicted within certain definite limits."

two main figures are connected only by means of the grounds of a hospital located at the tip of Goofy's business foot. Goofy's neck is a single parking lot.

This district was thrown out by the Pennsylvania Supreme Court in 2018 as an example of partisan gerrymandering gone too far: a victory for fair elections and roughly round shapes. A commonplace belief in the history of redistricting reform is that we can prevent the excesses of gerrymandering if we require districts to have "reasonable" shapes, thus limiting the ability of legislators to make mischief. Many state constitutions even have provisions directing mapmakers to avoid districts shaped like Disney melees; Wisconsin's, for instance, requires districts to be "in as compact form as practicable." But what exactly does this mean? Lawmakers have never settled on a consensus standard. And attempts to specify what shapes are "compact" sometimes make things even more muddled. In 2018, Missouri voters approved a referendum amending their constitution to require forevermore that "compact districts are those which are square, rectangular, or hexagonal in shape to the extent permitted by natural or political boundaries." First of all, a square is a kind of rectangle. And what has Missouri got against triangles, pentagons, and nonrectangular quadrilaterals? (My personal theory is that Missouri is self-conscious about being a trapezoid and is overcompensating.)

The discipline of geometry does offer some options for measuring the "compactness" of a shape. Your intuition is probably that a very complicated shape, like the erstwhile Pennsylvania Seventh, encloses its area very inefficiently, using a long, complicated, jiggery boundary. So maybe we want shapes whose perimeter isn't too large compared to their area.

Your first thought might be to use a ratio: How much area is there per mile of perimeter? A higher score would be better. Here's the problem with that idea. A tiny square district four miles north to south and four miles east to west will have a perimeter of 16 and an area of $4 \times 4 = 16$. So your area/perimeter score would be $16/16 = 1$. But what if you magnified the square district so that it was 40 miles on a side? Now the perimeter would be 160 and the area 1,600. The score has improved to 1,600/160, or 10.

This is an unpleasant state of things. How "compact" a square is

shouldn't depend on its size! Nor should it depend on whether we denominate its size in kilometers, miles, or furlongs! Whatever measure of "compactness" we use should be what geometers call an *invariant*;* it shouldn't change when the region is moved, or rotated, or enlarged, or shrunk. When we move or rotate a region, its perimeter and area don't change, so that's no problem. When we magnify it by a factor of 10, though, its perimeter grows by that same factor of 10 while its area grows by a factor of 100. This suggests that a better ratio to use is

$$\text{area}/\text{perimeter}^2$$

which doesn't change when you magnify or shrink the district. By the way, a very handy way to keep track of this sort of thing is to write everything with units of measurements attached! The perimeter of our 40-mile square is 160 miles, while its area is 1,600 *square* miles; so the area divided by the perimeter isn't 10, it's 10 *miles*, a length, not a number.

Redistricting types call the ratio above the Polsby-Popper score, after two lawyers who realized its relevance in the 1990s, but the notion is older than that. For a circle of radius r, the perimeter is $2\pi r$ and the area is πr^2, so the score is

$$(\pi r^2)/(2\pi r)^2 = \pi r^2/4\pi^2 r^2 = 0.079\ldots$$

and note that the answer doesn't depend on the radius of the circle at all! The r's cancel out. That's the invariance at work. It's the same deal for squares; if the length of a side is d, the perimeter is 4d and the area is d^2, so the Polsby-Popper score is

$$d^2/(4d)^2 = d^2/16d^2 = 1/16 = 0.0625$$

which doesn't depend on the length of the side. The square's score isn't quite as good as $1/4\pi$. In fact, it turns out that $1/4\pi$ is the very highest

* In particular, an invariant for the *similarities* discussed in chapter 3.

score any shape can possibly have! This jibes with what our intuition tells us about how large the area of a figure can be if you fix its perimeter. Put a loop of string on the table and try to "inflate" it by packing as much material inside the loop as possible; don't you feel like it would take on circular form? This fact was known and proved by Zenodorus, in the somewhat casual sense that most ancient mathematicians proved things, a century or so past Euclid. Mathematicians call it "the isoperimetric equality." It didn't get a proof up to the standards of modern geometers until the nineteenth century.

So you can think of the Polsby-Popper score as measuring how "circle-like" a district is, at which point you should start to wonder whether it's actually a good idea. Is a circular district actually better than a square? Is a longer rectangle like this one

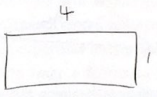

with a score of 4/100 = 0.04, really that much worse?

For that matter, what do we actually mean by perimeter? Real-life boundaries of districts are partly straight surveyor's lines but partly things like seacoasts, which are fractally wiggly at every scale, so they get longer and longer the more closely you measure every miniscule jig and jag. The quality of a district shouldn't depend on the size of your ruler!

Let's take another tack. In many ways, the most manageable geometric figures are those that are *convex*. Speaking vaguely, a convex figure is one that only bends outward,

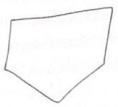

never inward:

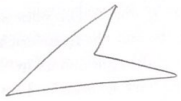

But there's a lovely official definition: a shape is convex when the line segment between any two points in the shape is entirely contained within the shape. (This definition makes sense in two dimensions, or three, or even some much larger number of dimensions well beyond your capacity to visualize "inward" or "outward" bending.) You can see how the latter shape above fails that line segment test:

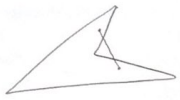

The *convex hull* of a shape is the union of every line segment joining every pair of points in the shape:

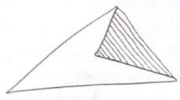

You can think of it as "filling in all the non-convex places," or, more physically, as wrapping plastic wrap around your shape as tightly as you can. The convex hull of a golf ball is a sphere; the dimples get filled in. Your own convex hull is very tightly wrapped around you if you press your legs together and your arms to your sides, but much greater in extent if you extend your limbs in all directions.

Anyway: the "Population Polygon" score of a district is the ratio between the number of people who live in the district and the number of people who live in its convex hull. The convex hull of Goofy and Donald Duck contains all those people *between* Goofy and Donald Duck, so the score of that district is going to be pretty bad.

Population Polygon is an improvement over Polsby-Popper, since it takes into account where people actually live. But there's a deeper problem with enforcing compactness as a brake on gerrymandering, which is that it doesn't work. Maybe in the days of paper maps people had to resort to wacky shapes to get just the portfolio of voters they wanted, but no longer. The map software that allows you to assess a million maps in an afternoon simultaneously lets you pick the ones that are nicely shaped *and* achieve your objectives. Those gerrymandered districts around Milwaukee are innocent-looking quasi-rectangles, which would get acceptable scores on any quantitative measures of compactness.

Sandra Day O'Connor once wrote* that when it comes to legislative districts, "appearances do matter": a salamandery district creates the *perception* that something other than democratic ideals is at work. If you ask me, those ideals aren't much shored up by replacing Goofy and Donald with a map that's just as partisan but less offensive to the eye. There are some good reasons to want a district to be compact, I guess—a shorter average drive to your state representative's office, a modestly increased alignment of political priorities among the constituents. But to the extent that compactness constrains blunt gerrymandering at all, it's just because they're constraints. The fewer choices map-drawers have, the less likely they are to be able to find options that are grievously rigged. It's not that there's anything inherently more fair about roughly round districts; it's that there are many fewer ways to break a state into districts if you have to make them roughly round.

What we now know is that traditional compactness measures just aren't enough to keep parties from stacking the deck in their favor, any more than the equal-population requirement of *Reynolds v. Sims* was. Of course, you could be stricter about the compactness measures to make the constraints more severe, or enforce state laws about breaking county boundaries, or just invent purely arbitrary rules ("the number of registered voters in each district must be a prime number") to limit the

* In the racial gerrymandering case *Shaw v. Reno*—the drawing of districts to guarantee, or prevent, minority representation is a whole other aspect of the districting story that this chapter isn't big enough to contain.

wiggle room available to legislators caught up in their decennial gerrymandering estrus. But arbitrary rules like this aren't really politically feasible. If the goal is to stop gerrymandering, the strategy is going to have to target gerrymandering directly. That means we need a measure of a map that tells us not how equally populated the districts are, or how round and plump they are, but *how gerrymandered* they are. That's a harder problem still. But geometry can get us there.

HIVE THE GRITS!

To paraphrase H. L. Mencken, almost every interesting question of applied mathematics has an answer which is simple, mathematically elegant, and incorrect. For districting, that answer is *proportional representation*: the principle that a party ought to get a share of seats in the legislature equal to the proportion of the popular vote drawn by its candidates. That's one straightforward quantitative answer to what it would mean for a district map to be "fair," and it's really popular. *The Washington Post* reported that, in the 2016 Wisconsin State Assembly elections, 52% of votes went to Republican assembly candidates, but the Republicans won 65% of the seats; the GOP, the paper wrote, "seems to have benefited from gerrymandering, given that discrepancy between votes won and seats held." The implicit suggestion is that something's smelly when those numbers don't match.

Proportional representation is the reason people tend to like Option 1 of the Crayola maps. The Purple party got 60% of the vote, and it gets 60% of the seats.

But would proportional representation actually be the outcome if maps were drawn fairly? Almost surely not! Take the Wyoming State Senate. Wyoming is by some measures the most strongly Republican state in America. Two-thirds of its voters picked Donald Trump in 2016, and the same proportion voted Republican in the governor's race in 2018. But the state senate isn't two-thirds Republican; there are twenty-seven GOP senators and only three Democrats. Should we really see

that as unfair? When a state's population is two-thirds Republican, it's pretty likely that almost every geographic chunk of the state is pretty Republican. In the extreme case of a state that's totally homogeneous politically, every neighborhood in every town having the same proportion of Democrats and Republicans, the party with the popular vote majority would win every single seat in the legislature. That's the scenario of Crayola Option 4. And the single-party legislature wouldn't be the result of gerrymandering, but of the state's strangely consistent voter distribution.

Idaho has two representatives in Congress, and so does Hawaii, and we don't think it's strange that Idaho's delegation has been fully Republican and Hawaii's fully Democratic for the past decade,* even though the proportion of voters who back the majority party in each state is closer to 50% than it is to 100%. I don't think a fair splitting of Idaho into two districts would return one Democrat and one Republican. I don't even think you *could* draw a non-ridiculous district that covered half of Idaho and had a Democratic majority.

And what about the plight of the Libertarians? The proportion of Americans voting for Libertarian candidates for the House of Representatives consistently hovers around 1%; but there has never been a U.S. representative elected from that party, let alone the three to five proportional representation would recommend, because there is no such thing as a libertarian city or even neighborhood (though it's fun to imagine!). In Canada, whose elections are structured very similarly to those in the U.S., the deviations are even starker; in the 2019 federal elections there, the New Democratic Party drew 16% of the vote against only 8% for the Bloc Québécois, but the Bloc, whose voters are concentrated in a single province, won substantially more seats in Parliament.

Canada, by the way, doesn't have a gerrymandering problem, despite having a U.S.-style legislature. That's not because Canadians are nicer than Americans. It's because Canada has assigned the drawing of districts (called *ridings* up north) to nonpartisan commissions since 1964.

* Each state had one representative from the minority party for one term after the 2008 election, to be fair.

Before that, district drawing was as politically motivated and dirty as it is in the United States. Canada's very first prime minister, the Conservative Sir John Macdonald, wielded the districting pen in a ruthless effort to diminish the power of his opponents in the Liberal Party, the so-called "Grits." The map used for the election of 1882 was so brazen as to inspire this poem published in the *Toronto Globe*, surely the clearest explanation of the principles of gerrymandering ever set to trochaic tetrameter:

> Therefore let us re-distribute
> What constituencies are doubtful
> So as to enhance our prospects;
> Hive the Grits where they already
> Are too strong to be defeated;
> Strengthen up our weaker quarters
> With detachments from these strongholds
> Truly this is true to nature
> In a mighty Tory chieftain!

Proportional representation is a perfectly reasonable system; many countries build it into their method of assembling legislatures. But it's not *our* system, and it's unreasonable to expect proportional representation to be the outcome of a U.S. election. Nonetheless, the specter of proportional representation still haunts the gerrymandering discourse. In a closed-door seminar advising Republicans on how to draw maps in their favor without running afoul of judges, secretly taped by a participant, the GOP election lawyer Hans von Spakovsky warned his audience about those who would try to overturn the maps in court:

> What they were arguing was that, if for example, the Democratic Party has a presidential candidate who gets 60% of the votes statewide in the presidential election, why then they're entitled to 60% of the state legislative seats, and 60% of the congressional seats.

This is false, though it's not clear to me whether Spakovsky knows it's false. Proportional representation isn't the standard reformers are arguing for. So what is?

MIND THE GAP

The 2004 Supreme Court case of *Vieth v. Jubelirer* put the problem of partisan gerrymandering into a curious legal limbo. Four justices felt that the practice of gerrymandering for partisan gain was entirely nonjusticiable; that is, that it was a purely political matter the federal courts were forbidden to interfere with. Four felt that the map insulted the right to representation so grievously as to constitute a constitutional violation.

Justice Anthony Kennedy, in this case as in many others the fulcrum of the court, joined the majority in upholding the gerrymandered map at issue, but disagreed with the authors of the majority opinion on the critical matter of justiciability. Courts *did* have the power and duty to stop partisan gerrymandering, he wrote, if only there were a reasonable standard judges could use to determine when a map is so bad it makes the Constitution puke.

We've seen that proportional representation isn't that standard, and neither are measures of geometric compactness. So a new idea was needed. Reformers got one from political scientist Eric McGhee and law professor Nicholas Stephanopoulos in the form of the "efficiency gap."

Remember: what makes a gerrymander work is that your party wins a lot of districts by a little and a few districts by a lot. You can think of that as an "efficient" allocation of your party's voters. Looking at Crayola Option 2 through that lens, we can see a massive failure of efficiency on Purple's part. What good does their 85,000–15,000 win in the 7th District do them? They would have been better off swapping 10,000 of those voters into District 6 and bringing 10,000 Orangists into District 7 in exchange; they'd still win the 7th by a commanding 75,000–25,000 margin, but now they'd win the 6th District 55,000–45,000 instead of losing it by the same amount.

Those extra Purples in District 7 are, from the point of view of their party, wasted. In Stephanopolous and McGhee, "wasted votes" are votes which are either

- votes cast in a district where your party loses; or
- votes above the 50% threshold in a district where your party wins.

In Option 2, the Purple party wastes a *lot* of votes. Here's a chart:

WASTED	Purple votes	Orange votes	WASTED
45,000	45,000	55,000	5,000
45,000	45,000	55,000	5,000
45,000	45,000	55,000	5,000
45,000	45,000	55,000	5,000
45,000	45,000	55,000	5,000
45,000	45,000	55,000	5,000
35,000	85,000	15,000	15,000
35,000	85,000	15,000	15,000
30,000	80,000	20,000	20,000
30,000	80,000	20,000	20,000

There are 45,000 votes wasted in each of the six losing districts; in each of the 7th and 8th, the 35,000 votes exceeding what Purple needs for a majority are wasted, too; and in the 9th and 10th Districts, 60,000 of the 160,000 Purple votes are wasted. That adds up to 6 × 45,000 + 70,000 + 60,000 or 400,000 wasted votes.

Orange, by contrast, is incredibly efficient. Only 5,000 of its votes are wasted in each of the first six districts; and in the districts where it loses, it loses *badly*, wasting only 30,000 votes in Districts 7 and 8 and 40,000 in Districts 9 and 10. That's 100,000 votes wasted in all, 300,000 fewer than Purple.

The *efficiency gap* is the difference between the number of votes wasted by the two parties* expressed as a percentage of the total number of votes cast. In the case of Option 2, that gap is 300,000 out of a million, or 30%.

* The *two* parties? But what if . . . yes, I know, I know. Quantifying gerrymandering when more than two parties are involved is a largely unexplored topic, and I encourage you to think about it!

That's a *huge* efficiency gap. In real elections, the figure is typically in the single digits. Some lawyers have suggested that any figure over 7% should be enough to induce a court to have a careful look-see.

Not all the options we laid out for Crayola are that gappy. Here's a chart for Option 1, the map that satisfies proportional representation:

WASTED	Purple votes	Orange votes	WASTED
25,000	75,000	25,000	25,000
25,000	75,000	25,000	25,000
25,000	75,000	25,000	25,000
25,000	75,000	25,000	25,000
25,000	75,000	25,000	25,000
25,000	75,000	25,000	25,000
35,000	35,000	65,000	15,000
35,000	35,000	65,000	15,000
40,000	40,000	60,000	10,000
40,000	40,000	60,000	10,000

Purple wastes 25,000 votes in each of the first six districts, 35,000 each in Districts 7 and 8, and 40,000 in Districts 9 and 10, for a total of 300,000. Orange also wastes 150,000 in the first six districts, but only 15,000 each in Districts 7 and 8 and 10,000 each in Districts 9 and 10, for a total of 200,000. So the efficiency gap drops to 100,000 out of a million, or 10%, still in favor of Orange. In Option 4, the map where Purple holds all the seats, the Orange Party wastes 40,000 votes in every single district, while Purple wastes only 10,000; so we get another huge efficiency gap of 30%, but this time favoring Purple. What about Option 3?

WASTED	Purple Votes	Orange Votes	WASTED
30,000	80,000	20,000	20,000
20,000	70,000	30,000	30,000
20,000	70,000	30,000	30,000
20,000	70,000	30,000	30,000
15,000	65,000	35,000	35,000
15,000	65,000	35,000	35,000
5,000	55,000	45,000	45,000
45,000	45,000	55,000	5,000
40,000	40,000	60,000	10,000
40,000	40,000	60,000	10,000

Now, each party wastes the same number of votes: 250,000. This map has efficiency gap *zero*; from the point of view of this measure, it's as fair as can be, despite its departure from proportional representation.

As far as I'm concerned, that's good! In practice, maps that are drawn by neutral arbiters rarely approach proportional representation, except for those cases where both seat share and popular vote share are close to 50–50. Instead, the seat share is usually farther from 50–50 than the vote share is. By the efficiency gap standard, an election where one party got 60% of the votes and seated 60% of the legislature might be evidence *for* gerrymandering, not against it.

The efficiency gap is an objective measure, it's easy to calculate, and reams of empirical evidence show that it jumps in maps we know were gerrymandered, like Wisconsin's. So it became a fast favorite of plaintiffs. It played a major role in the court case that threw out Wisconsin's maps in 2016, after years of legal dispute.

And here is where I pull the football away once again. The efficiency gap started to come under fire almost as soon as it got popular. It has flaws, severe ones. For one thing, it's badly discontinuous. Whether votes are wasted depends on who wins the district, which means that the efficiency gap can change wildly under very small changes in election outcome. If Purple wins a district by 50,100–49,900, Orange has wasted 49,900 votes there and Purple only a hundred. A tiny shift in

the vote gives Orange a 50,100–49,900 victory, and flips the wasted votes; now it's Purple that racks up nearly 50,000 more. That changes the efficiency gap by almost 10% all by itself! A good measure shouldn't be this brittle.

Another problem with the efficiency gap has more to do with law than math. To get a court to throw out a map, or even to have the case heard, the person bringing the case has to have standing; that is, the plaintiff has to show that they, personally and individually, were denied some portion of their constitutional rights by the state's map. When districts are wildly different in size, it's plain who was harmed: the person in the gigantic district whose vote counts for less. The standing claims in a gerrymandering case are much murkier. And the efficiency gap doesn't help much. Whose rights were denied, or at least meaningfully shaved down? It can't be everyone whose vote counts as "wasted"—that, for instance, would include everyone on the losing side of a close district election, and the voters in the most competitive districts certainly don't seem like the ones whose right to vote is being abridged. This issue of standing was precisely where the Wisconsin case foundered at the U.S. Supreme Court, which found unanimously that the plaintiffs had not done enough to establish that they, personally, were harmed by the gerrymander. The case was returned to Wisconsin for repair, but never made it back to the Supreme Court, which decided to use cases from North Carolina and Maryland to render its judgment on gerrymandering.

The efficiency gap also suffers from a certain over-rigidity. If the number of votes in every district is the same, as in our Crayola examples, then it turns out that the efficiency gap is simply the difference between

the winning party's victory margin in the popular vote

and

half the winning party's victory margin in seats.

So you get efficiency gap zero when the victory margin in seats is exactly twice the victory margin in the popular vote, and the closer you come to that standard, the smaller the efficiency gap is. In Crayola, the Purple party won the popular vote by a 20-point margin. So, as far as efficiency gap is concerned, the "right" margin of victory in legislative seats is twice that, or 40 points. That's exactly what happens in the efficiency-gap-zero Option 3, where Purple wins 70% of the seats. In Option 1, where Purple wins both the popular vote and the race for seats by a 20-point margin, the efficiency gap is 20% – 10% = 10%.

Courts don't like systems where there's is a single "correct" number of seats assigned to a given vote share. They smack of proportional representation, even when, as here, the formula is usually incompatible with proportional representation.

I say "usually" incompatible because there's one situation in which efficiency gap and proportional representation (and probably you, too) agree about what's fair. That's the scenario where each party gets exactly half the overall votes. Then there's a basic symmetry you might expect any map called "fair" to satisfy. If the population of the state is exactly evenly split, shouldn't the two parties share the legislature equally as well?

The Republican Party of Wisconsin would say no. And however I may feel about their districting skullduggery in the spring of 2011, I have to admit they have a point.

Crayola map 2 awards a majority of seats to the Orange Party, even as they get thumped by the Purples in the popular vote. But what if the Purples of the state congregate in a couple of dark-Purple metro areas, set against the background of a countryside that leans orange? It's possible you'd see results a lot like this, without any scale-thumbing by the map-drawers. Is this kind of asymmetry actually unfair if the Purple people are gerrymandering themselves?

Wisconsin's Republican attorney general Brad Schimel argued in an amicus brief to the Supreme Court that this scenario is exactly what's happening in Wisconsin. The assembly district in Madison where I live, AD77, delivered 28,660 votes to the Democrat Tony Evers. The GOP's man Scott Walker got only 3,935. In Milwaukee's District 10, Evers was

even more dominant, winning 20,621 to 2,428. There were no districts where Republicans won by even close to that much. That's not because gerrymandering packed those districts full of Democrats. Madison *just is* full of Democrats.

The superficially fair criterion that a 50–50 split of votes should result in a roughly 50–50 split of seats, Schimel argued, would actually be biased *against* Republicans, not just in Wisconsin but in every state whose dense cities were dominated by Democratic voters—which is to say, just about every state.

BIPARTISAN STATISTICAL MALFEASANCE

Not all the points in that legal brief are good ones. The Act 43 map was designed to be voter-proof by building in resistance to a uniform change in voter mood; if every district moves toward Democrats by the same fixed percentage, it takes quite a big shift to dissolve the GOP's engineered advantage. Schimel, whose task was to deny the effectiveness of the map his own party had worked so hard to create, points out that all ninety-nine districts do not, in fact, swing exactly in tandem.

There are many statewide statistical measures one could compute in order to test just how uniform those swings tend to be, and just how well the Act 43 map resists Democratic gains under a more realistic model of year-to-year variation. That would be an interesting and useful analysis. It's not what Schimel did.

Rather, he picked out one district, State Senate District 10, which had given the Republican candidate 63% in one election and 44% in the next, a 19-point swing. Is it really plausible that the whole state got that much more Democratic in such a short time? If it had, Schimel estimates, the Democrats were en route to "*winning 77 out of 99 districts*" (breathless italics Schimel's). I guess the gerrymander wasn't so bad after all!

What Schimel doesn't say is that the District 10 race won by Democrat Patty Schachtner (a bear-hunting grandmother of nine whose highest previous office was county medical examiner) was a special

election for an open seat, held in January, with turnout about a quarter of what you get in a normal election year. And in the election before that, the Republican candidate who got 63% was a popular incumbent who'd held the seat for sixteen years. You can't make the case that Wisconsin politics swings by 19 points in eighteen months unless you choose your data points very carefully to get you to that conclusion, which is exactly what Schimel did.*

Statistical malfeasance of this kind isn't limited to Wisconsin Republicans. The total number of votes cast for Democratic Wisconsin assembly candidates in 2018 was 1,306,878. Republican assembly candidates got just 1,103,505. So Democratic candidates got 53% of the popular vote for assembly—but only won thirty-six of the ninety-nine seats. That's not just a departure from proportional representation, which would be forgivable; it's an almost veto-proof majority seized by a party that got a minority of the votes. This statistic was shared everywhere. It appeared on Rachel Maddow's popular liberal TV show and was tweeted out by the head of the state Democratic Party, as proof positive of the rigging of Wisconsin's district map.

But *I* haven't mentioned it. Here's why. One of the gerrymander's chief effects is to pack Democrats into districts so homogeneous that Republicans don't have a speck of a chance there. In a year with a Democratic mood, like 2018, it's not even worth a Republican candidate's time to run. So there were thirty out of ninety-nine districts with no 2018 Republican candidate at all—my district in Madison was one of them, natch—against only eight districts lacking a Democrat. Each of those thirty uncontested races *would* have given some votes to a Republican candidate, if any had been willing to run. But that 53% number treats them as if they had no Republican sentiment at all.

Both Schimel's number and Maddow's number were correct. That somehow makes it worse! A false figure can be corrected. A true one chosen to make the wrong impression is a much harder poodle to muzzle. People often complain that no one likes facts and numbers and reason and science anymore, but as someone who talks about those things

* Added in proof: and indeed, in November 2020, her first try in a regularly scheduled election, Schachtner lost her bid for reelection by a 19-point margin.

in public, I can tell you that's not true. People *love* numbers, and are impressed by them, sometimes more than they should be. An argument dressed up in math carries with it a certain authority. If you're the one who outfitted it that way, you have a special responsibility to get it right.

WRONG QUESTIONS

If even the basic principle that even-steven votes should lead to even-steven legislatures is suspect, what hope is left of defining fairness? How are we to judge which of the four maps of Crayola is the right one? Is it Option 1, the one that satisfies proportional representation, where Purple gets six seats and Orange four? Is it Option 3, the one with zero efficiency gap, where Purple has a 7–3 edge? What about Option 4? It feels wrong for Purple to hold all the seats, but as we've seen, that's just what happens if the state of Crayola happens to be politically homogeneous, with the same 60–40 ratio of Purple to Orange in all four compass directions, in the cities and the hamlets. In that case, no matter how you draw the lines, each district will tilt 60–40 in Purple's favor and you get a monochrome legislature.

The Wisconsin GOP would suggest even Option 2 shouldn't be ruled out; if the concentration of Purples in Purpleopolis is strong enough, it may well be that any reasonable map will produce four highly Purple districts and six modestly Orange ones.

We seem to have hit an impasse; there is no clear way to look at those numbers and agree on which map is fair. This feeling of futility is welcomed by gerrymanderers, who prefer to do their dark work unconstrained. It is at the center of every argument put to the court in defense of the practice: maybe it's fair, maybe it's not, but tragically, Your Honor, there is just no way to judge.

Maybe not. But you and I are not judges. We are, for the moment, mathematicians. We're not bound by the limits of the law; we can use every tool at hand to try to figure out what's really going on. And if we're very lucky, we'll come up with something that will stand up in court.

The legal battles over gerrymandering came to a climax in March 2019, when the U.S. Supreme Court heard oral argument in two cases with the potential to finally open or close the constitutional door Justice Kennedy had left tantalizingly ajar. Kennedy himself wasn't there to hear the case; he had retired the year before, replaced on the Court by Neil Gorsuch. One case, *Rucho v. Common Cause*, came from North Carolina; the other, *Lamone v. Benisek*, was from Maryland. Both the disputed maps concerned districts for the U.S. House of Representatives; the North Carolina map, gerrymandered by Republicans, had arranged for ten of the state's thirteen seats to be reliably Republican, while in Maryland, the state's all-Democratic government had cut the number of plausible Republican seats in the state to just one out of eight. The Maryland mapmakers were advised by veteran Democratic congressman and House Majority Leader Steny Hoyer, who once told an interviewer, "Now let me make it clear, I am a serial gerrymanderer." Ironically, Hoyer's career in politics started when, as a twenty-seven-year-old political rookie, he won the 1966 race for Maryland State Senate seat 4C, a seat that had come into existence just that year after the Supreme Court threw out Maryland's unequally sized senate districts in the wake of *Reynolds v. Sims*. (Isaac Lobe Straus, sadly, didn't live to see it.)

The twin set of cases presented a perfect opportunity for the court to rule on gerrymandering without appearing to take partisan sides. The most high-profile gerrymanders in the country, in places like North Carolina, Virginia, and Wisconsin, had been drawn by Republicans, so the fight to reform redistricting was usually seen as a Democratic struggle; but high-profile Republican officials like Ohio governor John Kasich and Arizona senator John McCain also weighed in against gerrymandering, submitting amicus briefs to the court presenting their own woeful experiences with the effect of motivated mapmaking on democracy. Experts from around the country submitted their own briefs. There was a historians' brief, which quoted no fewer than eleven different *Federalist Papers*; there was a brief from a team of civil rights organizations addressing the impact on minority rights; there was a political scientists'

brief disputing Justice O'Connor's view that the problem of gerrymandering would take care of itself; and, for the first time in the history of the Supreme Court, there was a mathematicians' brief.* I signed it. In a few pages we'll get to what's in it.

Mathematicians are like Ents, the sentient trees in *The Lord of the Rings*—we don't like to get involved in the mundane conflicts of state, which are out of sync with our slow time scale. But sometimes (and I'm still inside this Ent simile, by the way), events in the world so offend our particular interests that we have to lumber in. Our intervention was necessary here because of some fundamental misapprehensions about the nature of the problem, which we hoped by means of our brief to correct. From the very beginning of the oral arguments, it was clear we hadn't fully succeeded. Justice Gorsuch, questioning the North Carolina plaintiff's lawyer Emmet Bondurant, cut to what he believed to be the chase: "How much deviation from proportional representation is enough to dictate an outcome?"

In math, wrong answers are bad, but wrong questions are worse. And this is the wrong question. Proportional representation, as we've seen, isn't what usually happens when districts are neutrally drawn. Yes, more than three-quarters of North Carolina's districts were firmly in Republican hands, even though Republican voters in North Carolina don't make up anything like three-quarters of the electorate. But that wasn't the actual problem the plaintiffs were asking the court to remedy.

It's easy to see why the judges would *want* that to be the request, because that would make their job easy; they could just say no. The case of *Davis v. Bandemer* had already established that a lack of proportional representation didn't make a map unconstitutional. But the actual issue in *Rucho* was subtler. To explain it, as is so often the case in math when we get really stuck, we need to go back to the beginning of the problem and start over.

* Technically, the "Amicus brief of Mathematicians, Law Professors, and Students," but most of us were mathematicians.