DRUNK DISTRICTING

We have tried to find a numerical standard for "fairness" and we have failed. The reason is that we've made a basic philosophical mistake. The opposite of gerrymandering isn't proportional representation, or efficiency gap zero, or adherence to any particular numerical formula. The opposite of gerrymandering is *not gerrymandering*. When we ask whether a district map is fair, what we really want to ask is:

> Does this districting tend to produce maps similar to the ones a neutral party would have drawn?

Already we are in a realm that makes lawyers tug nervously at their chins, because we're asking questions about a *counterfactual*: what would have happened in a different, fairer world? To be honest, it doesn't sound much like math, either. The question requires knowledge of the desires of the map-drawer. What does math know about desires?

A path out of this thicket was first chopped by the political scientists Jowei Chen and Jonathan Rodden. They were troubled by the problems with traditional measures of gerrymandering, especially the principle that 50% of the votes should yield 50% of the seats in the legislature. It was clear to them that concentration of one party in city districts was likely to produce what they called *unintentional gerrymandering* favoring the more rural party, even in maps drawn by disinterested actors. That's what we saw in Crayola; the party whose voters are crammed into a few districts is at an asymmetrical disadvantage when it comes to winning seats. But would that asymmetry be big enough to explain the disparities we observe? In order to find out, you need to get some neutral parties to draw maps for you. And if you don't know any neutral parties, you can just program a computer to act like one. The idea of Chen and Rodden, which is now central to the way we think about gerrymandering, is to generate maps *automatically*, and in large number, by a mechanical process that has no preference between the

parties because we didn't code it to have one. So we can rephrase our initial question:

> Does this districting tend to produce outcomes similar to those a computer would have drawn?

But of course there are lots of different ways a computer might have drawn a map; so why not leverage the computer's power to let us look at *all* the possibilities? That lets us rephrase the question in a way that starts to sound more like math:

> Does this districting tend to produce outcomes similar to a map *randomly selected* from the set of all legally permissible maps?

This suits our intuition, at least at first; one might imagine that a mapmaker truly indifferent to how many seats each party gets would be equally happy with any of the ways of scissoring up Wisconsin. If there were a million ways to do it, you could roll a million-sided die, read the tiny number off the top, pick your map, and relax until the next census.

But that's not quite right. Some maps are better than others. Some are downright illegal—if the districts are noncontiguous,* for instance, or if they violate the Voting Rights Act requirement of districts where racial minorities are likely to be able to elect representatives, or if the populations of the districts differ from each other by more than the rules allow.

And even among the maps that don't run afoul of statute, we have preferences. States want to reflect natural political divisions, avoid cutting up counties, cities, and neighborhoods. We want our districts to be reasonably compact, and in the same vein we want their boundaries not to be too snaky. You can imagine giving each district map a score that measures just how well it does with respect to these measures, which in legal terminology are called *traditional districting criteria*, but which I will call *handsomeness*. And now you choose a district at random from

* Except in Nevada, the one state with no contiguity requirement—hold that thought, we'll need it later!

among the lawful options, but in a way that's biased to favor the handsomest maps.

So let's try one more time:

> Does this districting tend to produce outcomes similar to a map randomly selected, with a bias toward handsomeness but no bias regarding partisan outcome, from the set of all legally permissible maps?

A question now presents itself. Why don't we just let our computer search and search until it finds the very handsomest of all maps, the one that best respects county boundaries and does the least non-convex squirming along its trim perimeter?

There are two reasons. One is political. People who actually work with state governments, in my experience, are unanimous in the opinion that elected officials and their voters *hate* the idea of a computer-drawn map. Districting is a task given to the people of the state, through some official body that's supposed to represent our interests. Delegating that task to an unauditable algorithm is not going to be acceptable.

If you don't like that reason, here's another one: it would be absolutely and definitively impossible to do that. A computer can pick the best map out of a hundred. It can pick the best map out of a million. The number of possible districtings is . . . a lot more than that. Remember 52 factorial, the astronomical number of orders you can put a deck of cards in? That number is like a tiny shriveled bean next to the colossal quantity of ways you can divide the state of Wisconsin into ninety-nine contiguous regions of roughly equal population.* Which means you can't simply ask your computer to assess each map's handsomeness and pick out the best one.

Instead, we may look at just a few possible maps, where by "just a few" I mean 19,184. You get a picture like this:

* No exact formula for this number, or even a decent approximation, is known. The number of ways to divide the 81 little boxes in a 9 x 9 square into nine equally sized regions, each one connected—the number of possible "Jigsaw Sudoku" pictures, if you're into that—is already 706,152,947,468,301. Wisconsin has 6,672 wards that you have to divide into 99 regions.

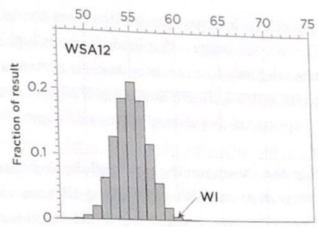

What you're looking at is what's called an *ensemble,* a set of maps generated at random by a computer. This particular computer was managed by Gregory Herschlag, Robert Ravier, and Jonathan Mattingly at Duke University. For each of those nineteen-thousand-some randomly generated maps, they take the Democratic and Republican votes cast in the actual 2012 Wisconsin State Assembly election and assign them to their new autogenerated district.* For each map, you count the number of districts in which Republicans got more votes. That's what you see in the bar chart above. The most common outcome, which happens in more than a fifth of the machine-generated maps, is that Republicans win fifty-five of the seats. Slightly less frequently, Republicans win either fifty-four or fifty-six. Together, these three possibilities cover more than half of the simulations. As you get further in either direction from that most frequent outcome† of fifty-five seats, the bar chart tails off; like so many random processes, it forms something vaguely bell-curvish, and outcomes very far from fifty-five are very unlikely. They are, to use a statistical term of art, *outliers.*

A districting that separates the 2012 voters into sixty districts with more Republicans and only thirty-nine with more Democrats is one of those outliers. That a map would yield such a good outcome for the

* There are wrinkles here, like: What do you do with people who live in wards where the real-life election was uncontested? You have to make your best guess as to how those voters *would* have voted, given a candidate in each party; you can do this by extrapolating from how the ward voted in the contested races for president, senator, and House representative going on at the same time.

† The value of a variable that occurs most frequently—the place where a bar graph peaks—is usually called the *mode,* which is yet another word invented by Karl Pearson.

GOP is highly unlikely, happening in less than one in two hundred of the computer's trials. Or rather—that kind of map is highly unlikely *if the map is chosen randomly by a person or machine without a partisan interest*. If the map is, instead, chosen by a group of consultants in a locked room with the explicit mission of maximizing GOP seats, it's the opposite of unlikely.

The ensemble also shows you the truth and the lie in the Wisconsin legislature's defense of its map. We can't help it, they say, if Democrats choose to congregate in cities among their own effete kale-eating liberal kind; that makes the legislature skew Republican even when the popular vote is split.

And that's true! But with the ensemble, we can estimate *how* true. In 2012, a typical neutrally drawn map, under conditions of near parity between Democratic and Republican assembly votes statewide, would have given the GOP a 55–44 majority. That's a lot less than the 60–39 majority they actually acquired. Six years later, in the 2018 election, Scott Walker won just under half the popular vote; nonetheless, a typical neutral map would have him coming out ahead in fifty-seven assembly districts. But the GOP-drawn map manages to create sixty-three Walker-favoring districts! The political geography of Wisconsin helps Republicans; the turbo-boost they get from the gerrymander goes above and beyond that.

At least, sometimes it does. In 2014, a midterm election year where the whole country was in a somewhat Republican mood, the GOP did well in Wisconsin, getting almost 52% of the statewide assembly vote. But they increased their assembly majority by only three, winning sixty-three out of ninety-nine seats. Filter that same election through the 19,184 random maps and it doesn't look like an outlier at all; in the 2014 election, it turns out, sixty-three Republican seats is just about what a random neutral map would likely have delivered.

What happened? Did the gerrymander lose its mojo in just two years? That would be evidence that gerrymandering doesn't need judicial intervention, but wears off on its own like a hangover. But it's not quite like that. It's more like Volkswagen. A few years ago it was revealed that the auto company had been systematically evading pollution

tests, installing software in its diesel cars to fool regulators into thinking the engines were meeting emissions standards. Here's how it worked: the software detected when the car was being tested, and *only then* did it turn on the antipollution system. The rest of the time, the car just sailed down the highway spewing particulate matter.

The Wisconsin map is a similarly audacious piece of engineering. And the ensemble method reveals it; because it gives you information not just about what happened in the state's elections, but what *might* have happened if the elections had gone a bit differently. What if we take the assembly election of 2012 and shift each of the 6,672 wards 1% in the direction of the Democrats or the Republicans? Does the gerrymander bend or does it break? This is the same flight into the counterfactual Keith Gaddie used when Republicans were designing the map in the first place. And it reveals something startling. In electoral environments where Republicans get a majority of the statewide vote, the gerrymander doesn't have much effect; those are elections where the GOP would get an assembly majority anyway. It's only in Democratic-leaning environments that the gerrymander really kicks in, acting as a firewall to maintain the Republican majority against prevailing popular sentiment. You can see the firewall in the plot on page 346: in years when the Republicans did well, the circles and the stars aren't far apart, but as the Republican popular vote share gets lower, the stars separate from the circles, staying stubbornly above the fifty-seat line that gives the GOP the majority.

The Duke team estimates through its ensembles that the Act 43 map does exactly what Gaddie predicted it would do. The map keeps the assembly in Republican hands unless Democrats win the statewide vote by 8 to 12 points, a margin rarely achieved in this evenly split state. As a mathematician, I'm impressed. As a Wisconsin voter, I feel a little ill.

I've left something out. There are six hundred kajillion possible maps; that's why we can't just pick the very best one. So how is it that we can pick nineteen thousand of them at random?

For that, we need a geometer. Moon Duchin is a geometric group theorist and a math professor at Tufts University in Massachusetts. Her Chicago PhD thesis was about a random walk on Teichmüller space.

Don't worry about what Teichmüller space is,* just focus on the random walk; that's the key. We saw it with Go positions and we saw it with card shuffling and in a minor way we even saw it with mosquitoes: a random walk, our old friend the Markov chain, is the way to explore an unmanageably large set of options.

Remember—to walk randomly among the district maps, you need to know which map you can stumble into from which other map, which is to say you need to know which maps are *near* which other maps. We are back to geometry, but geometry of a very high and conceptual kind: not the geometry of the state of Wisconsin, but the geometry of the collection of all ways to break up that geometry into ninety-nine pieces. That is the geometry the mapmakers explored in order to find their gerrymander, and it's the geometry mathematicians have to map out in order to show what a gruesome outlier the gerrymander is.

There's no controversy about what geometry to use on the state of Wisconsin itself. Madison is close to Mount Horeb, Mequon is close to Brown Deer. For the higher geometry of the space of all districtings, you have a lot of choices; and it turns out these choices matter. My favorite is the one that Duchin developed together with Daryl DeFord and Justin Solomon, called the *ReCom* geometry, short for "recombination." The random walk on that geometry works like this.

1. Randomly choose two districts in your map that border each other.
2. Combine those two districts into one double-sized district.
3. Make a random choice among the ways of splitting that double-sized district in half, yielding a new map.
4. Check if the map you made violates any legal constraints; if it does, go back to 3 and choose a new splitting.
5. Go back to 1 and start again.

The "split and recombine" (or "ReCom") move of steps 2 and 3 is to districting what shuffling is to a deck of cards. And, as with the cards, you can explore lots and lots and lots of different configurations with just a few moves. It's a small world. You can randomize a deck with seven shuffles. Seven ReComs is not, unfortunately, quite enough to explore the space of districtings. A hundred thousand ReComs seems to do the trick; that sounds like a lot, but it's tiny compared to the problem of sorting through *all* the districtings one by one. You can ReCom a hundred thousand times on your laptop in an hour. That gives you a sizable ensemble of neutrally drawn maps, to which you can compare the map you suspect of being gerrymandered.

* If you must know, it's a kind of geometry of all two-dimensional geometries named after, though by no means entirely developed by, one of early-twentieth-century math's most fervent Nazis.

The point of the ensemble method isn't to eliminate partisan gerrymandering entirely, any more than the point of *Reynolds v. Sims* was to require districts to have equal population down to the very last voter. Every decision made by a map-drawer, from incumbent protection to promotion of competitive races, might have partisan impact. The goal isn't to enforce an impossible absolute neutrality, but to block the very worst offenses.

Think back to Tad Ottman's speech to the Republican legislators, about the "obligation" the party had to seize the opportunity to cement control. If your job is to get and hold a majority of the legislature, and the law allows you to play as dirty as you like, then dirty is your duty. Blunting gerrymandering's power, establishing that there's some level of unfairness democracy won't tolerate, would have a healthy effect on the whole process. Politicians would be more likely to make reasonable compromises if the rewards of gerrymandering weren't so very great. If you don't want kids to shoplift, maybe don't leave *so* many candy bars *so* close to the front door.

THE TRIUMPHANT RETURN OF GRAPHS, TREES, AND HOLES

I could gloss over the part of ReCom where you split the double-sized district in two, but I won't, because talking about it lets me bring back two characters from earlier in the book. First of all, the voting wards in a district, like stars in movies or atoms in a hydrocarbon, form a network, or as James Joseph Sylvester would have called it, a graph: the vertices are the districts, and two vertices are connected just when the corresponding districts border each other. If the wards look like this:

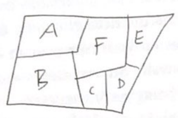

the graph looks like this:

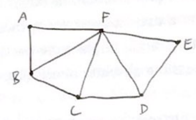

We need to find a way to split up the wards into two groups, and we need to make sure that each group of wards forms a connected network on its own.

Putting A, B, and C in one group and D, E, and F in the other works fine:

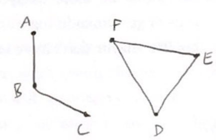

But grouping C, D, and F leaves you with A, B, and E, which don't form a connected district.

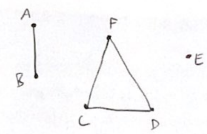

We're standing on the lip of a whole bubbling caldera of graph theory here. John Urschel, an offensive lineman for the Baltimore Ravens, dumped his professional football career in 2017 to work on this stuff, which is what he'd really wanted to do all along. One of his very first papers after leaving football was about how to split graphs into connected chunks using the theory of eigenvalues we met in chapter 12.

There are a lot of ways to split a graph. When the graph is as small as this one, you can list all the splittings, and choose one from the list at random. But listing all the possible splittings gets complicated as soon as the graph gets even a little bigger. There's a trick to picking one at random, and it involves more old friends. Suppose Akbar and Jeff play a game; they take turns removing an edge from the network, and whoever breaks it into disconnected pieces loses. In the graph above, Akbar might remove the edge AF, and then Jeff could remove DF, and then Akbar could remove EF (but not AB, because then he'd disconnect A and lose!) and Jeff could remove BF, and now Akbar is stuck; any edge he erases is going to break the graph into two disconnected pieces.

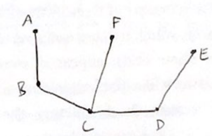

Could Akbar have played smarter and won? No, because this game has a secret feature: as long as neither player blunders and disconnects the network unnecessarily, it doesn't matter which moves you make; the game will always end after four turns, with Akbar the loser. In fact, no matter how big the network, the number of moves in this game is fixed. There's even a nice formula for it: it's number of edges – number of vertices + 1

At the beginning of the game, there are nine edges joining the six districts, so you get 9 – 6 + 1 = 4. When the game ends, with only five edges left, that number has dropped to 0. And what remains of the network has a very special form; there's no way to trace out a closed loop in the graph, as you could in the original graph by traveling in a cycle from A to B to F back to A. If there were a loop like that, you could remove one of the edges in that cycle without disconnecting the graph. What's left is a graph with no cycles at all; and a graph with no cycles is a tree.

How many holes are there in a network? That is, in its way, a confusing question, just like the question of how many holes there are in a straw or a pair of pants. But for this question I've already told you the answer; it's that very number just above, edges minus vertices plus one. Every time you snip an edge out of a cycle, you're getting rid of a hole. When you can't snip anymore, you've ended up with a graph with no holes in it at all: a tree. This is not just a metaphor; there is a fundamental invariant of any kind of space, called its *Euler characteristic*, which very, very, very roughly tells you the number of holes in it.* We've seen it once before, when we were counting holes in straws and pants. Straws have Euler characteristics and so do networks and so do string-theoretic models of twenty-six-dimensional spacetime; it's a unified theory that covers geometries from the modest to the cosmic.

So we're back to the geometry of trees. A tree like the one at the end of the edge-cutting game, which touches every vertex of our network, is called a *spanning tree*. These things appear all over math. A spanning tree of a square grid network like Manhattan streets is something you've seen before: it's called a maze. (In this picture, the white lines are the edges. If you've got a pencil you can convince yourself the maze is connected; you can draw a path from any point to any other without leaving the white lines. In fact, there's only *one* route you can take without backtracking. It's my book, I give you permission to write in it.)

* Only slightly less roughly, it's more like "the number of even-dimensional holes minus the number of odd-dimensional holes." If you have the appetite to learn what it *really* is, see Dave Richeson's book *Euler's Gem*.

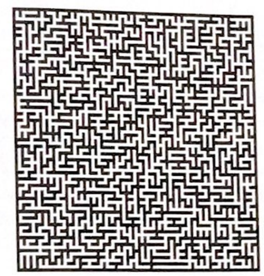

Or you can draw the spanning tree with vertices as dots and edges as line segments, more like the way we drew the districting graph:

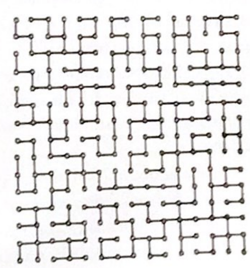

Most graphs of any decent size have a lot of spanning trees. The nineteenth-century physicist Gustav Kirchhoff worked out a formula that can tell you exactly how many, but that doesn't come close to answering all the questions they present, and a century later these trees are still an active research area. There is regularity and structure. For instance: How many dead ends does a random maze have? Of course there will be more and more dead ends the bigger the maze is, but what if we ask what proportion of locations in the maze are cul-de-sacs? A very cool theorem of Manna, Dhar, and Majumdar from 1992 shows that this proportion doesn't grow to 1 or drop to 0 as the maze gets large—instead, it just gets closer and closer to, for some reason, $(8/\pi^2)(1 - 2/\pi)$, just under 30%. You might think that the number of spanning trees of a random graph would be a more or less random number. Not so. My colleague Melanie Matchett Wood proved in 2017 that, if your graph is chosen at random,* the number of spanning trees is slightly more likely to be even than to be odd. To be precise, the chance the number of spanning trees is odd is the infinite product

$$(1 - 1/2)(1 - 1/8)(1 - 1/32)(1 - 1/128)\ldots$$

where the denominator of each fraction is four times that of the previous one. (Geometric progressions again!) The product comes to about 41.9%, quite a distance from 50/50. This asymmetry is the signature of a deeper geometric structure on the collection of all spanning trees; it turns out, for example, that there is a meaningful way of saying when a sequence of spanning trees forms an arithmetic progression!

But to explain this I would have to delve into the fascinating details of the "rotor-router process," and we still haven't saved democracy. So let's get back to districts.

Once you have a spanning tree in hand, there's an easy way to cut the network into two parts; just make the game-losing move and cut an edge, disconnecting the graph. Any choice you make will break the graph in two; with a little work, you can usually find an edge that makes the two pieces roughly equal in size. (If you can't, pick another tree and start over.) It looks like this, where the two sections are the part of the tree I've scribbled over and the part I haven't.

* In the sense of Erdős and Rényi from chapter 13.

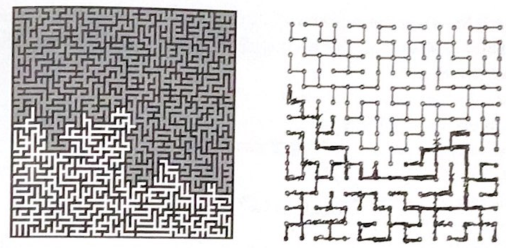

And now you know more or less how ReCom works.* You take your double-sized district, you choose a spanning tree at random—for example, by playing the edge-cutting game with randomly chosen moves†—pick a random edge in that tree, cut it, and the graph cleaves neatly into your two new districts.

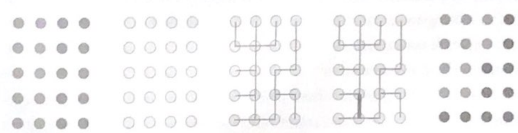

I'd better stop here for a caveat. There's a big difference between a random walk by ReCom on the space of maps and the random walk by shuffling on the space of ways to order a deck of cards. In the latter case, we have the seven-shuffle theorem, where by theorem I mean *theorem*; there is a mathematical proof that a certain number of shuffles (six!) is sufficient to explore every possible ordering, and what's more, that given just a few shuffles (seven!) every single order is just about equally likely.

When it comes to districtings, there are no theorems. We know

* And if you want to know more rather than less, see the 2021 book *Political Geometry*, edited by Duchin along with Ari Nieh and Olivia Walch.

† Though if you really want to get each spanning tree with equal probability, you have to be a little more intentional about how you choose which edges to cut; or you can follow DeFord, Duchin, and Solomon and use *Wilson's algorithm*, which is also a bit faster.

much less about the geometry of districtings than we do about the geometry of shufflings. The space of all districtings might, for instance, look like this:

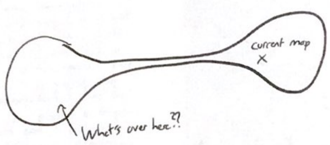

in which case, if you started at one end, you could conceivably randomly wander around for a long time before ever starting to explore what's on the other side of the isthmus. Or, for all we know, the space of all districtings could split into two separate connected pieces, or more than that. There might be an undiscovered country of possible maps of North Carolina, radically different from any ever contemplated by mathematicians, computers, or unscrupulous politicians, and perhaps among *those* maps, having ten out of thirteen Republican seats is quite typical. If we can't rule that out, do we still have the right to say the current gerrymandered map is an outlier?

Yes, as far as I'm concerned. We may not know with absolute certainty whether or not some secret reservoir of alternate maps exists; but we know that, in practice, if you start with the map the North Carolina legislature made, and mess with it, it gets less Republican no matter what you do to it. That experiment provides a strong indication, in any meaningful statistical sense, that the map was cooked. It is not a *proof* that the map-drawers set up to rig the game. For that matter, neither are the emails and memos from the map-drawers directly asserting they were trying to rig the game; after all, there is no Euclidean demonstration that they didn't actually mean to type "Let's get down to business and draw district maps that impartially capture the will of the people" but their fingers slipped and "Let's gerrymander this state so bad we can't possibly lose" came out instead. It's a proof in the sense of law, not in the sense of geometry.

UNITED STATES V. TUNA MELT

The ensemble of maps produced by random walks was at the very heart of the gerrymandering cases the Supreme Court heard in the spring of 2019. The point wasn't to prove that the maps had been drawn with partisan intent; that question wasn't in dispute. Thomas Hofeller, the North Carolina map's architect, had already testified that his aim had been "to create as many districts as possible in which GOP candidates would be . . . successful" and "to minimize the number of districts in which Democrats . . . [could] elect a Democratic candidate." The question was: Had the plan worked? You can't throw out a map for merely *trying* to be unfair. You have to prove it actually was.

The ensemble method is the best tool we have for that. Older ideas, like the efficiency gap, were largely absent from the plaintiffs' request. What they were asking the court was to recognize that the North Carolina map was an outlier, as out of place among its neutrally drawn counterparts as a warthog in a litter of piglets. This outlier analysis, they argued, is the "manageable standard" the court had been seeking. Jonathan Mattingly, a Duke mathematician and a member of the team that made the ensemble of Wisconsin assembly maps, had done the same for North Carolina's congressional districts; he testified that in his ensemble of 24,518 maps, there were just 162 in which Republicans won ten districts. The existing map packed North Carolina's Democrats so efficiently into three districts that the Democratic vote shares there were 74%, 76%, and 79%; not a single one of the 24,518 simulated maps created any districts that lopsided.

The mathematicians' brief made a similar point, though ours had prettier graphs.

And then came the oral argument, which for everyone who had been watching the case through a mathematical lens was a monumental bummer. It was as if years of progress and research on districting had never happened, and we were back to the stale question of whether 55% of the statewide vote should guarantee 55% of the seats in the legislature. Paul Clement, defending the North Carolina maps, got it started, telling Sonia Sotomayor: "I think you've put your finger on what my friends on the other side perceive to be the problem, which is a lack of proportional representation." Justice Sotomayor tried to tell Clement that her finger was elsewhere, but he pressed on, saying to Stephen Breyer, "[Y]ou can't talk even generally about outliers or extremity unless you know what it is you're deviating from. And I take it, implicit in your question and implicit in Justice Sotomayor's question, that what's bothering people is a deviation from a principle of proportional representation." "Actually—" Sonia Sotomayor broke in. "You keep saying that, but I don't quite think that's right," Elena Kagan objected. That didn't stop Clement, who had a point to make about Massachusetts. Republicans never have a congressional representative in Massachusetts, he observed, though Republicans make up a third of the state's population. "Nobody thinks that's unfair, because you really can't draw districts to do it because they're evenly distributed. It might be unfortunate for them, but I don't think it's unfair."

The plight of the Massachusetts Republican, as it happens, is also discussed in the mathematicians' brief. Our account of it basically matches Clement's, apart from one important detail; what he says the plaintiffs are asking to enforce is in fact what the plaintiffs are asking to forbid. It's *not* unfair that Republicans in Massachusetts don't get proportional representation. You can make an ensemble of maps, thousands of them, drawn with no nefarious partisan intent, and *every single one of them* will send nine Democrats and zero Republicans to Congress. That's why Common Cause wasn't asking the Supreme Court to guarantee proportional representation. Proportional representation is a lousy criterion for fairness. A map in Massachusetts that resulted in proportional representation wouldn't be immune from accusations of gerrymandering; it would, in fact, be a gerrymander just as bad as Joe Aggressive.

But many of the justices persisted in treating proportional representation as the issue they were being called on to decide. Neil Gorsuch worried that, if he decided against North Carolina, "We're going to have to, as part of our mandatory jurisdiction, in every single redistricting case, look at the evidence to see why there was a deviation from the norm of proportional representation. That's—that's—that's the ask?"

It was not the ask. This seemed hard for Gorsuch to accept. There's a really astonishing exchange near the end of the oral argument between Gorsuch and Allison Riggs, representing the League of Women Voters in opposition to the gerrymandered map. Riggs is explaining that her client is asking the court only to throw out the most freakishly outlying gerrymanders, whose partisan performance sets them apart from all but a handful of neutral alternatives. States would then still have a lot of breathing room to choose from the other 99% of all maps with a free hand, taking whatever nonpartisan criteria they liked into account. Gorsuch breaks in:

> ***Justice Gorsuch:*** —but with—with respect, counsel, and I'm sorry to interrupt, but breathing room from what?
>
> ***Ms. Riggs:*** Breathing room to—
>
> ***Justice Gorsuch:*** From—how much breathing room, from what standard? And isn't the—isn't the answer that you just—I understand you don't want to give it, but isn't the real answer here breathing room from proportional representation up to maybe 7 percent?
>
> ***Ms. Riggs:*** No.

After some more tussling, Gorsuch seems to concede that Riggs is not going to accept his enforced paraphrase. "We need a baseline," he says. "And the baseline, I still think, if it's not proportional representation, what is the baseline that you would have us use?"

He was asking a question that had been answered, a moment before, by Elena Kagan: "[W]hat's not allowed is deviation from whatever the state would have come up with, absent these partisan considerations."

Reading through the transcript of the argument, for a mathematician, feels like teaching a small seminar where only one student has done the reading. Justice Kagan gets it. She delivers a clear and succinct paraphrase of the quantitative argument she's being asked to consider. And then . . . everybody carries on as if she hadn't said a word. Sonia Sotomayor and John Roberts don't say much, but what they say is mostly right. Stephen Breyer has his own gerrymandering test, which neither side likes much. And Gorsuch, Samuel Alito, and to some extent Brett Kavanaugh, with help from Paul Clement, collaborate on building a fictional version of the case in which the plaintiffs are asking the court to impose some form of proportional representation on the states.

If you need a break from ensembles, random walks, and outliers, here's how the oral argument would have gone if it had been about ordering a sandwich instead:

> ***Ms. Riggs:*** I'd like a grilled cheese.
>
> ***Justice Alito:*** Okay, one tuna melt.
>
> ***Ms. Riggs:*** No, I said grilled cheese.
>
> ***Justice Kavanaugh:*** I hear the tuna melt's good.
>
> ***Justice Gorsuch:*** You want that tuna melt open face or closed?
>
> ***Ms. Riggs:*** I don't want a tuna melt, I want a—
>
> ***Justice Gorsuch:*** It seems like you don't want to just come out and say it, but don't you want a tuna melt?
>
> ***Ms. Riggs:*** No.
>
> ***Justice Kagan:*** She asked for a grilled cheese. That's not a tuna melt because there's no tuna in it.
>
> ***Justice Gorsuch:*** But if, as you say, you don't want a tuna melt, what do you want? Are we supposed to just make up a sandwich for you?
>
> ***Justice Alito:*** You come in here, you ask for a hot sandwich on toasted bread with cheese on it. That, to me, is gonna be a tuna melt.
>
> ***Justice Breyer:*** Nobody ever orders the chopped liver, but have they really given it a chance?
>
> ***Mr. Clement:*** The framers had every opportunity to make you a tuna melt, but they chose not to.

Maybe you already know how this ended up, and if you don't know, you can probably guess. On June 27, 2019, the Supreme Court ruled in a 5–4 decision that it was outside the scope of the federal courts to decide whether a partisan gerrymander was constitutional or not; in technical terms, the matter is "nonjusticiable." In lay terms, states can

gerrymander their legislative maps with unlimited wanton abandon. Chief Justice Roberts, writing for the majority, explained:

> Partisan gerrymandering claims invariably sound in* a desire for proportional representation. As Justice O'Connor put it, such claims are based on "a conviction that the greater the departure from proportionality, the more suspect an apportionment plan becomes."

Roberts does concede, far down in the decision, that proportional representation isn't what the *Rucho* plaintiffs are asking for, but much of what he writes is devoted to reiterating his opposition to that unasked-for requirement. The constitutional constraint that no one's voting power be diluted, he insists against no opposition, "does not mean that each party must be influential in proportion to the number of its supporters."

No, I won't make you a tuna melt. You know we don't serve tuna melt here!

I'm not a lawyer, and won't pretend to be one. Nor will I pretend the constitutional issues in this case were easy. Easy cases don't make it to the Supreme Court. So I'm not going to tell you that the majority ruled wrongly here as a matter of law. If that's what you're looking for, I recommend Justice Kagan's dissent, which is so sardonic and bleak it seems at times about to break out in bitter laughter.

For Roberts, it's crucial that a certain amount of partisan bias in redistricting has been explicitly allowed by courts in the past. The question before the court was whether, at some point, there is such a thing as *too much*. The majority in *Rucho* said no: if it's constitutional to do it, it's constitutional to overdo it. Or, more precisely, if the court can't find a clear, agreed-on universal line to draw between permissible and forbidden, then the court can't consider the matter at all. It's a legal version of the *sorites* paradox, which goes all the way back to Eubulides, an old sparring partner of Aristotle. The *sorites* paradox asks us to work out

* As best I can tell from my lawyer friends, "sound in" here means something in between "derives from" and "amounts to"—and people say *mathematicians* talk in impenetrable jargon!

how many grains of wheat it takes to make a heap (in Greek, a *soros*). A single grain isn't a heap, and neither is two grains, surely. In fact, no matter how much wheat is on the table, it's impossible to imagine a situation where adding one grain of wheat to something that's not a heap of wheat results in something that is. So three grains of wheat isn't a heap, and neither is four, and so on. . . . Carrying this argument to its limit shows there's no such thing as a heap of wheat; yet somehow heaps of wheat exist.*

Roberts sees gerrymandering as unavoidably soritical. Any line drawn between "acceptable gerrymandering" and "I'm sorry, but it's just too much" will inevitably be arbitrary, he says, and will eventually be complicated and case-dependent (he might be satisfied with a rule that ninety-nine grains isn't a heap but a hundred is, but not if the threshold depended on whether the grains were wheat or sand).

I take his point. And yet I keep thinking about Nevada. Alone among the fifty states, it has no contiguity requirement for its legislative districts at all. In principle, the legislature in this Democratic-leaning state could fill three of the twenty-one state senate districts entirely with registered Republicans, and balance the party composition of the remaining districts to be just about exactly 60% Democratic, making them almost certainly safe seats and locking in a veto-proof 18–3 supermajority in the upper chamber, which would persist even if the state swung quite a bit to the right and elected a Republican governor. By the reasoning of *Rucho*, there is no clear-cut way to identify such a plan as "too much." Sometimes legal reasoning—even if it's sound as a matter of law—parts ways from common sense.

The majority's ruling, in the end, turns on a technical point—that partisan gerrymandering is a "political question," which means that, even if the Constitution *has* been violated, the Supreme Court is forbidden to intervene. That the results of gerrymandering "reasonably seem unjust"—that they are, in fact, "incompatible with democratic principles"—isn't in dispute; those are in fact quotes from the majority decision! And the gerrymanderers' unconvincing protests that their

* The reader is invited to note the resemblance to the idea of proof by induction that showed up when we talked about Nim.

maps weren't really *that* good at locking in electoral advantage are dismissed almost without comment. But just because something is unjust and incompatible with democratic principles and fiendishly effective, Justice Roberts writes, doesn't mean it's within the purview of the court to find a constitutional violation. Gerrymandering stinks, but not so badly the Constitution can smell it.

You can feel the discomfort in the decision: not just in the concessions that gerrymandering impedes democracy, but in the fervently expressed wish that somebody other than the justices of the Supreme Court will do something about it. Maybe something will be discovered in state constitutions that forbids gerrymandering, Roberts writes. Or, if not, maybe voters in afflicted states will rise up and change the system by ballot referendum, if they happen to live in a state where the legislature can't immediately reverse the outcome of such a vote. Maybe the U.S. Congress will do something about it, who knows?

I imagine Roberts as a worker who, on his way out of the factory at 5:05, notices that the building has caught fire. There's a fire extinguisher right there on the wall—he *could* grab it and spray foam all over the problem, but wait a minute, there's a principle at stake here. It's after five and he's not on the clock. Union rules are pretty clear that he's not supposed to put in unpaid overtime. If he puts out *this* fire, he sets a precedent; now is he on the hook every time the building catches fire after the whistle blows? There's probably somebody around working late who can put it out. And there is, after all, a fire department—that's who's supposed to put out fires! Admittedly, there's no telling how long it'll take them to show up, and the truth is, this town's fire department is known to be pretty lax about showing up at all. But still—it's officially their job, not his.

"THAT ONLY A POLITICAL CONVULSION CAN OVERTHROW"

For opponents of gerrymandering, the Supreme Court's decision was not the hoped-for happy ending. But it could be a happy beginning.

Justice Roberts, after all, wasn't wrong that there are other possible avenues for reform. Within a year of the decision in *Rucho*, the North Carolina congressional districts were thrown out by a panel of state judges for violating the North Carolina constitution. The Pennsylvania Supreme Court had done the same in 2018 (at which point the governor brought in Duchin, of the ReCom algorithm, to help build new and fairer maps). The House has passed a bill, currently blocked by Senate leadership, that would create nonpartisan commissions for the drawing of U.S. House districts (but not of state legislative districts, over which Congress has no power).*

And the mere existence of a high-profile case brought gerrymandering much more sharply into the public's view. The HBO news-comedy show *Last Week Tonight* ran a full twenty-minute segment on redistricting. Three siblings in Texas's badly malformed Tenth Congressional District made their own gerrymandering board game, Mapmaker, which sold thousands of copies after gerrymandering archenemy Arnold Schwarzenegger boosted it on social media. More people know about gerrymandering than used to, and when people know about it, they don't like it. Fifty-five out of Wisconsin's seventy-two counties, some that lean Democratic, some Republican, have passed resolutions asking for nonpartisan districting.

Voters in Michigan and Utah approved new nonpartisan districting commissions by popular referendum. In Virginia, which had a legislative map gerrymandered by Republicans, a bipartisan group in the legislature managed to pass a constitutional amendment handing control of redistricting to an independent commission. But the state shifted so far left so rapidly that, in 2019, Democrats broke the gerrymander and took over both houses of the legislature. Many members of the newly minted majority, now in the driver's seat for the next census, were suddenly less gung ho for reform.

Justice Kagan's dissent argues that one can't expect too much from the political process. The political process is exactly what gerrymandering seeks to constrict. Maryland's congressional gerrymander, for instance, is still firmly in place, even with a Republican governor; Democrats hold a veto-proof majority in the state legislature and are expected to keep the map as it is.

* At least that's the conventional wisdom; but some lawyers have lately argued that Congress can regulate state legislative districts under its Fourteenth Amendment power.

How would Wisconsin get fairer maps? Wisconsin's state constitution has so little to say about district boundaries (and what it does say is so routinely ignored already) that it's hard to see a court challenge to the current maps succeeding.* Wisconsinites have no way to put a ballot initiative to a popular vote, unless the legislature initiates it, and the legislature likes things the way they are. Wisconsin could elect a new governor who would be sure to veto a GOP-gerrymandered map—in fact, it did just that in 2018. There are rumors the legislature plans to ask state courts to declare that redistricting is the job of the legislature and the legislature only, not requiring the governor's signature, and they may find sympathy in the state judiciary. If that were to happen, it's hard to see how the people of Wisconsin could gain any say in the matter.

In Michigan, the independent redistricting commission has been mired in legal challenges from the state's Republicans since the day it was voted into law by 61% of the state's voters. In Arkansas, redistricting reform group Arkansas Voters First gathered more than a hundred thousand signatures in the middle of a pandemic to get a constitutional amendment on the November ballot; the secretary of state declared the petitions void, because the certification that paid canvassers had undergone criminal background checks was misworded on the relevant form. State politics is full of veto points, so political factions with turf to protect have many ways to shield themselves from the public.

For all that, I'm an optimist. Americans used to shrug at legislative districts of wildly different sizes, saying that's just the way the game is played; now most people I talk to are shocked the practice was ever allowed. We are inclined to dislike what's unfair, and our ideas about

* Though when I once told a retired Wisconsin judge I didn't see how the text of the Wisconsin constitution was enough to support a legal challenge, he fixed me with a world-weary eye and said, "I see you are not a litigator."

fairness are never fully separated from mathematical thinking. Talking to people about the dark art of gerrymandering is a form of teaching math, and math has intrinsic purchase on the human mind, especially when it's twined up with *other* things we care deeply about: power, politics, and representation. Gerrymandering was a huge success when it was done behind locked doors. I'd like to believe it can't persist in an open, well-lighted classroom.

CONCLUSION

I Prove a Theorem and the House Expands

The British architect Herbert Baker, one of the designers of the Indian colonial capital of New Delhi, argued that the new city ought to be built on a neoclassical plan. Architecture of a more indigenous flavor wouldn't suit the empire's goals: "While in this style we may have the means to express the charm and fascination of India," he wrote, "yet it has not the constructive and geometrical qualities necessary to embody the idea of law and order which has been produced out of chaos by the British administration." Geometry can be marshaled as a metaphor for unquestioned-because-unquestionable authority, the mathematical analogue of a natural order centered on the king, or the father, or the colonial administrator. The monarchs of France spent untold *pistoles* laying out formal gardens whose perfect lines converging on the palace represented the unchanging order they took to be axiomatic.

Maybe the purest example of this point of view is the short novel *Flatland*, written by English schoolmaster Edwin Abbott in 1884. It is a story told by a square. (The first editions were published pseudonymously, with "A Square"* as the author.) The book takes place in a two-dimensional world, whose inhabitants, like Sylvester's bookworms, are

* Possibly a math pun—Abbott's full name was Edwin Abbott Abbott, so his monogram EAA could be styled algebraically as "E A squared."