

Blind Justice: Algorithms and Neutrality in the Case of Redistricting

Moon Duchin* Tufts University Department of Mathematics Medford, Massachusetts, USA mduchin@mggg.org Douglas Spencer University of Colorado School of Law Boulder, Colorado, USA douglas.spencer@colorado.edu



Figure 1: One vision of a race-blind VRA: how many majority-Black Congressional districts might be found by a naively race-blind search in each state? Blue mark shows the Black share of each state's 2010 voting age population; the X marks the share of majority-Black districts in the 113th Congress; the shaded range shows the share of majority-Black districts ever observed in a sample of 2 million plans; and the large dot shows the median. An extreme "race-blind" standard would pin VRA liability to comparisons with a typical map from such an ensemble. This would lay waste to the status quo, eliminating the cause of action entirely in all but five states. Figure reproduced from [8].

ABSTRACT

In several areas of law and public policy, there have been longstanding dreams that computers can secure decisionmaking that takes only some things into account, while remaining demonstrably neutral to other factors. In 2022, the U.S. Supreme Court will consider mandating race-neutrality in multiple domains, notably in college admissions and redistricting. In this piece, we clarify the real and imagined uses of computers in redistricting, considering their application for optimization approaches and, more recently, for representative sampling. The current pitch to the Court for a race-blind Voting Rights Act is discussed at length.

^{*}Author order is alphabetical



This work is licensed under a Creative Commons Attribution International 4.0 License.

CSLAW '22, November 1–2, 2022, Washington, DC, USA © 2022 Copyright held by the owner/author(s). ACM ISBN 978-1-4503-9234-1/22/11. https://doi.org/10.1145/3511265.3550440

CCS CONCEPTS

• Applied computing \rightarrow Law, social and behavioral sciences; • Mathematics of computing; • Theory of computation \rightarrow Randomness, geometry and discrete structures; Random walks and Markov chains;

KEYWORDS

election law, redistricting, race-neutrality, law and technology

ACM Reference Format:

Moon Duchin and Douglas Spencer. 2022. Blind Justice: Algorithms and Neutrality in the Case of Redistricting. In *Proceedings of the 2022 Symposium on Computer Science and Law (CSLAW '22), November 1–2, 2022, Washington, DC, USA.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3511265. 3550440

1 NEUTRALITY IN THE LAW

In what ways should the law be neutral? Lawmakers and judges have struggled to answer this question from the earliest days of our Republic. The U.S. Constitution explicitly prohibits "bills of attainder" and ex post facto laws (Art. I §9)—the former would target and punish individuals or groups by statute, without a trial, and the latter would punish acts that were legal at the time they

Moon Duchin & Douglas Spencer

were performed. The motivation for these prohibitions was a new federal Congress that might target its political opponents instead of legislating more neutrally. Indeed, the core attribute of good public policy is its effect on the entire public. In domains as varied as taxes (*McCulloch*), religious freedom (*Church of Lukumi*), and gun regulations (*Printz*), courts have equated generality with neutrality [11]. In other words, generally-applicable laws will tend to survive judicial review while laws that target population subgroups will face more scrutiny in the courts.

In practice, the U.S. Supreme Court has developed a jurisprudence that is highly deferential to *facially* neutral laws, even when the effects are shown to fall very differently on different subgroups. The 14th Amendment guarantees that all citizens of the United States shall be treated equally under the law. The Supreme Court has interpreted this promise as protecting citizens against discriminatory *treatment* but not against disparate *outcomes*. Courts have created barriers to challenging statutes, ordinances, and regulations so long as their text remains neutral (i.e., generally applicable), even when they are not neutral in practice.

The racially discriminatory effects of literacy tests (*Lassiter*), employment promotion policies (*Washington v. Davis*), zoning laws (*Arlington Heights*), voter ID laws (*Crawford*), criminal justice policy (*McKleskey*), and at-large voting systems (*Mobile v. Bolden*) were not sufficient on their own to invalidate problematic laws because the laws were written in general terms and were thus deemed to be adequately neutral by the Supreme Court.

If we accept that neutral rules are those that are general and do not target, then we might be led to conclude that their entailments that is, their consequences and downstream properties—constitute a *neutral baseline* of fair play. Finding the neutral baseline has been particularly elusive in vote dilution cases, where the question presented relates to to the weight of a vote. As Justice Frankfurter staked out his opposition in *Baker v. Carr* (1962): "Talk of 'debasement' or 'dilution' is circular talk. One cannot speak of 'debasement' or 'dilution' of the value of a vote until there is first defined a standard of reference as to what a vote should be worth."¹

Legal scholar Justin Levitt puts essentially the same point in more modern language: "Dilution depends on knowing what the baseline should be. You only know that a drink is diluted when you know it falls outside a normal range of what it should taste like. You only know that a district is diluted when you know it falls outside the normal range of what its composition should be."² Defining that baseline has proven difficult and controversial. But recent advances in mathematics and computing have generated some optimism that algorithms and computers will be able to discover race-blind baselines against which electoral systems can be judged for their dilutive effect on the voting power of racial minorities.

2 RACE-NEUTRALITY IN PARTICULAR

One of the more controversial aspects of the Supreme Court's neutrality jurisprudence is its implications for corrective legislation. Remedial laws are rarely neutral. Indeed, the very purpose of a remedial law is to identify and correct the discriminatory effects of prior laws on protected classes of citizens. For example, recognizing the long history of racial disenfranchisement, Congress enacted the Voting Rights Act (VRA), one of the most powerful remedial federal laws in America's history. The VRA prohibits any voting standard, practice, or procedure "which results in a denial or abridgment of the right of any citizen of the United States to vote on account of race or color." The statute is remarkable in two ways. First, the text is decidedly not neutral as it explicitly provides a remedy for some citizens based on their race. Second, as amended in 1982, the VRA provides a cause of action in cases where voting rules result in discrimination, even if the voting rules themselves are facially neutral. Both of these features of the VRA cut against the Supreme Court's neutrality jurisprudence and have thus made it clear that the Roberts Court-increasingly vocal about its desire for race-neutrality-will revisit current practice or even upend the law entirely in the coming years.

During its tenure, the VRA has survived several legal challenges that have raised questions about its non-neutral structure. The Supreme Court has repeatedly held that the VRA is a constitutional exercise of Congressional authority. The Court's position has drawn on both textual and historical elements.

- (1) The 15th Amendment has a specific focus on voting rights, explicitly granting that "Congress shall have power" to protect the rights of racial minorities "by appropriate legislation."
- (2) Race discrimination has an outsized role in American history, policy, and jurisprudence. Race has been a consistent blind spot for the Court, with signal cases like *Prigg v. Pennsylvania*, *Dred Scott, Plessy*, and *Korematsu* widely acknowledged to have been wrongly decided, to devastating effect.

Taken together, these create the conditions to conclude that, first, a law that explicitly acts on the basis of race in voting is less problematic than other facially non-neutral laws. And, second, that a strengthening of the tools to root out racial discrimination may be appropriate for the court, such as by providing a cause of action based on discriminatory effects (without a showing of intent).

This historical and textual confluence has led to a half-century in which race-consciousness is licensed but in extremely limited ways. The formalization of the limitations comes through the framework of *strict scrutiny*, the standard that requires race-conscious government action to be *narrowly tailored* to achieve a compelling state interest.³ These terms of art enter the discourse through a varied lineage of cases covering internment (*Korematsu*), school segregation (*Bolling v. Sharpe*), anti-miscegenation laws (*McLaughlin v. Florida, Loving v. Virginia*), government contracting (*Richmond v. J.A. Croson*), free speech (*Grayned v. Rockford*), and affirmative action (*Regents of the University of California v. Bakke*). Once strict

¹He went on to presciently warn that finding this baseline was a problem of some mathematical depth: "One of the Court's supporting opinions, as elucidated by commentary, unwittingly affords a disheartening preview of the mathematical quagmire (apart from divers judicially inappropriate and elusive determinants) into which this Court today catapults the lower courts of the country without so much as adumbrating the basis for a legal calculus as a means of extrication."

²Justin Levitt, Symposium: Clarity of the record should bring clarity of purpose, SCO-TUSblog (Feb. 11, 2019, 2:07 PM), https://www.scotusblog.com/2019/02/symposiumclarity-of-the-record-should-bring-clarity-of-purpose/

³Strict scrutiny is triggered by policy that curtails fundamental rights or is addressed to certain protected classes, mainly those based on race and citizenship.

scrutiny is invoked by the Court, it brings with it a very high burden of proof that a law meets a state interest in a precise and not over-broad manner.

The arguments justifying a narrowly tailored attention to race in VRA practice have proven successful in the past, but times are changing. The current Supreme Court has articulated a very strong commitment to neutrality principles in the law, and especially so when it comes to legislative efforts to remedy historical race discrimination. In 2007, Chief Justice Roberts memorably asserted that "The way to stop discrimination on the basis of race is to stop discriminating on the basis of race" (Parents Involved). The Court has since signaled a similar concern with the crux of the Voting Rights Act itself. In Shelby County (2013) the Court severely curtailed the scope of Congress's remedial authority under the 15th Amendment. The Court argued that proper use of the VRA must be more narrowly tailored than prior Courts had allowed-not just to legitimate state interests, but to the specific goals of the 15th Amendment itself. In Brnovich (2021) the majority chipped away at the effect standard. More narrowing looms.

The 1980s amendments and decisions paved the way for the use of the VRA in litigation about redistricting, or the formation of electoral districts. While the original case that set out the framework for voting rights litigation (*Gingles*) involved a claim of vote dilution in an at-large voting system, Justice O'Connor's concurring opinion in that case highlighted that the harm of vote dilution was a risk in *any* voting system:

"The phrase 'vote dilution,' in the legal sense, simply refers to the impermissible discriminatory effect that a multimember *or other districting plan* has when it operates to cancel out or minimize the voting strength of racial groups... Put simply, in order to decide whether an electoral system has made it harder for minority voters to elect the candidates they prefer, a court must have an idea in mind of how hard it "should" be for minority voters to elect their preferred candidates under an acceptable system. ...no reason appears why this test would not be applicable to a vote dilution claim challenging single-member as well as multimember districts" (emphasis added).

Starting in the 1990s with a series that is known as the Shaw line of cases, the Court held that lawmakers who engage in redistricting must avoid drawing districts where race is the predominant factor in deciding the district's shape. If a court determines that race was the predominant factor, powerful justification will be required to avoid a finding of excessive race-consciousness, and in practice the challenged district will almost certainly be struck down. The purpose of this rule was to prevent racially motivated gerrymandering, or the disadvantaging of some groups on the basis of race, but the logic has proved challenging under the VRA. If the votes of racial minorities were diluted by a redistricting scheme, how should one draw alternative plans for litigation? Whether drawn for demonstrative purposes or as potential remedies, wouldn't new plans tend to rely on race as a predominant factor? To date, redistricters have responded in two ways. The first response is that drawing a remedial district does not require reliance on race as the predominant factor because drawing remedial districts also requires

line-drawers to pay great attention to other principles such as population, compactness, contiguity, and the preservation of political subdivisions. The second response is that, even to the extent that race is the predominant factor in the design of a demonstrative or remedial district, courts should not reject the district because states have a compelling interest in complying with the VRA to eradicate racism in American politics.

These responses have proven to be less persuasive to the current Justices on the Supreme Court. On February 7, 2022, the Supreme Court heard an appeal from the State of Alabama arguing that it should not be required to draw a second Congressional district in which Black voters could plausibly elect a candidate of choice, because it would be impossible to do so without relying on race as the predominant factor.⁴ A lower court had held that the VRA required the state to draw a second district. The Supreme Court agreed to hear the appeal and stayed (paused) the lower court's ruling in the meantime. In the order granting the stay, three Justices signaled their disapproval of the above responses to the Shaw cases. Alabama advanced multiple arguments in asking for a stay, but one primary focus was the execution of a preliminary matter in voting rights cases, what is called the "Gingles 1" factor. Before a court investigates the social and historical conditions related to a discriminatory voting system, plaintiffs must demonstrate that it is possible to draw an additional majority-minority district in the jurisdiction. According to Alabama, this threshold demonstration was presented as excessively race-conscious. In essence, Alabama advanced the point of view that no tailoring is narrow enough; on this view, it is impermissible to use race data at all, even in a technical step that is used to advance a voting rights challenge.

At the heart of Alabama's argument was a flashy assertion: that plaintiffs' experts themselves had generated thousands, even millions, of districts with fully race-blind randomized computer algorithms, never finding any plans with two majority-Black districts.⁵ This moves the goalposts from producing an example by hand to producing an example by a fully transparent algorithm and suggests the possibility of asking for new goalposts that are so distant they can't be seen from the current state of the law: would mere examples from a fully race-blind process be enough, or would the new standard be to show that such examples are somehow typical? The suggestion that a demonstration plan should be typical of random, race-neutral plans—henceforth, "the Alabama ask"—breaks sharply with recent judicial skepticism. Algorithms, often mistrusted by the Court for their dubious manageability (see *Rucho*), are now called on to make good on blind justice.

⁴Note: one of us is the Gingles 1 expert for the Milligan plaintiffs in that case. ⁵From the State's application for a stay: "Well-established in the court below, no race-neutral map drawer would draw that map. In a sample of more than two million race-neutral maps generated by *Plaintiffs' own experts*, not even one contained two majority-black districts. There is no better evidence that the first precondition for a vote dilution claim has not been met here. See *Thornburg v. Gingles*, 478 U.S. 30, 50-51 (1986). A second majority-black district that can be drawn only by initially subverting race-neutral redistricting criteria to a "non-negotiable" racial target is not a "reasonably configured" district. *Cooper v. Harris*, 137 S. Ct. 1455, 1470 (2017). Accordingly, no invocation of the VRA can justify, much less require, the race-based redraw of Alabama's race-neutral map."

3 THE FIRST WAVE OF COMPUTER REDISTRICTING: FINDING EXAMPLES

In part, computers have seemed like tantalizing partners in the quest for neutrality because they demonstrably only compute with the information that they are provided. Before we can examine whether computers can deliver on dreams of neutrality, we trace key contributions to the history of computer discourse in redistricting.⁶ In the first wave of techno-optimism—which coincides with the Reapportionment Revolution and the passage of the VRA in the 1960s—nearly all of the focus was on what computer scientists call *heuristic optimization*. This means that the algorithms described by early authors relied on narrow quantification of the goals in order to task computers to find good, but not provably the best, plans⁷

The dreams were articulated even before the first rickety programs were constructed. Most authors trace the first call for an algorithmic turn to William Vickrey, writing in Political Science Quarterly in 1961 [13]. Vickrey advocated for randomized algorithms to play a role in the selection of a districting plan. But first he insisted that the algorithms could not find fairness, citing "the absence of any clear criterion as to what a fair result should be" and arguing against elevating random outputs to a definition of fairness. Nonetheless, "if there is to be any attempt at all to purify the electoral machinery in this respect," this means that "procedural fairness" demands a "completely mechanical" process with "no room at all for human choice."

Vickrey thought that randomness was important so that legislators could not simply work backwards from a desired result; he imagined that they should have to choose (and spell out, or *operationalize*) principles only, and then be forced to live with the unpredictable consequence of letting an algorithm follow their choices.

In 1963, James Weaver and Sidney Hess published *A Procedure for Nonpartisan Districting: Development of Computer Techniques* in the Yale Law Journal [14], describing the first working software to automate the drawing of districts.⁸ Weaver and Hess described their work as a method that reluctant courts could turn to to create remedial plans in the case that legislatures refuse to produce an acceptable redistricting plan. The appeal of computers, in their words, was that "a court may desire to limit its own discretion" in order to insulate itself from criticism. Interviewed years later about the project, Weaver described his work this way: "It hit me like the electric lights over people's heads on the comic pages that the computer could create districts which were blind to politics. Not only do they process numbers better than humans, they can't introduce politics if no party affiliation is in *their* memories." (emph. orig.)

In terms of the method itself, they were working in an optimization mentality and introduced an explicit quantitative goal or *objective function*—in this case, a new formula for compactness—that they sought to maximize.⁹ In modern computer science language, this was essentially a *k*-means algorithm employed in an era in which that term had not yet been popularized but the underlying ideas were already mainstreamed. Weaver and Hess wrote: "No available programs or computer techniques are known which will give a single, best answer to the districting problem, though such a solution seems possible if enough funds and efforts are put to the problem, especially considering the rapid advances in size and sophistication of available computers." In brief, then, Weaver and Hess thought that euristic optimization could be used to produce a high-quality plan with less intensive human decisionmaking.

In a slightly different form, the dream was also articulated early by Thomas Nagel in a law review article called Simplified Bipartisan Computer Redistricting from 1965, the year that the Voting Rights Act was passed [10]. Nagel was an academic political scientist and an ardent advocate for applying state-of-the-art computing in public policy. He wrote software designed to "transfer a set of agreed-upon values into a concrete plan" and thereby not only generate usable options for redistricting, but also facilitate compromise and illustrate tradeoffs in values to the stakeholders. Nagel's method was also heuristic optimization, but with a more flexible objective function, which he imagined might take into account multiple factors from compactness to political lean via simple linear combination, such as three times shape eccentricity plus two times partisan skew. His very practical hopes for his algorithm are clear from his article's repeated emphasis on the low cost for punch cards and technical assistance. The fact that he was not likely to locate a global optimum (the best plan of all) was clear to Nagel, but he argued that a local optimum was good enough for government work, so to speak: the algorithm "will not guarantee that the criterion is as low as mathematically possible, though it should be low enough to satisfy the political and judicial powers that be."

By the end of the 1960s, then, multiple visions of how computers could deliver healthy objectivity to redistricting had been sketched and even attempted. But it is notable that each of these formulations required users to specify goals in clear, quantitative terms, and each is geared to producing examples, not to any kind of systematic study of valid plans.

4 GINGLES AND THE ROLE OF EXAMPLES

The intent-vs.-results shift described above came in the early 1980s in the wake of the landmark case *Mobile v. Bolden*. In that case, plaintiffs had challenged a voting *system* (plurality at-large election), claiming it was structurally discriminatory. The Court ruled against plaintiffs for want of intent evidence. Congress sprang into action with a muscular amendment to the VRA in 1982, overriding *Mobile v. Bolden* and clearly spelling out the new effects standard.

In a hugely impactful academic paper the same year, litigators James Blacksher and Larry Menefee operationalized an effects test by proposing three threshold conditions that could be manageably verified [3]. These conditions were adopted wholesale by the Court

⁶The historical overview in this section owes a great deal to the excellent book chapters by historian Alma Steingart and computer scientists Amariah Becker and Justin Solomon [2, 12].

⁷ In particular, the authors generally used hill-climbing-style algorithms to make moves intended to search a state space for local improvements in the value of some objective function. They made no claims, let alone guarantees, of global optimality.

⁸Weaver was a civically engaged engineer; Hess later became a prominent figure in management science and operations research, cycling back and forth from industry to academic jobs. Weaver continued to publish on his efforts at computerized redistricting throughout the 1960s and into the early 70s. See https://library.udel.edu/special/ findaids/view?docId=ead/mss0756.xml;tab=content.

⁹"Compactness" refers to measurements of district shape, and to the preference for more plump and regular districts over more spindly and erratic ones.

just a few years later in *Thornburg v. Gingles* (1986). Since that decision, in order to launch a VRA lawsuit at all, plaintiffs must check three boxes. The first checkbox, now called "Gingles 1," requires a demonstration map that improves on the challenged plan. Having checked the box of Gingles 1, plaintiffs must demonstrate a meaningful divergence in the preferences of the minority and the majority that has served to block the will of the minority—this is called *racially polarized voting* ("Gingles 2-3"), and is usually handled with a small set of more-or-less routine statistical tests. These checkbox elements just get you in the door. If plaintiffs are able to check all three boxes, then the litigation proceeds and the court considers the totality of the circumstances.

For present purposes, Gingles 1 deserves close attention. Its requirement is to show that the minority group under consideration is "sufficiently large and geographically compact to constitute a majority in a single-member district." This has been interpreted as requiring a demonstration by plaintiffs that the state *could have drawn* additional majority-minority districts that are themselves reasonably shaped.¹⁰

But the language still leaves some room for interpretation: does the group have to have more than half of the population? Or does constituting a majority mean that the group can put together a majority of like-minded voters, so while they might be only 45% of a district, say, they have sufficient cohesion, plus crossover support from other groups, to have a meaningful opportunity to elect candidates of choice? The latter kind of district is called by many names: *effective, performing*, or *opportunity* district in view of historical voting patterns, as distinct from one with a *numerical majority*.

In 2009, litigants asked the Supreme Court for clarity about the Gingles 1 checkbox in a case called *Bartlett v. Strickland*: should the demonstration plan add an effective district or one based on a numerical racial threshold? The Court's answer was that the standard for the Gingles 1 trigger (and this trigger only), the new district must have 50% minority population plus one person—a straight-up racial headcount of voting age population. Justice Roberts wrote the decision, and Justice Alito joined him in the plurality that insisted on the 50% mark. Neither in this decision nor in any other did the Court ever downplay the salience of opportunity districts. Indeed, a performance analysis can shield an enacted plan from liability, and the remedial districts ultimately drawn must demonstrate effectiveness, and in both cases electoral performance functions with no special role for the 50% line.

The reference to "opportunity" is key here. Effective districts don't have to be a lock for minority-preferred candidates, but must only be configured so that voting patterns will not *lock out* the minority group's preferences. Once liability is established by the court, the parties can start looking into remedies. Despite the requirement in Gingles 1 to show that a majority-minority district would be possible, the need for majority-minority districts drops out in the remedial phase, and instead the remedy must include districts that demonstrably provide electoral opportunity. Both roles for a demonstration plan, whether focused on numerical majorities or on increased effectiveness, are met by examples, and so these legal needs are potentially quite well aligned with heuristic optimization. But to do this efficiently at scale would require several generations of improvement to the algorithms.

5 THE METHOD OF ENSEMBLES: TOWARDS REPRESENTATIVE SAMPLING

Computerized district generation has had a major second wave in the 2010s, and it has already broken through to litigation. By this time, a shift had taken place in the thinking about the relevance of random districting, with first softer and then more vocal attempts to design algorithms that could deliver a picture of the universe of possibility rather than providing one or a handful of model plans.

In a highly influential paper from 2013 [5], political scientists Jowei Chen and Jonathan Rodden provided the first example known to us where randomized plan generation—"automated districting simulations," in the authors' language—would be used to quantify the bias in a proposed or enacted plan.¹¹ That is, the pool of computer-made plans would provide the comparator, by virtue of being drawn with no partisan intent, delivering a neutral baseline from which to measure deviation. The framing of the paper was also hugely influential: by leveraging technology to explore the measurable properties of neutrally drawn plans, the authors sought a means to understand the effects of *political geography*. That is, residential patterns in voting have a controlling effect on the attributes visible in typical valid districting plans, and line-drawers who seem to be biased might well be perpetrating "unintentional gerrymandering" via the mere consequences of territorial sorting.¹²

The article's empirical content is generated with code that implements an idea quite different from those in Hess and Weaver (*k*-means) or Nagel (iterative flips and swaps). It introduces what we have called a *Petri dish method* (or *agglomeration*) that seeks to fill up a state with districts, either by a kind of flood-fill procedure that grows districts from seeds or by an iterative merging procedure that glues units together in a randomized fashion. In either case, the growth procedure continues until there are the right number of districts of the right size, with possible adjustments made if needed. Chen in particular brought his algorithms to expert work in a line of often-successful mid-decade redistricting cases. Time and again, courts seemed dazzled by the mere use of computers and willing to accept that a sample containing thousands of different maps must give a glimpse at the world of possibility.

Why has mere quantity seemed to be enough to get a sense of possibility? Partly because it is hard to get a handle on the sheer number of alternatives. For instance, Justice Alito questioned plaintiffs' attorneys in the oral argument for *Rucho v. Common Cause*, asking "So you've got – let's say you've got 100 maps or you might even have 25. I think you probably have thousands. So you

¹⁰Interestingly, the language from the Blacksher and Menefee article itself presents some ambiguity about whether the standard is just one example, or many. "[T]he relevant question should be whether the minority population is so concentrated that, if districts were drawn pursuant to accepted nonracial criteria, there is a reasonable possibility that at least one district would give the racial minority a voting majority." The authors even indicate that it would be imaginable, but not necessary, to address this inquiry with a computer.

¹¹Earlier notable work includes papers by Cirincione et al. and by Altman–McDonald from the late 90s to mid 00s. The Altman–McDonald works, particularly, emphasize that generating computer alternatives can help illuminate the intentions of line-drawers. ¹²In particular, it has frequently been observed, even back to Vickrey in 1961, that a very uniformly distributed minority group is sharply disadvantaged in redistricting, since they will tend also to be the minority in every district. For this line of thought, see also [7], which shows that Republicans were so uniformly distributed in Massachusetts between the 2000 and 2010 Census that they were mathematically fenced out of Congressional representation.

have all of these maps, and you have to choose among them." Just to give Justice Alito a size-check on the scale of this problem, let's point out that he is off by some orders of magnitude. For instance, we can precisely enumerate the number of ways to build a valid Minnesota state Senate plan by pairing up the state House districts from the 2010 redistricting: there are 6,156,723,718,225,577,984, or over six quintillion. That is from enormous building blocks; now imagine building from tiny atoms like census blocks, which gives staggeringly greater choice to the line-drawer. This alone should make it clear that when an expert brandishes a sample of plans, the seal of quality is not *how many*, but how they are distributed. And in retrospect, a glaring problem with the randomized agglomeration methods is the lack of any description of how the sampled plans are distributed in the state space of configurations.

A few years behind Chen and Rodden, teams with new methods and more computing power followed on, reviving Nagel's early use of flips and swaps in the form of Markov chain Monte Carlo, or MCMC, algorithms. At this stage it became readily possible to build ensembles with tens of thousands of alternatives, and to target particular distributions on plans that could control the extent to which, for instance, compact districts would be preferred to elongated ones. The introduction of a new tree-based Markov chain called *recombination* by DeFord et al. [6] brought another efficiency breakthrough, and by the time the 2020 Census data was released in August 2021, the recombination method could produce millions of legitimately different plans-at the realistic scale of state redistricting problems-in minutes, together with a solid explanation of how those plans are distributed.¹³ But out in the field, methods are still highly varied. In just the last few years, expert work in litigation has employed agglomeration methods, genetic algorithms, flip-based MCMC, tree-based MCMC, and a sequential Monte Carlo method based on importance sampling. The aims are subtly different, and the details of implementation are wildly different. It is clear that we need to take great care with how to think with ensembles.

6 BLINDING JUSTICE

During the 2020 redistricting cycle, lawmakers in several states repeatedly expressed their goal to draw districts without looking at race at all during the process. For example, the Texas state Senator who led the state's redistricting process said, "I've stated it, and I'll state it again—we drew these maps race blind. We have not looked at any racial data as we drew these maps, and to this day I have not looked at any racial data." The context for this insistence was a keen awareness of prior litigation where overt consideration of race had proved fatal to the districting plans in court. To preempt a similar legal challenge in 2020, these state legislators prophylactically argued that they were drawing districts without looking at racial data at all.

But of course it is not necessary for state legislators to actually look at racial data as such to understand the underlying demographics and the consequences of some boundary shifts. A retiring Republican state Senator in Texas called it "obvious" that the changes constituted a "renewed effort to dismantle SD 10" with Moon Duchin & Douglas Spencer



Figure 2: How many South Carolina state Senate districts have BVAP>50%? Blue: statistics for an ensemble of 17,243,492 plans generated by recombination from 2020 block groups, with no consideration of counties. Red: a similar ensemble of 17,868,528 plans generated from 2020 precincts, with strong county-preservation via MST weights. The frequency of observing 0 or 1 majority-BVAP district is nearly double in the blue ensemble. This example suggests that responsible ensemble practitioners should perform many checks that their results are well-aligned with the law and robust to implementation decisions.

goals that were simultaneously racial and partisan. Indeed, a similar split to Texas Senate District 10 in 2011 had been ruled to be an unconstitutional racial gerrymander.

Across the country, one of the experts hired by North Carolina to defend its maps made a similar point in cross-examination:

Q: And you agree that it would be difficult for elected officials to shut off what they know about race and partisanship of the people in the state when drawing maps, right?

A. It might be impossible unless, you know, you had a mind wipe or something.

Put another way, it is hard to imagine people who are qualified to make decisions on the basis of legitimate considerations like communities of interest and shared economic and infrastructural needs but are able to suppress knowledge of race.

Algorithms are not altogether different. One can remove race as an explicit column in the data that is used to generate an ensemble of plans. But race is woven deeply into the other traditional districting principles, from county and city lines to the human geography of communities. Returning to Alito's questioning in Rucho, he calls out "the so-called neutral criteria - compactness, contiguity, protecting incumbents, if that's really neutral, respecting certain natural features of the geography ... "So Alito clearly recognized that facially neutral criteria may have non-neutral impact. Alabama's recent argument went so far as to insist that preserving the cores of prior districts was a race-neutral priority of the highest order. But, as with incumbent protection, the advantage provided by core preservation depends heavily on the status quo. Far less obvious relationships are revealed by empirical study, as in Figure 2, where facially race-neutral decisions-building from block groups vs. precincts and respecting or ignoring counties-have an appreciable impact on the number of majority-Black districts.

¹³In its simplest form, recombination sampling targets the so-called *spanning tree distribution*, in which the likelihood of choosing one plan relative to another depends only on a certain measure of compactness of the plans. [6]

These points illustrate that it may be difficult to eliminate all racial proxies, but it is equally important to interrogate the assumption that such a move would be desirable. As one of us noted on the stand in recent Pennsylvania litigation, "if you would like a districting plan that doesn't split counties, no one would propose that you shouldn't know where the counties are." As absurd as it might sound to draw county-blind and expect good respect for county boundaries, it is equally unreasonable to expect that ignorance of race data confers racial fairness.

Other visions are possible for harnessing the new methods of district generation. Neutral ensembles can be very useful for highlighting whether a plan behaves as though drawn by articulated rules, and that has proved useful for flagging outliers in partisan properties or in demographic composition. Beyond that, the Markov chain approaches for ensembles can be repurposed to build new heuristic optimizers, searching for interesting examples with vastly more power than before. This suggests a promising use of computational redistricting for demonstrative maps, whether focused on numerical majorities or on effectiveness [1, 4]. But the use of a neutral baseline to root out all signs of race-consciousness may still prove ideologically appealing to the Court.

7 SYSTEMS HAVE CONSEQUENCES

Much of the practice around litigating racial issues in redistricting has been stable for quite some time. The original VRA dates to 1965; the notion that it can function without a showing of racist intent dates to 1982; the Gingles checkboxes date to 1986; racial predominance caveats were amplified in the 1990s; and the 2009 Bartlett decision affirmed that the trigger for litigation includes a district passing a 50% demographic line. The Alabama ask before the court could be intensely destabilizing, perhaps more than most observers realize.

It may be no surprise that blind districts are structurally unfriendly to minority representation, because a group that is in the minority of a state overall is also, all things being equal, likely to be in the minority of an individual district within the state. As Vickrey recognized in 1961, "it is extremely unlikely that any random drawing of the electoral boundaries would produce proportional representation." Political scientist Bernard Grofman amplified this prediction in 1982 in a piece evocatively titled *For single-member districts, random is not equal* [9]. He wrote: "Unfortunately, except under very special circumstances, unlikely to ever be achieved in practice, random districting will not yield proportionality..."

But these authors did not have the means to test these effects in the field, and as it turns out, they massively underrated the fenceout effects of blind districts. It is not just that blind districts fall short of proportionality, but indeed, in the most extreme version of the Alabama ask, they will typically fail to secure *any representation at all.* To see this, we can turn to empirical data to study the mere mathematical consequences of blindly dividing residential territory into compact, equi-populous chunks around the country.

Figure 1, reproduced from [8], summarizes the results of districting with 2010 data at the Congressional level, where runs of length two million have been used to explore the blind baseline.¹⁴ The

twenty states featured in the figure include all states with at least 15% Black population overall. For those states, the proportional share of representation would amount to 40 seats in Congress. If a blind Gingles 1 standard were adopted, keyed to the most majority-Black districts ever encountered by a naively race-blind process, then the number of such districts would drop to just 14. But in the most radical version of events, a mere existence proof will no longer suffice to check the box for VRA liability, and courts will demand to see that additional majority-minority districts are typical rather than possible. For instance, we might imagine requiring that the median number of majority-minority districts exceeds the challenged plan. In this blind Gingles world, the cause of action for a VRA case nearly disappears: only in five states could any VRA lawsuit have gone forward on behalf of Black voters. That is, the total number of Black opportunity districts subject to any legal protection by the VRA would drop to not 40, not 14, but just 5 districts nationwide. The median standard would foreclose the use of the VRA to challenge securely all-White districting plans in Florida, Louisiana, Mississippi, New Jersey, North Carolina, South Carolina, Virginia.... and Alabama.

This outcome would truly represent a turning on its head of the logic recommending the use of districts in a representative democracy. Even Justice Frankfurter, long-term skeptic of the justiciability of redistricting abuse, articulated the animating principle of districts in a way that makes this clear. In his words from *Colegrove* in 1946: "The upshot of judicial action may defeat the vital political principle which led Congress, more than a hundred years ago, to require districting. This requirement, in the language of Chancellor Kent, 'was recommended by the wisdom and justice of giving, as far as possible, to the local subdivisions of the people of each state a due influence in the choice of representatives, so as not to leave the aggregate minority of the people in a state, though approaching perhaps to a majority, to be wholly overpowered by the combined action of the numerical majority, without any voice whatever in the national councils."

Representative democracy is incompatible with the total fenceout of minorities, and this must guide the selection of *systems* of election. Single-member districts holding plurality elections constitute a system of election; we have only recently gained the ability to understand the consequences of benchmarking districts to raceblind drawing—that is a fundamentally empirical question that depends heavily on the detailed patterns of clustering and dispersal across the landscape of each state. But now we know: in nearly every state, a race-blind standard would put plans that reliably eliminate minority representation entirely out of the reach of the law.

ACKNOWLEDGMENTS

Our sincere thanks go to Parker Rule, Max Fan, and the rest of the MGGG Redistricting Lab. We also thank Guy Charles, Sam Hirsch, and anonymous referees for extremely helpful comments.

¹⁴These are runs using block group units with a recombination Markov chain that enforces 2% population balance and contiguity while up-weighting compact plans.

As should be clear from Figure 2, there are many parameters to vary to get a fuller picture, but the basic finding that blind districting fences out even sizeable minorities is quite robust. See [8] for pointers to the codebase and many more details about the algorithm.

CSLAW '22, November 1-2, 2022, Washington, DC, USA

Moon Duchin & Douglas Spencer

REFERENCES

- Amariah Becker, Moon Duchin, Dara Gold, and Sam Hirsch. 2021. Computational Redistricting and the Voting Rights Act. *Election Law Journal* 20, 4 (2021), 407– 441.
- [2] Amariah Becker and Justin Solomon. 2022. Redistricting Algorithms. In *Political Geometry*, Moon Duchin and Olivia Walch (Eds.). Birkhäuser Books, Chapter 16, 303–340. http://mggg.org/gerrybook
- [3] James U. Blacksher and Larry T. Menefee. 1982. From Reynolds v. Sims to City of Mobile v. Bolden: Have the White Suburbs Commandeered the Fifteenth Amendment. Hastings Law Journal 34, 1 (1982), 1–64. https://repository.uchastings.edu/ hastings_law_journal/vol34/iss1/1/
- [4] Sarah Cannon, Ari Goldbloom-Helzner, Varun Gupta, JN Matthews, and Bhushan Suwal. 2020. Voting Rights, Markov Chains, and Optimization by Short Bursts. preprint (2020). https://arxiv.org/abs/2011.02288
- [5] Jowei Chen and Jonathan Rodden. 2013. Unintentional Gerrymandering: Political Geography and Electoral Bias in Legislatures. *Quarterly Journal of Political Science* 8, 3 (2013), 239–269.
- [6] Daryl DeFord, Moon Duchin, and Justin Solomon. 2021. Recombination: A Family of Markov Chains for Redistricting. Harvard Data Science Review 3, 1 (31 3 2021). https://doi.org/10.1162/99608f92.eb30390f https://hdsr.mitpress.mit.edu/pub/1ds8ptxu.

- [7] Moon Duchin, Taissa Gladkova, Eugene Henninger-Voss, Ben Klingensmith, Heather Newman, and Hannah Wheelen. 2019. Locating the Representational Baseline: Republicans in Massachusetts. *Election Law Journal* 18, 4 (12 2019), 388–401.
- [8] Moon Duchin and Douglas M. Spencer. 2021. Models, Race, and the Law. Yale Law Journal Forum 130 (March 2021), 744–797. https://www.yalelawjournal.org/ forum/models-race-and-the-law
- [9] Bernard Grofman. 1982. For Single-Member Districts, Random is Not Equal. In Representation and Redistricting Issues, Bernard Grofman, Arend Lijphart, Roberts McKay, and Howard Scarrow (Eds.). 55–58.
- [10] Stuart S. Nagel. 1965. Simplified Bipartisan Computer Redistricting. Stanford Law Review 17, 5 (1965), 863–899. http://www.jstor.org/stable/1226994
- [11] Jeffrey M. Shaman. 2102. Rules of General Applicability. First Amendment Law Review 10, 3 (2102), 419-464.
- [12] Alma Steingart. 2022. Law, computing and redistricting in the 1960s. In *Political Geometry*, Moon Duchin and Olivia Walch (Eds.). Birkhäuser Books, Chapter 8, 163–177. http://mggg.org/gerrybook
- [13] William Vickrey. 1961. On the Prevention of Gerrymandering. Political Science Quarterly 76, 1 (1961), 105–110. http://www.jstor.org/stable/2145973
- [14] James B. Weaver and Sidney W. Hess. 1963. A Procedure for Nonpartisan Districting: Development of Computer Techniques. *The Yale Law Journal* 73, 2 (1963), 288–308. http://www.jstor.org/stable/794769