Gauging Gerrymandering in Pennsylvania

A Monte Carlo Approach Using Methods from Spatial Statistics

JAMES RUSSELL
Muhlenberg College

BENJAMIN LIEBERMAN
University of South Carolina

There is currently no widely accepted standard method to determine whether gerrymandering has occurred. To determine a cutoff for unreasonable gerrymandering, simulating collections of districting plans in the absence of partisan bias has been proposed. In simulation-based methods, real-world election outcomes are compared to results from simulated districting plans. Here, a simulation method that creates possible districts in continuous space is proposed. Existing methods use preliminary spatial discretization of the state to perform simulations. This spatial discretization can result in biased estimates, which could lead to inaccurate conclusions regarding gerrymandering. We use our continuous-space method to analyze the political districts in Pennsylvania. All of our simulated elections result in fewer than 13 Republican seats, indicating that the districting plan used in Pennsylvania prior to 2018 was likely gerrymandered. This finding agrees with and confirms the results of simulation-based discrete-space gerrymandering studies without the presence of discretization bias.

Introduction

The term "gerrymandering" was coined in 1812, named after Massachusetts Governor Elbridge Gerry. Gerry, using an unconventional method of redrawing congressional district lines, concocted a plan so that his party, which was

COMMONWEALTH, Volume 22, Issue 1 (2023).
ISSN 2469-7672 (online). https://doi.org/10.15367/com.v22i1.640.

a statewide minority, held 29 out of the 40 congressional seats. Due to the strange shape of the district plan, which resembled that of a salamander, his opponents and the media dubbed the district plan "the gerrymander." Simply put, gerrymandering is the process of intentionally drawing district boundaries to grant a political advantage. The question considered here is: how much gerrymandering is considered "too much"? First, we briefly discuss the legal grounding behind gerrymandering in the United States.

Davis v. Bandemer (1986) was a case brought before the U.S. Supreme Court to question partisan gerrymandering in Indiana. The court could not define a standard in the case, which gave way to the later case of *Vieth v. Jubelirer* (2004). That being said, *Davis v. Bandemer* did establish the legal precedent that racial gerrymandering would be unconstitutional under the Equal Protection Clause as stated by the court: "Racial gerrymandering claims are justiciable because of the greater warrant the Equal Protection Clause gives the federal courts to intervene for protection against racial discrimination."

Vieth v. Jubelirer (2004) serves as additional background to the legality of gerrymandering in the United States. The case was filed in federal district court in Pennsylvania on the grounds that the U.S. House districts granted unfair political advantage to the Republican Party by giving them control of a majority of the congressional seats with a minority of registered voters. Vieth's lawyers argued that the map was a violation of equal protection under the 14th amendment and violated Article I of the Constitution. The United States Supreme Court ruled against Vieth, noting that there was no objective method for ruling on matters of gerrymandering or for determining if gerrymandering was significant. Their reasoning was that there is no existing standard for "adjudicating political gerrymandering claims." However, the court did note that in the future, given more evidence—such as some quantitative metric—they would be willing to rule on matters of partisan gerrymandering to determine if a districting plan had indeed crossed the proverbial line.

Currently, redrawing congressional district lines to grant one's own party a political advantage is generally acceptable, as seen in *Vieth v. Jubelirer* and *Davis v. Bandemer*. However, along with the claim that partisan gerrymandering is acceptable come some caveats. A district plan may be rejected if it displaces or separates "special interest groups," preventing them from having a voice in the political process. Special interest groups are simply groups of underrepresented people, such as minority groups, who would be disproportionately harmed if broken up in districting plans. *Gill v. Whitford* (2017) was a case born out of the 2011 redistricting plan in Wisconsin. The plan was created under the direction of the Republican Party and thus appeared to be drawn to give Republican voters a disproportionate amount of representation.

In *League of Women Voters v. Commonwealth* (2018), the Pennsylvania Supreme Court decided that the existing Pennsylvania districting plan was unconstitutional under the state's constitutional guarantee of "free and equal" elections. The ruling called for "a process assuring that a redistricting plan would be in place for the 2018 elections." Following the ruling, the Pennsylvania Supreme Court implemented a new districting plan to replace the previous Republican plan that the court had found to be gerrymandered. This new districting plan was used during the 2018 election, resulting in 9 Republican seats (compared to 13 Republican seats in the previous election).

Since the *League of Women Voters* decision, the most notable gerrymandering case has been *Rucho v. Common Cause* (2019). On June 27, 2019, the U.S. Supreme Court rendered a long-awaited decision which settled the constitutionality of gerrymandering as it pertains to the role of the U.S. Supreme Court. In a 5–4 decision, Chief Justice Roberts stated that the court "conclude[d] that partisan gerrymandering claims present political questions beyond the reach of the federal courts." However, this opinion was not shared by all members of the court. In a lengthy dissent, Justice Kagan asserted that democracy will indeed suffer as a result of the court's decision. She continued by saying, "the practices challenged in these cases imperil our system of government. . . . None is more important than free and fair elections." Kagan's dissent underscores the importance and clear ramifications of extreme gerrymandering. Gerrymandering has the ability to shape election outcomes but, more importantly, allows the party in control to abuse its power and maintain control without fear of losing reelection. The court's decision, however, boiled down to one major argument: is it constitutional for courts to weigh in on issues surrounding gerrymandering? Ultimately, the court decided that it was not. The Chief Justice wrote later in his opinion that "there are no legal standards discernible in the Constitution for making such judgments. . . . Federal judges have no license to reallocate political power between the two major political parties . . . [with] no legal standards to limit and direct their decisions."

The decision by the court has grave consequences for lawmakers and voters around the country. Gerrymandering is a matter that can no longer be handled in the Supreme Court. As such, it is up to the populace to ensure that their representatives are engaging in fair election practices. This is particularly relevant to Pennsylvania voters since the 2020 census triggered a new districting plan in Pennsylvania in 2021. By using tools to determine the degree of partisanship of redistricting lines (such as the algorithms described below), voters will, with a degree of statistical certainty, understand the degree to which their representatives have manipulated the borders for their political gain.

Mathematical and Statistical Methods for Measuring Gerrymandering

Gerrymandering can be extremely difficult to measure due to complex geographic and cultural features. One of the primary difficulties in regulating gerrymandering is the inability to assess whether gerrymandering has actually occurred. In previous cases, judges have noted the lack of any manageable standard of measurement for gerrymandering (Wang 2016). In an attempt to overcome this difficulty, many methods for testing for and measuring gerrymandering have been proposed. Examples of some of these methods include mathematical measures (Chambers and Miller 2010; Dube and Clark 2016; Fan et al. 2015; Hodge, Marshall, and Patterson 2010; Powell, Dube, and Clark 2017), statistical tests (Wang 2016), and simulation-based techniques (Bangia et al. 2017; Chen and Cottrell 2016; Duchin 2018; Magelby and Mosesson 2018; Powell, Clark, and Dube 2015). This section provides a brief summary of several mathematical and statistical methods for measuring gerrymandering.

One common mathematical measure used to assess gerrymandering is the compactness of a shape. In informal terms, a shape is compact if the average distance between spatial locations within the shape is minimized. Thus, in a mathematical sense, a circle is the most compact shape. Irregular and elongated shapes tend to be less compact. Therefore, irregular districts (which may be the result of gerrymandering) tend to be less compact than more regularly shaped districts.

Although there is no one recognized measure for the compactness of a shape, various compactness measures can be used to explore district patterns. For example, Fan et al. (2015) used both shape-based compactness measures and inertia-based compactness measures to analyze changes in compactness in California and North Carolina. Dube and Clark (2016) used graph compactness to assess the level of gerrymandering, which allows for weighting to adjust for population size of different areas. The inertia-based methods of Fan et al. (2015) were adjusted to account for spatial population distributions. Powell, Clark, and Dube (2017) used an alternative measure that considers the length of the district borders to measure the compactness of political districting plans. Simulation methods were then used to compare the measured compactness of actual plans to simulated district plans. Further details on geographic measures of compactness can be found in MacEachren (1985).

A similar mathematical measure that can be used to measure the "irregularity" of a shape is convexity. A shape is considered convex if, for any two points located within the shape, the shortest path between those two points lies entirely within the shape. In geographic terms for a political district, a district is convex if one can travel in a straight line between any two points in that district without leaving the district. Strict convexity, however, is a difficult measure to use because it is binary. Either a districting plan is convex or it is not. Chambers and Miller (2010) presented a more flexible measure of convexity that uses the probability that the line between any two points within a district lies entirely within that district. The resulting value, the convexity coefficient, is a number between zero and one that indicates how convex a shape is. Hodge, Marshall, and Patterson (2010) simplified and adjusted this measure and applied the method to U.S. congressional districts. To facilitate computation, Monte Carlo simulation was used to estimate the convexity coefficient.

Various tests based on traditional statistical methods have also been proposed. For example, three different statistical tests for gerrymandering were presented in Wang (2016). The first test involved a simulation-based measure to estimate a reasonable national districting norm. The election results in a particular state were compared to the distribution of results drawn from random samples of national district outcomes in an election. If the state's results were too extreme compared to the simulated distribution of outcomes, the state was deemed to be gerrymandered. Wang's (2016) second test used a t-test to compare the winning vote shares for the districts won by each of the parties. If the difference in mean votes for the two parties was statistically significant, this may have indicated partisan gerrymandering, which it did in Wisconsin and Maryland in Wang's test. For a third test, the reliable wins test, either the mean-median test for vote share or a comparison of the standard deviations of the vote shares for the two parties in the respective winning districts were compared.

The efficiency gap proposed by Stephanopoulos and McGhee (2015) is a relatively simple measure that can also be useful in determining whether a district plan is gerrymandered. The efficiency gap is obtained by the difference of the party's wasted votes divided by the total votes cast in the election. Stephanopoulos and McGhee (2015) define wasted votes as those "that don't contribute to victory for candidates, and they come in two forms: lost votes cast for candidates who are defeated, and surplus votes cast for winning candidates but in excess of what they needed to prevail." The efficiency gap has several advantages compared to others that have been proposed (Stephanopoulos and McGhee 2018). However, several flaws have been noted. For example, the efficiency gap threshold for "too extreme" a gap is arbitrary and may yield different results in geographically diverse states (Chambers, Miller, and Sobel 2017). In addition, political parties can still successfully gerrymander a state's political districts without violating the efficiency gap criteria (Cover 2017). Using some clever, rather simple mathematics, it can be seen that a completely non-gerrymandered state can actually fail the efficiency gap test. If voters for a political party are densely packed and clustered into urban areas, that party might waste more votes in this state simply due to the spatial distribution of their voters. This would result in a large efficiency gap even if no gerrymandering had taken place.

The measures described above provide useful information on gerrymandering, but often they do not account for the relationship between spatial location and the registered political affiliation of voters. A state that has not been gerrymandered may be deemed as gerrymandered solely based on the spatial distribution of a political party's registered voters. Simulation-based methods provide a tool for automatically accounting for the spatial dependence of voter's political affiliation. Several methods presented above involve Monte Carlo simulation for estimation and comparison of measures (Chambers and Miller 2010; Hodge, Marshall, and Patterson 2010; Powell, Clark, and Dube 2017; Wang 2016). Further, various methods entirely based on simulation have been proposed. The "simulation-based" concept involves generating random district plans from a hypothetical collection of all possible districting plans. This method is complicated by the difficulty in defining which districting plans should be included in this hypothetical collection. This method, however, is very useful because it automatically incorporates the specific features of the particular state of interest (Duchin 2018).

Various methods have been proposed to generate random districting plans in the absence of gerrymandering. These random nonpartisan plans are then used as a benchmark to assess gerrymandering in various states. Powell, Clark, and Dube (2015) constructed random districting plans using census tracks as the building blocks. The randomly generated districts were combined with demographic information on a census-block level to predict election outcomes. Chen and Cottrell (2016) used a similar method based on a grid of "similarly populated polygons" within any particular state. Linear regression models in real-world districting plans were compared to those with gerrymandered elections to assess the level of gerrymandering for each state. Bangia et al. (2017) also generated random districts in the absence of partisan gerrymandering using discrete spatial regions in the state of North Carolina. In their method, Voter Tabulation Districts (VTDs) were randomly grouped into 13 congressional districts. The vote counts in each of the VTDs were used to simulate elections in the absence of gerrymandering and compared to the real-world election results. They found that the number of seats in 2012 and 2016 were not typical election outcomes in non-gerrymandered redistrictings. In addition, Magelby and Mosesson (2018) described a computationally efficient approach to simulating random districting plans using graph-partitioning algorithms. They tested their algorithm on Mississippi, Virginia, and Texas and found all three states were likely gerrymandered.

Although simulation methods show promise for measuring gerrymandering across diverse geographic and demographic regions, one drawback of existing methods is the use of discrete spatial blocks. Whenever continuous space is split into discrete blocks, bias can be introduced. Changing the size and shape of the discrete blocks can have a significant impact on the result of the analysis (Aster, Borchers, and Thurber 2012). In addition, selecting the discrete block structure can be a challenging problem (Kotsiantis and Kanellopoulos 2006). In the case of the existing simulation methods described above, the discrete blocks are selected based on the available data. This can make it difficult to assess the impact or possible bias of the preliminary spatial divisions that are used. When analyzing districting plans, this bias could result in incorrect classifications for gerrymandered and non-gerrymandered states and inconsistent results across different methods of discretization. The Center for Range Voting (Smith and Kok 2005) describes one method of partitioning space, the shortest-splitline algorithm (Smith and Ryan 2007), which produces one particular redistricting plan with the goal of avoiding partisan gerrymandering. We propose a novel method for measuring gerrymandering by constructing random districting plans in continuous space. Although these simulations in continuous space are generally more complex, they provide an unbiased tool for assessing gerrymandering.

Describing the Data Set

In Pennsylvania, voter registration information is available through the Commonwealth's Department of State (https://www.pavoterservices.pa.gov/). The data include sex, birthdate, address, political party, and the date last voted. To perform our analysis, the spatial location of each individual, rather than the address, is needed. Geocoding is used to convert addresses to coordinates in latitude and longitude. We used "geocodio" (https://geocod.io/) to convert our data.

One difficulty in obtaining the spatial coordinates was the large size of the initial data set, which consists of every registered voter in Pennsylvania. To minimize set-up cost and computational burden for the probabilistic

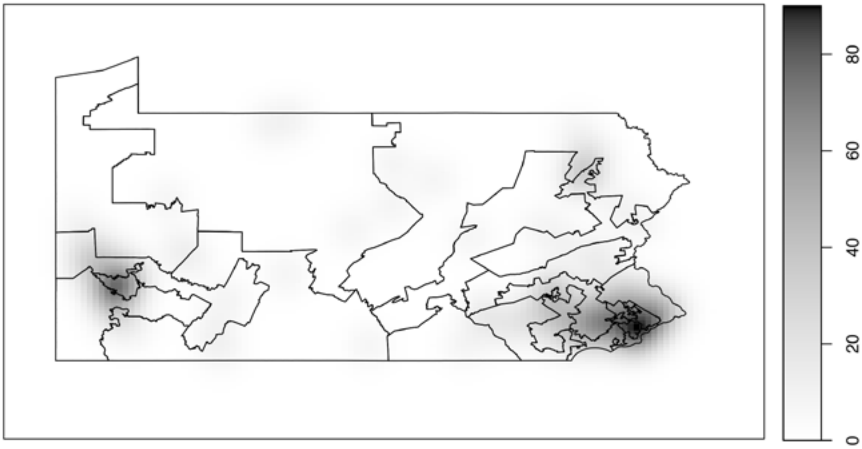

Figure 1. Pennsylvania Population Density. (*Source: Compiled by the authors using data from the Pennsylvania Department of State.*)

algorithm, a random sample of 10,400 registered voters was selected. Of the 10,400 registered voters, 47.3% were registered with the Democratic Party, 38.7% were registered with the Republican Party, and the remaining 13.9% were not registered with either of the two major political parties. To gain a better understanding of the data, a kernel density estimate of voter locations was plotted over the area of Pennsylvania. The kernel density estimate was calculated using the base stats package in R (Ihaka and Gentleman 1996). The plot for voter density is given in Figure 1. In the plot, the dark regions indicate high population density, and the light regions indicate low population density. The plot reveals the uneven population distribution across the state, as two cities with high population density, Philadelphia and Pittsburgh, are apparent. This illustrates that our sample reflects the population distribution of registered voters in Pennsylvania.

Method Overview

This section provides an overview of the algorithm we use to simulate districting plans in continuous space. The simulated districting plans are used to conduct a randomization test for partisan gerrymandering in Pennsylvania. In general, a randomization test uses a random simulation to assess a claim. Several simulated outcomes are generated, assuming some property is true. The outcomes from simulated trials are compared to actual data. If the actual data is significantly different from the simulated outcomes, the original property that was assumed must be incorrect. For example, I may want to test the claim that a coin is even (the probability of heads is equal to 0.5). To test this claim I could flip the coin 10 times. To generate the simulated outcomes in this example, I would flip a fair coin 10 times (with probability of heads equal to 0.5), record the number of heads out of the 10 flips, and repeat many times. If the outcome from our original coin flip (e.g., 9 heads) is significantly different than the simulated outcomes using the fair coin (e.g., a range of anywhere from 2 to 8 heads), I can reject the assumption that the original coin is fair.

Our goal is to simulate elections in the absence of gerrymandering. To do so, we split the state into 18 congressional districts using a random algorithm that does not consider political party. This is done using statistical point processes and spatial tessellations (see Appendixes 1 and 2 for additional details). The algorithm is computationally expensive due to the fact that an acceptance-rejection probabilistic algorithm is used to generate simulated districts that account for the congressional districting criteria of equal population and contiguity (see Appendixes 3 and 4 for additional details on generating districting plans that meet the requirements). Once the 18 random congressional districts are specified, logistic regression is used to predict voter turnout in a simulated election. The political affiliations and results of recent elections are used to generate the vote totals for the two major political parties in the simulated election. The results of the simulated election can then be compared to the observed results in a recent election.

Generating a Random Districting Plan in the Absence of Gerrymandering

To generate a single realization of a potential districting plan in the absence of gerrymandering, we use four steps:

1. *Randomly generate points across the entire state. These points will serve as "bases" to which all of the land in the state can be assigned.* To generate random points, an inhomogeneous Poisson point process is used. An inhomogeneous Poisson point process is a collection of random dots in space where the dots are not necessarily evenly distributed. This allows us to account for the varying population density across Pennsylvania. More details on an inhomogeneous point process are in Appendix 1.
2. *Assign every point in space in the region (all of the land in the state) to one of the "bases" generated in step 1.* This is done using a Voronoi tessellation on the point locations from step 1 to cut up the state into "tiles." A Voronoi tessellation is a mathematical tool to

divide a plane (in this case the land in Pennsylvania) into smaller tiles. See Appendix 2 for more details on Voronoi tessellation.

3. *Once we have a large collection of tiles, each tile is assigned to one of the 18 districts.* This is done using a random number generator and an assignment algorithm. More details on assigning the tiles to the 18 districts are in Appendix 3.
4. *Randomly shuffle individual tiles between adjacent districts until the specified districting criteria are met.* An accept-reject algorithm is used to control the shuffling of the tiles. More details on the accept-reject algorithm are in Appendix 4.

Steps 1 and 2 in this method consist of a novel technique for randomly splitting up continuous space in the state rather than using predefined discrete spatial regions such as VTDs or census blocks. Steps 3 and 4 are similar in concept to the method used by Magelby and Mosessson (2018) which used existing census blocks. In our method, since the tiles that form the basis for the random districting plan are randomly generated in steps 1 and 2, any potential bias introduced by preliminary discretization using preexisting blocks is avoided.

Simulating an Election

The random districting plan in the absence of gerrymandering is then used to simulate an election. The first step in simulating an election is determining which of the registered voters will vote. A logistic regression model is used to predict voter behavior. The data set contains the election in which each voter most recently voted. This variable is converted to either 0 or 1 depending on whether or not each particular individual voted in the 2016 Pennsylvania general election. Each registered voter's gender, political party, age, and county are used to predict voter turnout. One weakness of this method is that it assumes behavior is the same across elections. Further work could improve this model, but selecting alternative voter turnout prediction models had a minimal impact on the simulation results.

Finally, after determining who will vote in the simulated election, we calculate which political party will win in each of the 18 simulated districts. To do so, we determine who each voter will vote for using political affiliation. Voters registered as Republicans and Democrats are assigned to vote for their registered political party. Although voters may not always vote for candidates of the party in which they are registered, this is the only feasible option with the currently available data, since information including spatial location is inherently non-anonymous and thus individual voters' behavior in any given election is not available. Future work and data collection may allow for methods to predict individual voter decisions using a more accurate methodology. To determine the votes of individuals registered without a Democratic or Republican affiliation, county-level voting results in the previous election are used. For example, if a voter registered as Independent is from Adams County, the proportion of votes for Republicans in that county during the most recent election is used as the probability that the particular individual will vote Republican. Next, the vote counts in each of the 18 districts are calculated and the winning political party in each simulated district is determined.

Results

In the 1,000 simulations, the total number of seats won by Republican candidates ranges from 6 to 12, with a mean of approximately 9 Republican districts. The percentage of simulated elections with each respective number of Republican winners is presented in Table 1. In *League of Women Voters v. Commonwealth* (2018), the Pennsylvania Supreme Court ruled that the state had been gerrymandered. The prior election using the old district plan resulted in 13 Republican seats. All of our simulated elections generated in the absence of gerrymandering result in fewer than 13 Republican seats, indicating that the pre-2018 Pennsylvania districting plan was significantly different than any of the simulated plans. In 2018, with a new court-ordered district map, Republican candidates won 9 districts. This matches the mean number of districts won by Republicans in our simulations. Thus, our simulation rejects the pre-2018 map as gerrymandered, but does not reject the 2018 redrawn map as a gerrymandered plan.

Table 1. Results from Simulated Election Results for 1,000 Non-Gerrymandered Pennsylvania Districting Plans

Number of Districts Won	Simulated Occurrences
6	1.2%
7	9.6%
8	32.7%
9	32.5%
10	18.8%
11	4.7%
12	0.5%

Source: Created by the authors.

These results align well with the generally accepted view that the congressional districts were intentionally drawn to benefit Republican voters prior to 2018. The results are fairly similar to the results of Chen (Lapowski 2018), which found that in 500 simulations, 7, 8, 9, and 10 seats were won by the Republican candidate 6.4%, 36.2%, 55.5%, and 2% of the time, respectively. Chen's simulations involved discrete spatial blocks whereas our results involved continuous space. Our more flexible districting method allowed for more variability in the number of seats, but the simulation results still indicate that the number of seats held by Republican candidates prior to the 2018 redistricting (13) is not in line with simulation results. The simulation can easily be updated and universally extended to additional states.

Discussion

We examine the prevalence of gerrymandering in Pennsylvania through Monte Carlo simulation. To generate the simulations, the area of Pennsylvania is divided using Voronoi tessellation with the center of each tile sampled from an inhomogeneous Poisson point process. To ensure that each of the districts constitutes an approximately equal share of the population, a large number of tiles are generated and assigned to a particular district using an acceptance-rejection sampling algorithm. The vote of each individual is simulated using voter registration information, and the results from the simulated election are compared to the current distribution of congressional seats. In every simulated election, the number of seats won by the Republican candidates was less than the number of seats held by the Republican Party prior to the 2018 redistricting.

The simulation method presented here provides an extremely flexible tool for analyzing gerrymandering. It allows for the creation of random political districts in the absence of political gerrymandering. The algorithm can be adjusted for different districting criteria and can be updated to analyze any state where voter registration data is publicly available. One of the primary novel features of our method is the use of continuous space. Previous simulation methods for assessing gerrymandering assume the state has been discretized into small blocks. Whenever a continuous space is binned into blocks before analysis, bias can be introduced. In addition, changing the size and structure of the bins that are used can lead to inconsistent results. Using our method, the potential bias from discretization is avoided and a wider variety of possible districting plans is possible. However, even with the additional variability in potential districting plans, the simulated elections in Pennsylvania still resulted in fewer than 13 seats won by the Republican Party in every case.

Working with continuous space has several drawbacks. Often continuous space is more difficult to manage, resulting in additional computational and mathematical complexities. When running thousands of simulations, the computation time becomes a limiting factor for the scalability of the method. Additionally, the lines of the generated districts do not follow any existing townships, census blocks, or geographic features within the state. Therefore, it is not recommended to use our continuous space algorithm for generating a districting plan that would be used in a real-world election. Our method does not produce an "optimal" plan but instead creates a large collection of reasonable plans. Thus, the results from the simulations should be used for comparison purposes only. Although our method is computationally expensive and more difficult to implement in practice, the results are similar to existing versions in discrete space that are more computationally feasible such as that of Chen and Cottrell (2016). Thus, a primary contribution of our findings is to support existing methods of generating random redistricting plans that use preexisting discrete spatial blocks that are generally more practical to implement, have more readily available voting data, and are more computationally feasible to simulate realizations from than continuous space methods.

Using voter registration data also has several drawbacks. Often voter registration databases are out of date. In addition, the need for anonymity limits the ability to use important covariates in modeling election results. Simulating election results, in general, can be a very difficult problem. Each election is different and inherently draws a different subset of voters. Since each election occurs in particular circumstances and with particular candidates, using data from a previous election to predict voter turnout may yield different results. In the future, we hope to improve the regression model for voter turnout. Several characteristics available in other data sets, such as the American Community Survey (ACS), could potentially improve the performance of a simulated election. In future work, the additional information available in the ACS data set could be combined with the exact spatial location available in Pennsylvania voter registration data to yield a more powerful tool for predicting election results. Combining ACS data and results from previous elections with voter registration data may allow us to simulate elections by predicting voter decisions (who they will vote for) rather than relying on a voter's registered party and predicting only voter turnout. Further, the model for voter turnout can be improved to incorporate spatial dependence in voter locations. In this case, our simulation results were robust across different models for predicting voter turnout, so basic logistic regression was used for simplicity. In the future, methods from spatial statistics, such as spatial logistic regression and spatial autoregressive models (Bivand, Pebesma, and Gómez-Rubio 2008), can be used to analyze voter turnout data for various elections, thus enabling improved predictions of voter behavior.

Various options exist for assessing the simulated districting plans and election results once they have been generated. The exact measure for comparison can be selected on a case-by-case basis, but several options include: number of seats won, compactness (Fan et al. 2015), graph compactness (Dube and Clark 2016), convexity (Hodge, Marshall, and Patterson 2010; Miller 2007; Powell, Clark, and Dube 2017), and the efficiency gap (Stephanopoulos and McGhee 2015). In the future, it would be useful to calculate several of these measures in Pennsylvania to provide a mathematically rigorous baseline value for when a districting plan is extreme enough to conclude that gerrymandering has occurred.

Appendix 1. Inhomogeneous Poisson Point Process

A spatial point process consists of a region of interest and spatial locations within that region of interest. The spatial location of points is considered the random component. Typically, the number of points is considered random as well. One example of a spatial point process is the locations of trees in a forest. The region of interest is the area within the borders of the forest and the random points are spatial coordinates of the locations of the trees. One important aspect of a spatial point process is that the location of the point is random; if the location of the point is fixed, a different branch of spatial statistics would be necessary.

The Poisson point process (Miles 1970) is one of the most basic examples of a spatial point process. In a Poisson point process, the number of points observed is random and follow a Poisson distribution. The Poisson distribution is a common discrete distribution used in statistics for count data. Conditional on the number of points, in a basic (homogenous) Poisson point process the points are randomly spread uniformly throughout the region of interest. This means that each point has an equal probability of being found in any subsection of the space of equal size. In addition, for a basic Poisson point process the location of each point is independent of all other points.

The population density varies for different areas of Pennsylvania. Therefore, a homogenous point process would not be an accurate tool for representing voter distribution. Instead, we use the inhomogeneous Poisson point process (Bivand, Pebesma, and Gómez-Rubio 2008). The inhomogeneous Poisson point process incorporates an additional feature called an intensity surface. The intensity surface gives the relative probability of finding random points in a particular location. Subregions that have a higher intensity are more likely to have a higher number of points. For example, a region with a higher population density within a state, such as a large city, can correspond to a spatial location with a higher relative intensity surface than a region with a lower population density, such as a rural town. Thus, we use an intensity surface with larger values in urban areas (such as Philadelphia and Pittsburgh) and lower values in rural areas. The intensity surface of Pennsylvania can be estimated using the density of voter distributions presented in Figure 1.

After estimating the intensity surface for voter locations, we can generate an inhomogeneous Poisson point process. A random point process with 2,000 points is generated using the Spatstat package (Baddeley and Turner 2005) in R. The number of points can be adjusted for computational efficiency. Different numbers of points were considered and the resulting random districts were similar as long as the number of points was sufficiently large. If too few points were selected, the algorithm did not converge, and as the number of points was increased the computation time for the remainder of the algorithm increased significantly.

Appendix 2. Voronoi Tessellation

A Poisson point process can be combined with Voronoi tessellation to generate a random collection of shapes or tiles that span the original region of interest (Hinde and Miles 1980). Once the random points in the state have been selected, a Voronoi tessellation is used to split the area of Pennsylvania into tiles. Voronoi tessellations have been proposed to find an "optimal" districting plan (Miller 2007; Ricca, Scozzari, and Simeone 2008) but, to our knowledge, have not been combined with random simulations to measure gerrymandering. For a collection of *N* points, a Voronoi tessellation assigns each location in space to the nearest point using Euclidean distance. The result is a collection of *N* tiles that, when combined, make up the entire original area. An example of a Voronoi tessellation is given in Figure 2.

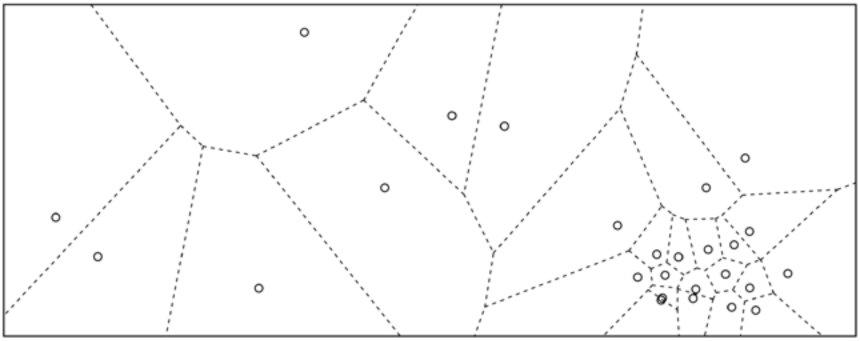

Figure 2. Sample Voronoi Tessellation. (*Source: Authors.*)

Once the tiles have been selected, each voter in our random sample is assigned to a tile based on their spatial location. The total number of voters in each tile is then calculated. Tiles with fewer than 6 residents (from our sample of 10,400) are discarded by removing the associated spatial point and recomputing the tessellation. Other minimum resident-size cutoffs are considered and yield similar results with significantly different computation times. A large number of points are typically disregarded in this manner, making it easier to assign tiles to districts. The total number of points in the updated point process ranges from around 500 to 570. The mean number of tiles is 536.9 and the standard deviation is 11.8.

Appendix 3. Assigning the Tiles to Districts

Once the original Voronoi tessellation has been finalized, the component tiles must be assigned to the 18 districts. One tile is randomly selected to belong to each of the districts as an original or "seed" tile. Once the 18 initial "seed" tiles are selected, each district is iterated through one by one. In each iteration of the loop, all unassigned tangent tiles to the district are input into a list and one tile is randomly sampled. The sampled tile is added to the district. If there are no tangent tiles that have not been assigned to a district, the current district is passed over. The algorithm continues to loop through the districts until all of the tiles have been assigned to one of the 18 districts.

Appendix 4. Acceptance-Rejection Algorithm for Tile Shifting

After each of the tiles has been assigned to a district, all of the districts are contiguous, but the population in each simulated district can vary substantially. The districts therefore need to be adjusted to meet the equal population criterion for congressional districts. The districts are adjusted using a random accept-reject algorithm. Accept-reject algorithms work in two steps. In the first step, a shift or adjustment is proposed. Afterward, the proposed change is either accepted or rejected with a probability that depends on the desired criterion. In this case the acceptance probability is determined by the congressional districting criteria. The difference in the population size of the relevant districts is calculated before and after the proposed shift. The calculated populations for the respective districts are used to calculate the probability that the proposed tile shift will be accepted. A random number generator, available in R (Ihaka and Gentleman 1996), based on the calculated probability, is then used to determine if the shift is accepted. If the shift is accepted, the tile allocation is updated to reflect the shift. Otherwise, the shift is rejected and the tile allocation is not adjusted.

The following algorithm is rerun in a loop until the difference in the proportions of the population in the largest and smallest of the districts is within 1%, which is a smaller spread than the current congressional districting plan. The algorithm can be broken down into five main components.

1. Randomly sample 1 of the 18 districts.
2. Randomly sample 1 of the tiles on the border of that district.
3. Randomly sample 1 of the districts that the selected tile is adjacent to, and propose moving that tile to the new district.
4. Calculate the acceptance probability (*p*) for the proposed shift:
   a. Test whether the proposed shift would create any noncontiguous districts. (This is done via an *N* x *N* adjacency matrix.)
   b. If the tile creates noncontiguous districts, set the acceptance probability to 0.
   c. If the tile preserves district contiguity, calculate the acceptance probability based on the change in population difference in the districts before and after the proposed shift.
5. Accept the proposed shift with probability *p*.

After the algorithm is completed, the tiles have been shifted so that they are contiguous and the variance in population between the largest and smallest districts is less than 1%.

REFERENCES

Aster, Richard C., Brian Borchers, and Clifford H. Thurber. 2012. *Parameter Estimation and Inverse Problems*. Cambridge, MA: Academic Press.

Baddeley, Adrian, and Rolf Turner. 2005. "Spatstat: An R Package for Analyzing Spatial Point Patterns." *Journal of Statistical Software* 12 (6): 1–42.

Bangia, Sachet, Christy Vaughn Graves, Gregory Herschlag, Han Sung Kang, Justin Luo, Jonathan C. Mattingly, and Robert Ravier. 2017. "Redistricting: Drawing the Line." *arXiv* (May 8). https://arxiv.org/abs/1704.03360.

Bivand, Roger S., Edzer Pebesma, and Virgilio Gómez-Rubio. 2008. *Applied Spatial Data Analysis with R*. 8th ed. New York: Springer.

Chambers, Christopher P., and Alan D. Miller. 2010. "A Measure of Bizarreness." *Quarterly Journal of Political Science* 5 (1): 27–44.

Chambers, Christopher P., Alan D. Miller, and Joel Sobel. 2017. "Flaws in the Efficiency Gap." *Journal of Law and Politics* 3 (1): 1–33.

Chen, Jowei, and David Cottrell. 2016. "Evaluating Partisan Gains from Congressional Gerrymandering: Using Computer Simulations to Estimate the Effect of Gerrymandering in the U.S. House." *Electoral Studies* 44 (December): 329–340.

Cover, Benjamin Plener. 2017. "Quantifying Partisan Gerrymandering: An Evaluation of the Efficiency Gap Proposal." *Stanford Law Review* 70 (April): 1131–1233.

Davis v. Bandemer. 478 U.S. 109. 1986. https://tile.loc.gov/storage-services/service/ll/usrep/usrep478/usrep478109/usrep478109.pdf, pg 111.

Dube, Matthew P., and Jesse T. Clark. 2016. "Beyond the Circle: Measuring District Compactness Using Graph Theory." Paper presented at the Annual Meeting of the Northeastern Political Science Association.

Duchin, Moon. 2018. "Gerrymandering Metrics: How to Measure? What's the Baseline?" *arXiv* (January 6). https://arxiv.org/abs/1801.02064.

Fan, Chao, Wenwen Li, Levi J. Wolf, and Soe W. Myint. 2015. "A Spatiotemporal Compactness Pattern Analysis of Congressional Districts to Assess Partisan Gerrymandering: A Case Study with California and North Carolina." *Annals of the Association of American Geographers* 105 (4): 736–753.

Gill v. Whitford. No. 16-1161 (U.S. July 2017).

Hinde, A. L., and R. E. Miles. 1980. "Monte Carlo Estimates of the Distributions of the Random Polygons of the Voronoi Tessellation with Respect to a Poisson Process." *Journal of Statistical Computation and Simulation* 10 (3–4): 205–223.

Hodge, Jonathan K., Emily Marshall, and Geoff Patterson. 2010. "Gerrymandering and Convexity." *The College Mathematics Journal* 41 (4): 312–324.

Ihaka, Ross, and Robert Gentleman. 1996. "R: A Language for Data Analysis and Graphics." *Journal of Computational and Graphical Statistics* 5 (3): 299–314.

Kotsiantis, Sotiris, and Dimitris Kanellopoulos. 2006. "Discretization Techniques: A Recent Survey." *GESTS International Transactions on Computer Science and Engineering* 32 (1): 47–58.

Lapowski, Issie. 2018. "The Geeks Who Put a Stop to Pennsylvania's Partisan Gerrymandering." *Wired* (February 28). https://www.wired.com/story/pennsylvania-partisan-gerrymandering-experts/.

League of Women Voters v. Commonwealth. 2018. 178 A.3d 737.

MacEachren, Alan M. 1985. "Compactness of Geographic Shape: Comparison and Evaluation of Measures." *Geografiska Annaler: Series B, Human Geography* 67 (1): 53–67.

Magelby, Daniel B., and Daniel B. Mosesson. 2018. "A New Approach for Developing Neutral Redistricting Plans." *Political Analysis* 26 (2): 147–167.

Miles, Roger E. 1970. "On the Homogeneous Planar Poisson Point Process." *Mathematical Biosciences* 6 (1): 85–127.

Miller, Stacy. 2007. "The Problem of Redistricting: The Use of Centroidal Voronoi Diagrams to Build Unbiased Congressional Districts." Senior project, Whitman College. https://people.math.sc.edu/Burkardt/classes/urop_2016/miller.pdf.

Powell, Richard J., Jesse Tyler Clark, and Matthew P. Dube. 2015. "Determining an Expected House Majority Using Pattern Analysis." Paper presented at the Annual Meeting of the Northeast Political Science Association.

———. 2017. "Mathematical Characteristics of District Boundary Lines as Indicators of Partisan Gerrymandering in U.S. House Elections." Paper presented at the Annual Meeting of the Midwest Political Science Association.

Ricca, Federica, Andrea Scozzari, and Bruno Simeone. 2008. "Weighted Voronoi Region Algorithms for Political Districting." *Mathematical and Computer Modelling* 48 (9–10): 1468–1477.

Rucho v. Common Cause. 2019. No. 18-422, 588 U.S. ___.

Smith, Warren D., and Jan Kok. 2005. "Splitline Districtings of All 50 states + DC + PR." https://www.rangevoting.org/SplitLR.html.

Smith, Warren D., and Ivan Ryan. 2007. "Gerrymandering and a Cure—Shortest Splitline Algorithm." https://rangevoting.org/GerryExamples.html.

Stephanopoulos, Nicholas O., and Eric M. McGhee. 2015. "Partisan Gerrymandering and the Efficiency Gap." *The University of Chicago Law Review* 82 (2): 573–1071.

———. 2018. "The Measure of a Metric: The Debate Over Quantifying Partisan Gerrymandering." *Stanford Law Review* 70 (May): 1503–1568.

Vieth v. Jubelirer. 2004. 541 U.S. 267 (2004). https://supreme.justia.com/cases/federal/us/541/267/

Wang, Samuel S.-H. 2016. "Three Tests for Practical Evaluation of Partisan Gerrymandering." *Stanford Law Review* 68 (June): 1263–1321.

Warrington, Gregory S. 2018. "Quantifying Gerrymandering Using the Vote Distribution." *Election Law Journal: Rules, Politics, and Policy* 17 (1): 39–57.

James Russell is an associate professor at Muhlenberg College in the Department of Mathematics and Computer Science. He received his PhD in statistics in 2016 from Pennsylvania State University. He is interested in real-world applications of statistics, especially with data involving spatial information.

Benjamin Lieberman graduated from Muhlenberg College and received an M.S. in statistics from the University of South Carolina. He is currently a statistician at Freddie Mac in the Single Family Division. He is interested in applying statistics to real-world problems and communicating his findings to a broad audience.