Meddling Metrics: the Effects of Measuring and Constraining Partisan Gerrymandering on Voter Incentives

BRIAN BRUBACH*, University of Maryland, College Park
ARAVIND SRINIVASAN, University of Maryland, College Park
SHAWN ZHAO, Montgomery Blair High School

Gerrymandering is the process of drawing electoral district maps in order to manipulate the outcomes of elections. Partisan gerrymandering occurs when political parties use this practice to gain an advantage. Increasingly, computers are involved in both drawing biased, partisan districts and in attempts to measure and regulate this practice. Several of the most high-profile proposals to measure partisan gerrymandering involve the use of past voting data. Prior work primarily studies the ability of these metrics to detect gerrymandering. However, it does not account for how legislation based on the metrics could affect voter behavior or be circumvented via strategic voting. We show that even in a two-party election, using past voting data can affect strategyproofness. We further focus on the proposal to ban "outlier maps," which appear biased toward a particular party when compared to a random sampling of legal maps. We introduce a game which models the iterative sequence of voting and redrawing districts under the restriction that outlier maps are forbidden. Using this game, we illustrate strategies for a majority party to increase its seat count by voting strategically. This leads to a heuristic for gaming the system when outliers are banned, which we explore experimentally. Applying a version of our heuristic to past North Carolina voting data shows that these strategies can be found for real states under some stricter assumptions. Finally, we address some questions from the recent US Supreme Court case, Rucho v. Common Cause, that relate to our model.

CCS Concepts: • **Theory of computation → Algorithmic game theory and mechanism design**; • **Applied computing** → *Law, social and behavioral sciences*; *Voting / election technologies*.

Additional Key Words and Phrases: gerrymandering, redistricting, voting, social choice theory, elections, fairness

ACM Reference Format:
Brian Brubach, Aravind Srinivasan, and Shawn Zhao. 2020. Meddling Metrics: the Effects of Measuring and Constraining Partisan Gerrymandering on Voter Incentives. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC '20), July 13–17, 2020, Virtual Event, Hungary.* ACM, New York, NY, USA, 19 pages. https://doi.org/10.1145/3391403.3399529

1 INTRODUCTION

Computer algorithms and automated systems have become crucial players in the game of drawing and evaluating US electoral districts. Governments have increasingly utilized software to draw districts that influence the outcomes of elections (e.g., to favor a particular political party, incumbent politician, or racial group). Conversely, academics and enthusiastic citizens have proposed algorithms that purport to be fair alternatives [16, 23, 26], a familiar promise of technology that

*Authors listed in alphabetical order of last name.

EC '20, July 13–17, 2020, Virtual Event, Hungary

ACM ISBN 978-1-4503-7975-5/20/07…$15.00
https://doi.org/10.1145/3391403.3399529

does not always hold true even with the best of intentions. More recently, major US court cases have highlighted computational approaches to evaluate whether a district map is gerrymandered to favor a particular political party [13–15, 21]. These legal challenges include a Pennsylvania Supreme Court case which led to the redrawing of congressional districts in that state [3] and a recent landmark US Supreme court case [4].

At the heart of algorithms claiming to draw fair, unbiased districts or evaluate existing maps, we find a series of metrics and constraints that attempt to define what makes a district fair, unbiased, or even legal. Our goal in this work is to illustrate unexplored or under-explored aspects of how our choice of metrics and constraints can influence voter behavior and how regulations based on these metrics can be circumvented. In the context of evaluating maps for partisan gerrymandering, we show that using past voting data to evaluate maps can incentivize strategizing among voters even in two-party single-member district plurality systems.

In some sense, strategic voting in an election is not inherently bad. The seminal works of [19, 27] showed that elections with more than two candidates are not strategyproof. We see this play out regularly in US elections when a voter prefers a third party candidate, but chooses to vote for their favorite among the two major-party candidates. However, there are some clear negative effects. First, strategizing reduces the effectiveness of gerrymandering regulations if players can get around them even for single election cycles. Second, the ability to strategize may not be fairly distributed, disenfranchising voters who are less able to strategize. The work of [22] describes this disparate effect in terms of classifiers. In the specific case of redistricting and the strategic act of gerrymandering itself, a political party with a rural voter base can have more power to gerrymander in its favor than an opposing party with an urban base [9].

The ruling in the most recent US Supreme Court case on partisan gerrymandering has essentially left it to states to decide how they will address the issue, if at all [4]. In doing so, state governments will need to consider the downstream effects of how they choose to measure and regulate partisan gerrymandering. To that end, we initiate the study of how measuring gerrymandering using past voter behavior can incentivize strategic voting to circumvent regulations. We show that careful scrutiny should be given to any measurement which uses past voter behavior to evaluate and affect the choice of district maps.

2 PROBLEM DESCRIPTION AND DEFINITIONS

In this section, we outline many problems and definitions associated with districting and gerrymandering, including proposed methods to measure gerrymandering or draw fairer districts. We begin with some basic definitions, then introduce the growing role of computer science in combating gerrymandering. Finally, we summarize recent attempts to evaluate or draw districts. Throughout the paper, we restrict our discussion to the United States political system, where gerrymandering has become a highly contentious issue, as a guiding example.

2.1 Basic districting and gerrymandering definitions

In US politics, many representatives in both federal and state governments are elected via single-member district plurality systems where voters are partitioned into districts which each elect representatives with a plurality vote. A *plurality vote* awards a seat in government to the candidate who has received the most votes even if that candidate has not received a majority of the total votes cast. Election districts may be redrawn every 10 years following a population census. This ensures that districts respond to population shifts over time and remain reasonably balanced in terms of the number of voters per district to avoid vote dilution.

Historically, state governments have been charged with drawing new maps, and the party which holds a majority of seats in a state's government may control the outcome of this process. More recently, some states have transitioned to a system where an "independent" commission draws the maps. Such commissions have been controversial in terms of both how their independence can be guaranteed and whether they violate the US Constitution which assigns this duty to state governments. Nevertheless, the US Supreme Court has upheld the right of voters in a state to give redistricting authority to an independent redistricting commission [1]. The present work mainly explores models where the party controlling the majority of seats also controls the drawing of districts subject to some restrictions. However, our results may be useful in understanding how to regulate independent commissions as well.

The strategy of gerrymandering is to influence the outcome of an election through the process of drawing districts. This can be done by *packing* many voters from one group into a single district where they will cast more votes than needed to win, or *cracking* voters from that group into multiple districts where they will cast votes for losing candidates in each district. In both packing and cracking, the idea is to make the opposing group "waste" as many votes as possible while making one's own group waste fewer votes.

The type of gerrymandering depends on the type of groups being targeted and the intended outcome. *Partisan gerrymandering* attempts to favor one political party over another. Similarly, *racial gerrymandering* attempts to favor one racial group over another. *Incumbent gerrymandering* is slightly different in that it creates a bias toward re-electing an incumbent candidate. The primary focus of this paper is partisan gerrymandering. However, these different types are often entangled such as when there is a correlation between racial demographics and party membership or when partisan gerrymandering leads to the creation of "safe districts" where incumbents have an advantage. In fact, laws that prevent the cracking of racial groups have been used to justify packing members of a racial group for the apparent purpose of partisan gerrymandering as in Florida's famously snake-like District 5.

The space of legal district maps is restricted by a set of (sometimes competing) constraints and objectives that vary from state-to-state. The most common restrictions are contiguity, community integrity, population balance, hole-freeness, and compactness. *Contiguity* simply means that districts should be contiguous spaces although in the more extreme cases, they may only be connected by narrow paths. *Community integrity* refers to the objective that districts should avoid splitting defined communities (e.g. counties, towns, etc.) if possible. However, communities are routinely split ostensibly to meet other objectives. *Population balance* is the objective that the number of voters in each district should be as balanced as possible in order to give roughly the same weight to each person's vote. The degree to which population balance is violated can depend on other considerations such as community integrity, and congressional district population sizes within a state can vary by as much as 897,080 in Texas District 22 to 713,480 in Texas District 13 [10]. There are also several single-district states where the district size is determined solely by the state population and may be larger than the national average or smaller districts. Montana's at-large district has an estimated population of 1,050,493 compared to one of the smallest districts, Rhode Island's District 2 with 520,389 [10]. *Hole-freeness* states that no district should be completely surrounded by one other district. Finally, *compactness* is perhaps the least consistently defined goal with many possible definitions of compactness existing. These include k-median-like objectives minimizing the average distance a voter has to travel to a center point in their district and objectives minimizing ratio of area to perimeter (see Section 2.5 for a discussion of algorithms that use these objectives).

2.2 Using computer science to combat partisan gerrymandering

There are two main directions where computer science has become involved in combating partisan gerrymandering: verifying whether a given map is unfairly gerrymandered and drawing fair maps.

Crucial to both directions is the question of how to measure gerrymandering and define fairness in this context. This is clearly true for a verifier, and an algorithm for drawing districts must have some notion of where the line is between fair and unfair in order to avoid crossing it.

Unfortunately, it is not simple to measure whether partisan gerrymandering has occurred nor is it straightforward to say that partisan gerrymandering is "unfair" in the US legal context (e.g. unconstitutional). While gerrymandering may often seem obvious just by looking at a map, this so-called "eyeball test" is not robust [17]. It is possible that a strange looking map is actually subject to the geography of a state with respect to both natural and artificial features that perturb distances between voters such as mountains or highways. On the other hand, it is also possible for a reasonable looking map to be gerrymandered. Then, supposing we have some test to show that a map is gerrymandered, we must further show that the practice we measure violates the law in some way. In the US legal system, the mere practice of gerrymandering is not illegal even though the concept of candidates choosing their voters instead of the other way around may violate many citizen's sensibilities. For example, the US Supreme Court has ruled that racial gerrymandering is unconstitutional while incumbent gerrymandering is allowed [5, 6].

2.3 Approaches to measurement

Here, we define and discuss some of the most well-known approaches to measure partisan gerrymandering. These include the look test, outlier detection, proportionality, competitiveness, the efficiency gap, and compactness. Our work focuses primarily on outlier detection and proportionality, but is also relevant to any measurement approach using past voting data.

One seemingly obvious standard which may use past voting data is *proportionality*. This is the idea that the proportion of seats assigned to each party should be close to the proportion of votes received by each party. However, there are challenges to this standard as well. Drawing a map which achieves proportional representation is impossible for states like Massachusetts where voters of the two parties are too evenly distributed [18] or single-district states. More importantly in the US, courts have rejected proportionality tests. For the sake of analysis and comparison to other tests, we consider an achievable proportionality concept of maximally proportional maps. A *maximally proportional map* is one which comes as close as possible to allocating seats proportionally among all *known* maps (given the large search space, we may not know the true most proportional maps).

More central to our work is the recent study of outlier maps. An *outlier map* is one which is abnormal by some measure of representation. In this paper, we use the natural measure of seat count awarded to each party for a map based on past voting data. Thus, a map may be an outlier if it awards 3 out of 11 total seats to a given party while almost every other legal map awards 4 or 5 seats to that party.

A key challenge in detecting an outlier map is that it is difficult to determine what an average map is. The space of all possible maps is too large to check each one. This has led to approaches that aim to approximately randomly sample from the set of all possible legal maps using Markov Chain Monte Carlo techniques [13, 15, 21]. One can generate thousands of random maps and use past voting data to determine how many seats each party would win on each map assuming voters were to vote the same way they have in the past. We can compare a given map to this random sample to determine if it is an outlier. The seat count measure of representation we consider is used in [21]. Other works use more fine-grained measures that capture smaller changes in the distribution of voters to districts [13].

Additional measures which use past voter data include competitiveness and the efficiency gap. The *competitiveness* of a district attempts to capture the extent to which it might be won by either of two parties. The *efficiency gap* measures the difference in the number of wasted votes between two parties [28]. Generally speaking, any vote cast for a losing candidate is wasted, and for a winning candidate, a wasted vote is any vote beyond 50% (or alternatively the minimum needed to win in some variants). This method was used unsuccessfully in a US Supreme Court case challenging gerrymandering in the state of Wisconsin [2] and critiques of it can be found in [8, 25]. Nevertheless, we conjecture that an efficiency gap standard might incentivize participation and truthful voting in contrast to the outlier detection method discussed here.

Finally, we address measures that explicitly do not take past voting behaviour into account. These typically involve some mathematical definition of what a good cluster should look like. For example, a ratio of perimeter to area, average distance from all voters to single meeting point, or even the average number of neighboring districts. They often arise in the discussion of algorithms for drawing districts, and we save a closer examination for Section 2.5.

2.4 Redistricting subject to gerrymandering regulations

In this paper, we consider models where humans control the redistricting process (possibly using any algorithms they please), but they may be restricted by gerrymandering metrics. In particular, we focus on a model we call the *majority party draw model* where the political party currently in power in a state controls the redistricting process. Thus, we assume a biased, partisan agent is drawing the maps to favor one party over the other. Districts are drawn by the party holding the majority of seats in order to maximize that party's utility subject to any constraints. We constrain the drawing party with the typical restrictions that districts must be contiguous and balanced in terms of population. In addition, we include a regulation proposed to reign in partisan gerrymandering, *banning outlier maps*. This regulation prohibits the drawing of outlier maps with respect to seat counts awarded to each party based on the votes cast in the most recent previous election.

2.5 Other approaches to drawing districts

While this work primarily studies approaches to measuring gerrymandering and their potential effects, we briefly discuss the topic of drawing districts as one might argue, "Why not just use one of the existing tools for drawing fair districts?" We describe two broad lines of work devoted to drawing districts without partisan gerrymandering and sketch why these existing techniques, by themselves, may not be sufficient or desirable. We group the techniques by whether or not they use past voting data.

The first general class of algorithms embraces the natural idea of drawing districts without considering voting preferences at all. We describe some of these works and then argue why it is still important to test them for partisan bias.

Many of these proposed tools attempt to optimize some measure of compactness. A k-median-based algorithm is featured in visualizations of different districting schemes on the website FiveThirtyEight [11]. The similar objective of balanced centroidal power diagrams is used in [16]. There are also methods that focus on simple, achievable objectives like the shortest splitline algorithm which recursively finds the shortest line dividing the population in half until the desired number of equally sized districts is found [26].

Although tools in this category are often self-described as unbiased, testing their output for gerrymandering via unintentional bias (e.g., using measures from Section 2.3) is still needed. In the classic fairness example of classifiers, we know that sensitive features such as race can be redundantly encoded in other features even if race itself is omitted from the feature set. Similarly, party membership (or race) may be encoded in other features of population data that are used by these algorithms.

In the US, features such as party membership and race are correlated with urban versus rural residence. Thus, any clustering algorithm which is sensitive to the density of points being clustered could potentially be biased. For example, a bias based on population density could be present in algorithms which minimize average distance from voters to a central meeting point such as those based on the k-median objective or balanced centroidal power diagrams. This is not to diminish the value of these algorithms or claim that they are biased, but rather to show that they should be tested for biases. This evaluation requires a measure like those described in Section 2.3 and studied in our present work.

We note that this concern about hidden bias is not new. Other works have explored the effects of an urban/rural party split [9, 12] and the following statement expressed by Justice Scalia in [7] was discussed in [14].

> Consider, for example, a legislature that draws district lines with no objectives in mind except compactness and respect for the lines of political subdivisions. Under that system, political groups that tend to cluster (as is the case with Democratic voters in cities) would be systematically affected by what might be called a "natural" packing effect.

In the second category, are approaches that explicitly use voting data. Some of these use measures from Section 2.3 to guide the algorithm. An evolutionary algorithm called PEAR uses a combination of objectives including competitiveness based on past voting data and compactness measured as 4π times the area of a district divided by its perimeter squared [23]. There are also cake cutting approaches which allow two parties to divide up a state [24]. This is shown to meet a definition of fairness to the two major parties, but may not be fair to other groups such as third parties, geographic regions with shared interests (e.g., farming communities), or racial groups. Allowing the two major parties to collaborate on drawing a map may also lead to undesirable incumbent gerrymandering as in California in the past [20]. While our work does not directly address these tools, it implies that they could be susceptible to strategic voting.

2.6 Other related work

Here, we summarize some additional related work on social choice theory and fairness.

The classic Gibbard-Satterthwaite Theorem [19, 27] established that at least one of the following must hold in ordinal voting systems electing a single candidate.

(1) The system is a dictatorship; one voter chooses a winner.
(2) The are only two candidates.
(3) The system is not strategyproof and inspires tactical voting.

While our result shows strategizing in a two party election with single candidates elected per district, we note that in our model, a single voter's strategy can affect the outcomes of multiple elections. Even though each district elects a single candidate, we show how a vote in one district can affect all districts in future elections. Thus, our work highlights a situation where the Gibbard-Satterthwaite Theorem may appear to apply, but surprisingly does not.

In the context of fairness, [22] considered the disparity that results from different groups having different abilities to strategize. Our work shows how certain regulations can create additional opportunities for strategic manipulation in voting systems and opens the question of whether additional disparity results from those regulations. This is relevant in elections where political parties are scrutinized to determine if they are leveraging such disparities to disenfranchise groups of voters (e.g., using voter id laws to target communities living in urban areas where drivers' licenses are not ubiquitous or restrictions on early voting to target people with less flexible working hours).

3 OUR CONTRIBUTIONS

We illustrate how using past voting data to regulate the drawing of district maps can incentivize voters to strategize and vote untruthfully. In particular, we show that a single-member district plurality system under the policy of banning outlier maps with only two parties is not strategyproof.

Under the policy of banning outliers, we show examples of how a party holding the majority of seats can vote and draw districts strategically to increase the number of seats they win. For banning outliers, we provide a heuristic for a party controlling the districting process to identify pure strategies that lead to winning more seats and test this heuristic empirically. Finally, we use grid graph models to explore questions from the US Supreme Court case Rucho v. Common Cause [4] relating to outlier maps.

While not the primary pursuit of this work, our observations may also be relevant to the areas of election security and machine learning. In the case of election security, voter fraud, and tampering with election results, we essentially reveal a scenario wherein a political party could gain in the long run by generating votes for the opposing party. Such a scenario may be difficult to detect if one is assuming that a cheating party would only assign votes to itself. In relation to machine learning, we can view recent computational efforts to detect gerrymandered maps as classifiers. Thus, our work explores how these classifiers interact with a broader system.

4 MODELING THE GAME OF VOTING AND REDISTRICTING

We introduce a simple game to model the cycle of drawing districts, voting, redrawing districts, and voting again. The drawing of districts is constrained by a regulation which uses past voting data to decide if a district map is legal or not. The goal of the model is to clearly illustrate how a regulation can incentivize voters to vote strategically rather than truthfully and better understand how regulations might be circumvented.

In this game, there are two political parties, *red* and *blue*, which we denote R and B, respectively. Voters are vertices in a graph $G = (V, E)$ with the number of vertices $|V| = n$ and the edge set E signifying neighboring voters. We define a map m as a partition of G into k districts. Each district must be a connected component in G with size equal to n/k (we assume for simplicity that k is odd and n/k is an odd integer). In this way, we enforce the rules that districts must be contiguous and perfectly balanced in terms of population. In Sections 5 and 7, we further restrict G to be a grid graph. Thus, in those sections, each district will be an equal-sized, non-overlapping polyomino as seen in the maps of 3×3 grids in Figure 1.

Each voter v has a true preference for one of the two parties. The set of true preferences for all voters $P \in \{R, B\}^n$ is known to both parties at the start of the game and does not change over the course of the game. Having true preferences known to the parties captures the fact that parties may know more about their voters (including how they have influenced them) in comparison to the regulation which only "sees" how people have voted. Without loss of generality, we assume the red party holds the majority of seats at the start of the game. We also use $q_v \in \{R, B\}$ to denote the most recent vote by each voter v and let $Q \in \{R, B\}^n$ be the known set of votes from the most recent election. We reiterate that set Q of votes from the previous election is known to both parties and the regulation at all times.

Finally, we have a regulation $\psi\colon m \mapsto \{legal, banned\}$ which determines whether a given map m is legal to use or banned and cannot be used. For the remainder of this work, we focus on the regulation of *banning outliers*. In this case, ψ also takes as input the previous votes Q, set of all possible maps M (or a set of maps sampled from M in Section 8), and a threshold $\tau \in (0, 1]$. If the fraction of maps awarding the same number of seats to the red party as m is less than τ, then m is banned. Otherwise, it is legal. In other words, m is a banned outlier if the number of maps awarding the same number of seats to red as m is less than $\tau|M|$.

At the start of the game, we let $Q = P$ assuming that in a prior election voters cast votes according to their true preferences. The reason for this choice is two-fold. First, it is the simplest and easiest to analyze case which still addresses our major question of whether strategic voting can be incentivized by gerrymandering regulation. Second, this models the adoption of a new regulation. We can assume that voters have been voting their true preferences in the past and the game starts at the moment when the regulation is imposed.

The game proceeds in four rounds so that we can observe the effect of a round of voting on which maps can legally be chosen. The objective of each party in this zero sum game is to win as many total seats as possible.

Round 1: The majority party (red) draws a map m subject to a gerrymandering regulation ψ. The voters' true preferences are used as the past voting data for the purpose of regulating this first round (i.e., $Q = P$).

Round 2: All voters vote simultaneously, but voters of the same party may collude. The votes are tallied, and each district's seat is awarded to whichever party won the majority of votes in that district.

Round 3: The party which won the majority of seats in Round 2 draws a new map m subject to ψ. However, Q is now the set of votes from Round 2 and this may affect ψ.

Round 4: Again, all voters vote simultaneously, but voters of the same party may collude. The votes are tallied, and each district's seat is awarded to whichever party won the majority of votes in that district.

While this game only captures two election cycles, we will show that it is able to reveal an incentive to vote strategically. Natural extensions to more cycles or more rounds of voting before redistricting will be addressed briefly in Section 10. We further note that in this short game, we can assume voters will vote their true preferences in Round 4 since these final votes are only used to determine the outcome of a single election.

4.1 Simplifying assumptions

Here, we outline and discuss a number of simplifying assumptions in our model and analysis. Further discussion of these assumptions and related future directions appears in Section 10.

To simplify the analysis of collusion, we assume each party controls all of its voters and chooses how they will vote. Furthermore, we require each voter to vote for one of the two parties. They do not have the option to abstain or vote for some third party. Therefore, in our model, a candidate winning a plurality of votes also wins a majority. To further guarantee clear majorities without ties, we use an odd number of districts with an odd number of voters in each district.

We consider the utility of a party to be a linear function of seat count. This reduces the space of strategies to explore since a party cannot gain utility by sacrificing a seat in one round in order to gain a seat in another round. However, one could also envision a more complex model with a nonlinear function that captures real world effects. For example, in many cases in the US system, their is a large added benefit to holding a 2/3rd majority. In a nonlinear model, sacrificing a seat in one round to win a seat in another could be beneficial.

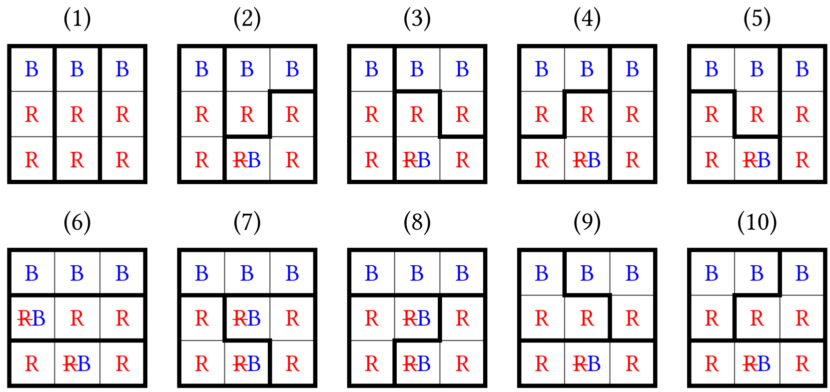

Fig. 1. All 10 possible maps of a 3×3 grid into 3 contiguous districts of equal size. Blue B's indicate voters who prefer the blue party and red R's indicate voters who prefer the red party. Squares with a crossed-out ~~R~~ followed by a B represent red party voters who can vote for the blue party in Round 2 in order to make map 1 appear to be a non-outlier map for the next round of drawing districts. Map 1 is the only map which awards 3 seats to the red party under true preferences, making it the favorite map for red, but also an outlier.

We note that in the US system, there are typically multiple elections between the rounds of redistricting that can occur following each decennial population census. Thus, our abstraction replaces a series of elections with a single voting round.

5 A SIMPLE EXAMPLE GAME ON A 3×3 GRID

To illustrate our game from Section 4 and the effects of regulation, we start by looking at the regulation of banning outliers applied to a specific set of voter preferences on a 3×3 grid with 3 districts of size 3. In this setting, we can clearly visualize an exhaustive set of maps.

A 3×3 grid admits 10 maps of 3 districts when contiguity and population balance are the only restrictions. Figure 1 shows all 10 maps along with a set of true voter preferences and strategies. The top row of voters prefer the blue party, while the bottom two rows prefer the red party. We can see plainly in Figure 1 that the red party would prefer the first map which partitions the voters into three columns. This map cracks the blue party so that red wins all 3 seats and it is the only map in which red wins 3 seats as opposed to 2.

For this simple example game, we consider the regulation of banning outliers with a threshold of τ that is strictly greater than 0.1, essentially the smallest meaningful threshold for this graph (we consider smaller, more realistic values of τ in later sections). Thus, the first map in Figure 1 will be banned with respect to the voters' true preferences P and therefore banned in Round 1. Because this is the preferred map for the red party, but it cannot be chosen given the regulation ψ and voting history $Q = P$, we call map 1 the target map for red. In other words, this is the map that red would like to draw in Round 3 in order to win an extra seat in the game.

Now, suppose one or more of the red voters were to vote for blue in Round 2 without giving up a seat in that round. This could make it appear that the first map will award 1 seat to blue in future elections when in fact the red voters could then vote truthfully to award all seats to red. We may then ask if blue can respond with its own strategic voting, but in this case red has a pure strategy for any choice of starting map that blue cannot respond to as stated in Observation 5.1.

Observation 5.1. *For any choice of starting map the red party has a pure strategy (illustrated in Figure 1) which leads to winning 5 seats total over the course of the two voting rounds in the game defined in Section 4.*

Note that in several maps (2, 3, 4, 5, 9, and 10) in Figure 1, red can flip a single voter to blue in order to make it appear that the middle column district of map 1 will be awarded to blue. In addition, this will not cause red to lose a seat in Round 2 voting compared to voting true preferences. To potentially counter this effect, blue could flip its voter in the middle column to red. However, this would cause blue to lose that voter's district in Round 2, immediately giving 3 seats to red. Conversely, in other maps (6, 7, and 8) blue can afford to flip one voter to red without losing a district and this forces red to flip two voters from two separate districts to guarantee that map 1 will not be an outlier when drawing a map in Round 3. Thus, given any legal starting map (all maps besides map 1), red can vote strategically to win 5 seats overall in the game, whereas truthful voting would only yield 4 seats overall. In Section 6, we explore this phenomenon more by establishing a set of conditions that permit a majority party to strategize in more general settings.

6 CONDITIONS FOR STRATEGIZING AGAINST BANNING OUTLIERS

In the previous section, we saw the existence of a set of preferences which incentivized strategic voting to gain seats when outlier maps are banned. Now, we define a set of conditions under which a pure strategy exists for the majority party to gain at least one seat by carefully drawing maps and influencing its supporters to vote strategically. Taken together these conditions are sufficient to incentivize strategic voting, but may not be necessary, especially when mixed strategies are considered.

In order to strategize, the majority party must be able to find two maps, a *starting map* which it can legally choose in Round 1 and a *target map* which awards the party more seats than the starting map, but is an outlier under true preferences. Additionally, given a starting map and target map, we identify a set of majority party voters called the *shills* who will vote for the opposing party in the first vote of Round 2, but then revert to supporting the majority party in the second vote of Round 4. We also require a set of minority party voters called the *accomplices* whose truthfull votes in Round 2, combined with the shills' fake votes, can be used to construct a district which appears to favor the minority party while actually favoring the majority party under true preferences. We note that the accomplices from the minority party are, of course, unwilling accomplices whose votes will be used against their own party no matter how they vote. The majority party will put them in the following position. If the accomplices vote truthfully in Round 2, their party will ultimately lose a seat in Round 4, while if they vote untruthfully in Round 2, their party will immediately lose a seat in Round 2. Given these definitions, the conditions below are sufficient for the majority party to strategize.

(1) There exists a starting map which is not an outlier with respect to true preferences.
(2) There exists a target map which awards more seats to the majority party than the starting map, but is an outlier with respect to true preferences.
(3) There does not exist a non-outlier map which, if chosen in all rounds, awards as many or more seats than the combination of choosing the starting map followed by the target map under true preferences.
(4) In the starting map, the shills are in a different district from the accomplices.
(5) In the starting map, the majority party does not lose a seat if the shills all vote for the minority party.
(6) In the starting map, the minority party will lose a seat if any one of the accomplices votes for the majority party.

(7) In the target map, the shills are in the same district as the accomplices.
(8) In the target map, the district containing the shills and accomplices will go to the majority party under true preferences, but appears to go to the minority party if the shills and accomplices vote for the minority party in the first round.
(9) The target map is not an outlier in Round 3 if the accomplices and all majority party voters besides the shills vote truthfully in Round 2.

Using these conditions, we can identify opportunities to strategize as described in the next section.

7 HEURISTIC FOR STRATEGIZING AGAINST BANNING OUTLIERS AND EXPERIMENTS

Based on the conditions from Section 6, we devise a heuristic for finding pure strategies for the majority party subject to the regulation of banning outliers. To test this heuristic, we use a grid graph model with a 5×5 grid and 5 districts of size 5.

7.1 The heuristic

We implemented a simple heuristic which takes a set of preferences as input and searches for starting maps, target maps, accomplices, and shills satisfying the conditions from Section 6. Considering all potential (starting map, target map) pairs takes at most $O(\tau|M|^2)$ time since there are $O(|M|)$ potential starting maps and only $O(\tau|M|)$ possible target maps given that target maps are outliers by definition. Searching a pair of maps to find shills and accomplices takes $O(n)$ time with the appropriate preprocessing of each map. Thus, the entire process requires $O(n\tau|M|^2)$ time per preference set. However, the practical running time is much faster since we can stop as soon as we find a satisfying pair of maps.

7.2 Experiment design

We test our heuristic in the scenario of a 60/40 split in the true preferences of voters. Our "state" is a 5×5 grid with 5 districts of size 5. For this grid, we consider all possible sets of true preferences where the majority party has 15 of 25 total voters (~3.2 million preference sets). A benefit of using this simple model is that we can consider all 4,006 contiguous and balanced district maps in order to identify outliers for a given set of preferences. This divorces the analysis of outlier regulation policies from the questions of how to sample maps and detect outliers that are the focus of [13, 15, 21].

Outlier threshold. We choose a threshold τ of 2% for banning outliers. In other words, given a true preference set P or voting history Q, any particular seat count awarded to the majority in less than 2% of the 4,006 possible maps is considered an outlier and any maps awarding that many seats would be banned. Aside from being a "natural looking" threshold, this choice is well-suited to our scenario. For a large number of true preference sets, the 2% threshold bans maps that award all 5 seats to the majority party (Figure 2). On the other hand, it allows maps that award 4 seats to the majority party for nearly all true preference sets (Figure 3). Awarding 4 seats to the majority may be seen as a reasonable deviation from proportionality due to the geography of the voter population. Allowing this amount of deviation is one way for regulators to be clear that they are not enforcing proportionality.

During the two district drawing rounds (1 and 3), the majority party is allowed to choose any of the non-outlier maps. During the voting rounds (2 and 4), both parties' voters may vote strategically, but our heuristic draws maps based on pure strategies where either only the majority party has an incentive to vote strategically or no party votes strategically. For each set of true preferences, we

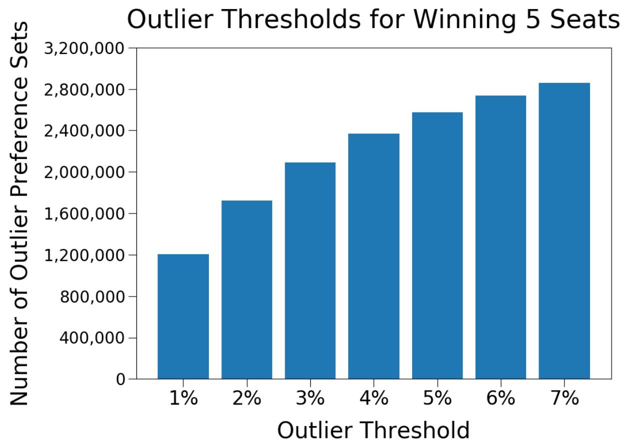

Fig. 2. The number of preference sets out of 3,268,760 total where a map awarding 5 seats to the majority party is an outlier according to percentage thresholds (1% - 7%).

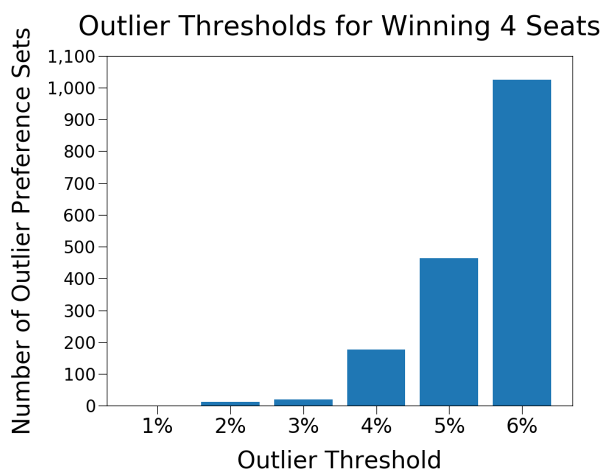

Fig. 3. The number of preference sets out of 3,268,760 total where a map awarding 4 seats to the majority party is an outlier according to percentage thresholds (1% - 6%).

test whether our heuristic can find a strategy to gain more total seats than could be gained through voting truthfully.

7.3 Experimental results

Figure 4 shows the results of our experiments. We see that for roughly half of the 3,268,760 preference sets, maps awarding all 5 seats to the majority party are not outliers. In those cases, no strategizing is needed. The majority party can simply pick a map which awards it 5 seats and use that map for the entire game. However, on most of the remaining preference sets, the majority party is limited to choosing maps in Round 1 that award fewer than 5 seats under true preferences, but our heuristic is able to find a pure strategy which leads to winning an additional seat in Round

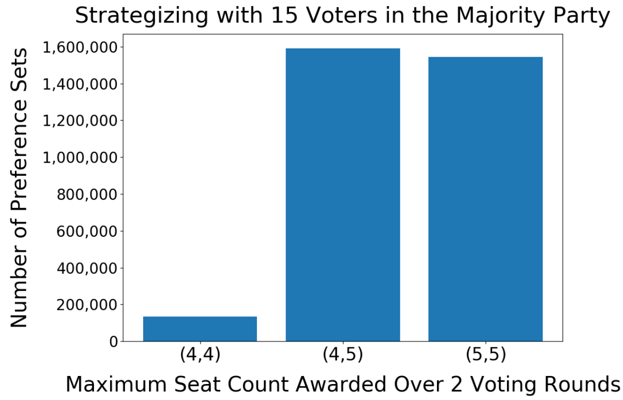

Fig. 4. Illustrating the ability of our heuristic to find pure strategies on all 3,268,760 preference sets with 15 majority party voters. (4,4) indicates preference sets where non-outlier maps award 4 seats to the majority party under true preferences and our heuristic is unable to strategize for more seats in the second vote. (4,5) indicates preference sets where non-outlier maps award 4 seats to the majority party under true preferences and our heuristic finds a pure strategy to win 5 seats in the second vote. (5,5) indicates preference sets where non-outlier maps award 5 seats to the majority party under true preferences and thus, there is no benefit to strategizing. This figure omits 12 preference sets where the best non-outlier map awards only 3 seats under true preferences and we can strategize to win 4 seats in the second voting round.

4. For fewer then 200,000 preference sets, the best non-outlier map awarded 4 seats and we were not able to find a strategy. We do not know if any pure or mixed strategies exist for these preference sets.

8 APPLICATION OF OUR HEURISTIC AND MODEL REAL VOTING DATA

8.1 Background

North Carolina is often highlighted as one of the most gerrymandered states in the nation. In the 2012 North Carolina congressional election, over half of the total votes went to Democratic candidates, yet only four of the thirteen congressional representatives were Democrats [21]. In fact, court cases have struck down North Carolina's 2012 and 2016 congressional maps for partisan gerrymandering and a case involving North Carolina's map recently went all the way to the Supreme Court in [4]. In one effort to address this issue, the "Beyond Gerrymandering" project sponsored by the Duke Center for Political Leadership, Innovation, and Service brought together an independent commission of 10 bipartisan retired judges to redraw North Carolina's congressional map without the use of past political data or election results with the intention to generate a more fair district map. The hypothetical map that was produced as a result of this summit will be referred to as the *judges' map*.

8.2 Our results in brief

In order to lend credence to our model and heuristic, we apply simplified versions to real North Carolina voter data. In the more restricted model, we find one possible strategy for the Republican party to circumvent an outlier ban by using the judges' map as a starting map and the 2016 North Carolina congressional map as a target map. The changes outlined in the next section make the problem of finding a strategy much more tractable for the larger, messier real world problem.

However, we also describe how this more restricted model can still inform our understanding of real scenarios.

8.3 Modified model

We retain concepts of our original game such as banning outliers, a two-party system, and four-round drawing-voting-drawing-voting cycle, as outlined in Section 4. We treat each of the 2,692 voter tabulation districts (VTDs) in North Carolina [21] as points on a planar graph and districts as a partition of this graph into connect subgraphs. Each VTD is encoded with Democratic and Republican votes. Crucially, we modify our original game such that the strategizing is one-sided. While the majority party can strategize, the minority party must vote truthfully. Thus, this model prevents the minority party from counter-strategizing. In our previous model, each basic unit was an equivalent voter. In this case, each basic unit (a VTD) does not carry the same weight and the number of potential reactionary strategies from the minority party increases dramatically.

While this restricted game captures fewer real scenarios, we highlight two interesting observations. First, our model is motivated by the idea of a "surprise attack" in which one party decides to manipulate votes to gain an advantage and the other party either does not expect the attack or knows, but cannot mobilize a counter strategy. In practice, this surprise attack could even be a secret attack via hacking electronic voting systems. Second, although the majority party strategy we find may not be an optimal pure strategy in the original game, it can be used to illustrate the existence of some mixed strategy in the original game. Thus, it still shows an incentive for strategic voting on real data in a less restricted game.

8.4 Conditions for strategizing on the modified model

We modify our conditions from Section 6 to account for the inability of the minority party to react. The *shills* and *accomplices* are now groups of voters within VTDs. However, they still perform the same role as in our original model. Taken together, these conditions are sufficient, but may not be necessary, for a party to strategize in the modified model of this section. Note that we have maintained the numbering from Section 6 in the list below to facilitate easier comparison between the two lists.

(1) There exists a starting map which is not an outlier with respect to true preferences.
(2) There exists a target map which awards more seats to the majority party than the starting map, but is an outlier with respect to true preferences.
(3) There does not exist a non-outlier map which, if chosen in all rounds, awards as many or more seats than the combination of choosing the starting map followed by the target map under true preferences.
(5) In the starting map, the majority party does not lose a seat if the shills all vote for the minority party.
(7) In the target map, the shills are in the same district as the accomplices.
(8) In the target map, the district containing the shills and accomplices will go to the majority party under true preferences, but appears to go to the minority party if the shills and accomplices vote for the minority party in the first round.
(9) The target map is not an outlier in Round 3 if all voters besides the shills vote truthfully in Round 2.

In comparison to the conditions from Section 6, conditions (4) and (6) are no longer needed and (9) is slightly relaxed. This is because we are assuming in the modified model that the minority party accomplices vote truthfully. Thus, we do not need to meet conditions that discourage the accomplices from strategic voting or account for that possibility.

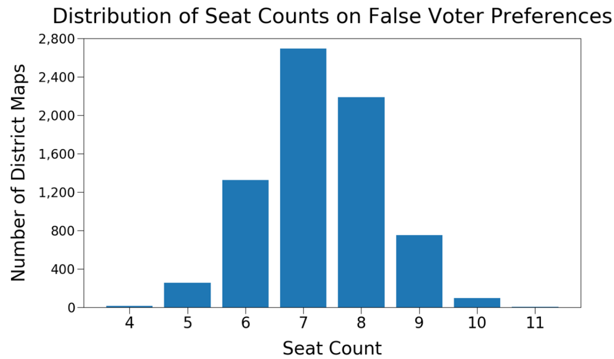

Fig. 5. Plot of Republican seat distributions of the false voter preferences after all of the Republican votes in the judges' map district 9 VTDs have been flipped to Democrat votes. We determined the seat counts of the false voter preferences on 7319 total maps generated by Monte Carlo Markov Chain that we retrieved from [21]. We note that under our false voter preferences, the 2016 map awards 9 Republican seats which is not an outlier for the false preferences. However, the true voter preferences award 10 seats on the 2016 map, which is an outlier.

8.5 Simplified heuristic

Because there is now an enormous number of possible maps and a large set of voters, we cannot search exhaustively like our procedures from our grid model. In addition to drastically narrowing our search space to specific starting (judges' map) and target (2016 map) maps, we implement a more targeted heuristic for identifying shills. For the most competitive districts (least margin between voters of the two parties) in our target map not already won by the majority party, we obtain the list of VTDs in those districts. We then flip as many majority voters from those VTDs as possible without losing any districts on the starting map, such that those VTDs entirely consist of minority votes in Round 2. We then compare the new set of voter preferences on the starting map and target map within the sampling distribution of maps to make sure neither are outliers. We perform this procedure on each district in our target map from most competitive to least competitive until we find a successful strategy or have exhausted the set of losing districts.

8.6 Experiment design

We note that the judge's map awards 9 out of 13 seats to the Republican party and is therefore not an outlier, while the 2016 map awards 10 out of 13 seats to the Republican party and is an outlier [21]. We test our modified model over North Carolina 2016 House voting data. We set the judges' map as our starting map and the 2016 map as our target map. We set the Republican Party as the majority party and the Democratic Party as the minority party. Voter data, the judges' map, the 2016 map, and a sampling distribution of maps were retrieved from the Github provided in [21].

8.7 Results

We find that flipping all of the Republican votes to Democrat votes in VTDs found in district 9 of the judges' map causes district 2 of the 2016 map to seem to be a Democratic district without the Republicans losing any districts in the judges' map. However, according to the real voting history, district 2 of the 2016 map is a Republican district. Thus, under the false preference sets, the 2016 map seems to award only 9 seats to the Republicans. From Figure 5, we note that 9 seats is not an outlier under the flipped voter preferences. This enables the Republican party to draw the 2016 map in Round 3 and secure 10 seats in Round 4. Thus, we have shown the existence of a successful strategy in our modified model to go from the bipartisan judges' map to the gerrymandered 2016 map which was shown to be an outlier [21].

9 QUESTIONS RAISED IN RUCHO V. COMMON CAUSE

In this section, we address several questions raised during the recent US Supreme Court case Rucho v. Common Cause [4]. All of these issues were discussed in the opening oral arguments of that case. We consider them as they relate to the regulations and models explored in this paper.

9.1 Banning outliers versus the proportional allocation objective

Skeptics of the outliers metric and banning outliers regulation have argued that it is similar to or a proxy for a proportionality rule. Using basic grid models models, we illustrate where these two rules diverge in valuing and restricting maps.

First, we note that Justices Alito, Gorsuch, Kavanaugh, and Roberts have asked in some way whether a rule which includes outlier detection (among other tests) amounts to a proportionality rule in the oral arguments of [4]. This is not an unreasonable concern. The concept of banning outliers contains important features which are open to manipulation such as

(1) Which metric do you use to compare maps (e.g., seat count[21] or variance in the proportions of one party's voters among the districts [13])?
(2) How do you set the threshold for what kind of map is an outlier?
(3) How do you choose among multiple non-outlier maps which may nevertheless assign more seats to one party or another?
(4) How do you define the space of legal maps that you are sampling from? In other words, how closely can your mathematical constraints on what constitutes a legal map approximate the legal definitions in state constitutions?

We show here that despite these concerns, the concept of banning outliers differs from enforcing or favoring proportionality in important ways. On our 5×5 grid model, we observe what percentage of maps award proportional representation for given sets of preferences. Figure 6 captures all possible preference sets on a 5×5 grid with 20 majority voters and 5 minority voters. Figure 7 captures all possible preference sets on a 5×5 grid with 15 majority voters and 10 minority voters. It is clear that the seats awarded by the most proportional map can differ greatly from the most likely randomly chosen maps that respect the voter geography. In the case of 20 majority voters, we found that a proportional map was the most likely random map for only 6.9% of the preference sets and for the case 15 majority voters a proportional map was most likely in only 41.5% of preference sets. Comparing Figures 6 and 7 also suggests that proportional maps are generally more rare when the minority party represents a smaller proportion of the population, while proportional maps are more common for preference sets with more even splits between the parties.

Perhaps the most convincing argument for the difference between outlier detection and proportionality comes in Figure 6. We can see that for many preference sets, proportional maps exist, but are rare and could be explicitly forbidden under a regulation banning outliers. In other words, for certain arrangements of voters, proportional maps could actually be outlawed by banning outliers.

9.2 Individual harm

Central to the argument that partisan gerrymandering is unconstitutional is the notion of individual harm. We may ask if an individual's vote was diluted by a map based on their political preference. A line of inquiry in [4] focused on the following situation. Suppose we can show that for a given

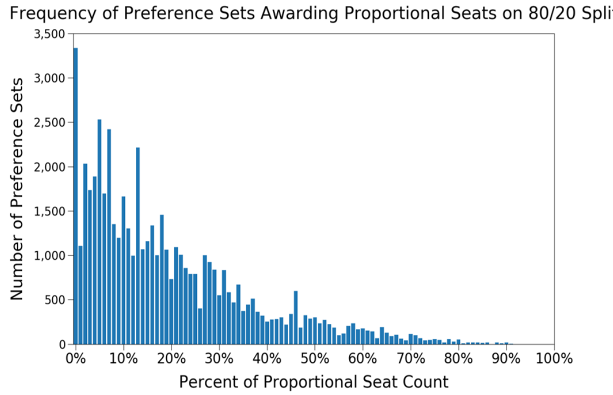

Fig. 6. Charting all possible preference sets on a 5×5 grid with 20 majority voters and 5 minority voters. We show the numbers of preference sets in which a given percentage of maps achieves proportional representation. The large bar on the left at 0% indicates preference sets where no map awarded proportional representation. The second bar from the left indicates the number of preference sets in which only 1% of maps awarded proportional representation (i.e. preference sets in which proportion maps exist, but are very rare).

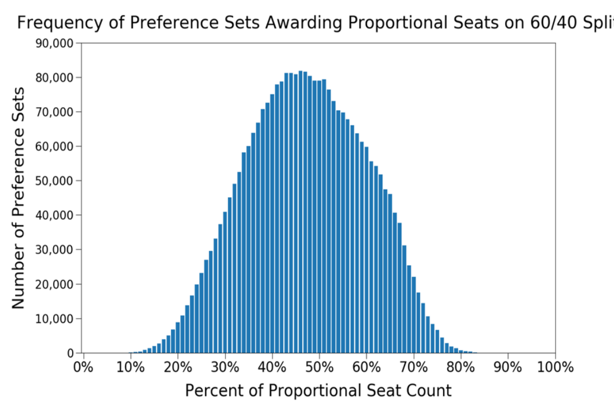

Fig. 7. Charting all possible preference sets on a 5×5 grid with 15 majority voters and 10 minority voters. We show the numbers of preference sets in which a given percentage of maps achieves proportional representation.

individual, most maps place them in a district where their chosen party wins, but the proposed map is one of a few maps which does not. Can this test be used to show individual harm?

Here, we give evidence of a problem with this test. Suppose there are many such individuals and no single map is fair to all of them. In this case, no deterministically drawn map could be considered fair. However, it could still be argued that a randomly drawn map is fair by this standard (this could trivially be achieved by picking a map uniformly at random from the set of sampled maps).

As evidence of this issue, we consider a 5×5 grid where the majority party has 13 voters in a checkerboard pattern. In this simple example, each voter can be placed in a district where their party wins in over 64% of maps. However, there is no single map which provides this opportunity to all 13 of those voters. Under a philosophy that places value on the most likely maps for a given geography, this raises the question of whether a deterministically drawn district can truly be called fair.

10 CONCLUSION, RECOMMENDATIONS, AND FUTURE DIRECTIONS

We have shown first and foremost that careful scrutiny should be given to any measurement which uses past voter behavior to evaluate and affect the choice of electoral district maps. The model we presented primarily serves to illustrate the basic phenomenon of strategic voting in the presence of gerrymandering regulation. To better understand this issue, more complex models should be considered.

One of the most unrealistic assumptions in this work is that a political party is able to totally control all of its voters. An obvious concern with the current model is whether it is feasible to organize a large enough group of voters from one party to cast votes for the opposing party in some, but not all elections on a single ballot. Perhaps the only real analog close to this would be the hacking scenario in which people's votes are changed illegally and without their knowledge. Thus, if we wish to consider the possibility of actual strategic voting in practice, we must modify the model in some way. A natural extension would be to consider the realistic influence that political parties do have over a voter's decision between voting for their preferred party or abstaining. It would be interesting to model and explore whether parties can distort gerrymandering measurement by choosing where to spend their limited budgets on "get out the vote" efforts using the power of modern tools such as targeted advertising.

Other useful directions would be to add noise to the model or change the number rounds. The total number of rounds as well as the number of voting rounds between redistricting could be arbitrary. Voter preferences might have some probability to change between rounds. While our four-round model is sufficient to show a basic incentive to vote strategically, richer behavior might evolve in a longer process. For example, majority party voters in our model vote truthfully in Round 4 to capitalize on their choice of a favorable map in Round 3. However, minority party voters might then wish to voter strategically if Round 4 were not the final round.

Regarding the problem of finding strategies to circumvent regulation, the heuristics presented here were fairly simple and lightweight. It is likely that more sophisticated algorithms and more computing power could be employed to greater effect especially on real data. As with any problem in algorithmic game theory there is the two-fold challenge of figuring out what the optimization problem is as well as how to solve it. We essentially identified one type of strategy and how to execute it. However, there are likely other approaches, especially when considering different models.

Finally, it is worth considering how gerrymandering metrics that use past voting data can be less susceptible to strategizing. We note that a metric which uses voting data from multiple elections would likely be harder to trick. In the US system in particular, data from senate and gubernatorial races could be especially useful since the voters within a state are not partitioned into districts for these elections and that could have a confounding effect on incentives.

ACKNOWLEDGMENTS

The authors wish to thank the anonymous reviewers for useful feedback as well as Jonathan Mattingly and Alan Frieze for fruitful discussions. Research supported in part by NSF Award CCF-1749864 and research awards from Adobe, Inc., Amazon, and Google.

REFERENCES

[1] Arizona State Legislature v. Arizona Independent Redistricting Commission, No. 13-1314, 576 U.S. ___ (2015).
[2] Gill v. Whitford, No. 16-1161, 585 U.S. ___ (2018).
[3] League of Women Voters of Pennsylvania v. Commonwealth of Pennsylvania, No. 159 MM (2018).
[4] Rucho v. Common Cause, No. 18-422, 588 U.S. ___ (2019).
[5] Shaw v. Reno, No. 92-357, 509 U.S. 630 (1993).

[6] Thornburg v. Gingles, No. 83-1968, 478 U.S. 30 (1986).
[7] Vieth v. Jubelirer, No. 02-1580, 541 U.S. 267 (2004).
[8] Mira Bernstein and Moon Duchin. 2017. A formula goes to court: Partisan gerrymandering and the efficiency gap. *Notices of the AMS* 64, 9 (2017), 1020–1024.
[9] Allan Borodin, Omer Lev, Nisarg Shah, and Tyrone Strangway. 2018. Big city vs. the great outdoors: voter distribution and how it affects gerrymandering. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*. AAAI Press, 98–104.
[10] United States Census Bureau. 2017. My Congressional District. https://www.census.gov/mycd/. [Online; accessed 15-August-2019].
[11] Aaron Bycoffe, Ella Koeze, David Wasserman, and Julia Wolfe. 2018. The Atlas Of Redistricting. https://projects.fivethirtyeight.com/redistricting-maps/. [Online; published 25-January-2018; accessed 15-August-2019].
[12] Jowei Chen, Jonathan Rodden, et al. 2013. Unintentional gerrymandering: Political geography and electoral bias in legislatures. *Quarterly Journal of Political Science* 8, 3 (2013), 239–269.
[13] Maria Chikina, Alan Frieze, and Wesley Pegden. 2017. Assessing significance in a Markov chain without mixing. *Proceedings of the National Academy of Sciences* 114, 11 (2017), 2860–2864. https://doi.org/10.1073/pnas.1617540114 arXiv:https://www.pnas.org/content/114/11/2860.full.pdf
[14] Wendy K Tam Cho. 2019. Technology-Enabled Coin Flips for Judging Partisan Gerrymandering. *Southern California law review* 93 (2019).
[15] Wendy K Tam Cho and Yan Y Liu. 2018. Sampling from complicated and unknown distributions: Monte Carlo and Markov chain Monte Carlo methods for redistricting. *Physica A: Statistical Mechanics and its Applications* 506 (2018), 170–178.
[16] Vincent Cohen-Addad, Philip N. Klein, and Neal E. Young. 2018. Balanced Centroidal Power Diagrams for Redistricting. In *Proceedings of the 26th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (SIGSPATIAL '18)*. ACM, New York, NY, USA, 389–396. https://doi.org/10.1145/3274895.3274979
[17] Moon Duchin. 2018. Geometry v. Gerrymandering. *Scientific American* 319 (10 2018), 48–53. https://doi.org/10.1038/scientificamerican1118-48
[18] Moon Duchin, Taissa Gladkova, Eugene Henninger-Voss, Ben Klingensmith, Heather Newman, and Hannah Wheelen. 2018. Locating the representational baseline: Republicans in Massachusetts. *arXiv preprint arXiv:1810.09051* (2018).
[19] Allan Gibbard et al. 1973. Manipulation of voting schemes: a general result. *Econometrica* 41, 4 (1973), 587–601.
[20] Corbett A Grainger. 2010. Redistricting and Polarization: Who Draws the Lines in California? *The Journal of Law and Economics* 53, 3 (2010), 545–567.
[21] Gregory Herschlag, Han Sung Kang, Justin Luo, Christy Vaughn Graves, Sachet Bangia, Robert Ravier, and Jonathan C Mattingly. 2018. Quantifying Gerrymandering in North Carolina. *arXiv preprint arXiv:1801.03783* (2018).
[22] Lily Hu, Nicole Immorlica, and Jennifer Wortman Vaughan. 2019. The disparate effects of strategic manipulation. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. ACM, 259–268.
[23] Yan Y. Liu, Wendy K. Tam Cho, and Shaowen Wang. 2016. PEAR: a massively parallel evolutionary computation approach for political redistricting optimization and analysis. *Swarm and Evolutionary Computation* 30 (2016), 78 – 92. https://doi.org/10.1016/j.swevo.2016.04.004
[24] Wesley Pegden, Ariel D Procaccia, and Dingli Yu. 2017. A partisan districting protocol with provably nonpartisan outcomes. *arXiv preprint arXiv:1710.08781* (2017).
[25] Benjamin Plener Cover. 2018. Quantifying partisan gerrymandering: An evaluation of the efficiency gap proposal. *Stan. L. Rev.* 70 (2018), 1131.
[26] Ivan Ryan and Warren D. Smith. [n.d.]. Splitline districtings of all 50 states + DC + PR. https://rangevoting.org/SplitLR.html. [Online; accessed 15-August-2019].
[27] Mark Allen Satterthwaite. 1975. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of economic theory* 10, 2 (1975), 187–217.
[28] Nicholas Stephanopoulos and Eric McGhee. 2015. Partisan Gerrymandering and the Efficiency Gap. *University of Chicago Law Review* 82 (2015), 831.